Jonsson, M., Section Introduction – Overview of Real-Time Communication Research Internationally, in ARTES – A network for Real-Time research and graduate Education in Sweden 1997-2006, Editor: Hans Hansson, pp. 333-351, Mar. 2006, ISBN 91-506-1859-8.

Chapter 14

Introduction

By **Magnus Jonsson**
School of Information Science
Computer and Electrical Engineering
Halmstad University
Email: Magnus.Jonsson@ide.hh.se

14.1 Introduction

More and more real-time systems are complex and distributed systems consisting of many sub-systems that must cooperate. In other words, those sub-systems must communicate with each other and to do this, still fulfilling the overall system specification, real-time communication must be supported. Moreover, real-time communication is also becoming widely needed in networks like the Internet, spanning over large distances. In for example Internet, the support for different traffic classes is often described as supporting different QoS (Quality of Service) levels and does normally include some kind of specification to meet real-time demands.

Real-time communication often relies on some kind of scheduling like EDF (Earliest Deadline First), but there are some important differences compared to standard single-processor scheduling. First, the transmission of a packet is, in most cases, non-preemptive. The ongoing transmission, and possibly also some of the already queued messages, can then not be interrupted when a more important message (e.g., with a shorter relative deadline) arrives or is generated. Moreover, a network, instead of a single link, makes the situation much more complex and one must consider things like medium access method, topology, multiple users on multiple nodes, non-deterministic access delay etc.

Regarding approaches to handle real-time communication, one is to schedule the traffic according to, e.g., relative deadline (also called laxity or delay bound) or priority. In other words, dynamic (on-line) communication scheduling is used. Without any further analysis this only gives a best-effort service, i.e., no real-time guarantees are given. Secondly, for networks with a deterministic scheduling/behavior, a schedability analysis can be made through static (off-line) scheduling to ensure a timely behavior. A third approach is to have a semi-static scheduling where admission control with a schedability analysis is used when a new logical channel for real-time traffic is requested. The new channel is accepted only if the network can guarantee the specified QoS, not violating the granted QoS for existing connections. As an example, any of these three approaches can be used in a network with priority scheduling:

- Approach 1: Nothing more than prioritizing important traffic (no guarantees).
- Approach 2: Analyze the network behavior at system design, given the behavior of the communicating processors/processes.
- Approach 3: Admission control (like Approach 2 but analysis during run-time, each time the traffic characteristics changes).

Below, we will give an overview of real-time communication characteristics and representative methods to give real-time communication support in different kinds of networks. After that, the appended papers are shortly introduced.

14.2 Application and traffic characteristics

Real-time systems and applications are often divided into classes depending on how strict timing requirements they have. One example is the division into critical, essential, and nonessential timing requirements, while another is the division into hard and soft real-time requirements. It is, however, not always the case that the real-time communication for an application can inherit the same classification as the application in whole. As an example, IP telephony can be regarded as a soft real-time application since it still might be possible to understand what is said even though some data is lost or not delivered in time. Despite this, a single message is useless if not delivered in time and might therefore be considered as having a hard deadline. A compromise is to state a probability that a message is delivered in time.

Since the classification into soft and hard real-time communication is somewhat unclear, real-time communication support is normally described more in terms of the service offered. For example, the communication support for a certain traffic class with real-time requirements might have both a guaranteed (minimum) throughput and a bounded delay. Some common such parameters to describe traffic requirements and QoS levels together with more general traffic classifications are:

- The traffic can be of either aperiodic (asynchronous) or periodic (synchronous) nature.
- A relative deadline (delay bound) might be stated. In case of periodic traffic it can be shorter, equal, or longer than the period which, depending on the specific network, can have significant influence on the real-time methods and analysis.
- Rate characteristics such as minimum inter-arrival rate and maximum message-length are parameters describing the desired throughput in a more precise manner.
- The delay jitter is important to be bounded for some applications and is defined as the difference between the minimum and the maximum end-to-end delay.
- The reliability in terms of, e.g., loss rate is used when describing services without hundred percent guaranteed timely delivery.

Sometimes, the traffic characteristics or network requirements are in an interval like [required, desired] or by both minimum and average values. A good example of this is when having multimedia traffic with dynamically changing throughput requirements due to variable bit-rate compression.

The linear bounded arrival process model is sometimes used to characterize and bound periodic but bursty traffic, i.e., traffic that sometimes comes in large bunches of packets instead of always being regularly spaced in time. The intended amount of data from the source is specified as:

S = maximum packet size
R = maximum packet rate
W = maximum work ahead

where W allows for short violations of R (bursts). In any interval of duration t, at most $W + t \cdot R$ packets may arrive (be generated).

Different applications can have fundamentally different traffic characteristics and real-time demands. For example, automotive applications like steer-by-wire have very strict demands where safety-oriented designs with high reliability must be considered. Methods like repetitive transmissions and/or fault-tolerant network architectures with hardware redundancy must be used to ensure correct and timely delivery. In automation industry, the traffic is often of periodic nature with a master-slave communication pattern where one or several master nodes periodically retrieve sensor values from slave nodes.

Data and telecommunication equipment like large distributed routers in the Internet are characterized by having requirements on short delays to not add too much to the end-to-end delay and not contribute too much to large queuing requirements. Regarding a specific Internet application, IP telephony should, although it is not always technically possible, have a maximum end-to-end delay of 150 ms to avoid that a half-duplex session is entered where one talks and the other listens [Hassan et al. 2000]. Interactive video applications have higher throughput requirements but the same delay requirements as IP telephony [Wolf and R. Steinmetz 1997]. For non-interactive multimedia applications, longer delays might be acceptable but the jitter can be a problem. Playback buffers can be used to cope with the jitter but rather large buffers might be needed if the network not includes mechanisms to reduce the jitter. Other things worth mentioning about multimedia communication are that: (*i*) related streams (e.g., sound and video) must be tightly synchronized (less than 80 ms skew), (*ii*) there is normally no time for retransmissions but 100 % reliability are not often required, and (*iii*) support for multicast is often desired. When evaluating networks regarding their capability of transporting multimedia traffic, traced traffic (e.g., from an MPEG compressed movie) is often used.

The area of sensor networks is currently getting a lot of attention in the research community and real-time communication is one subarea [Stankovic et al. 2003]. New protocols and methods to handle real-time communication in dynamic wireless low-power constrained networks must be developed. Other application areas where real-time communication is important include avionics, signal processing applications, process control, and remote surgery.

14.3 Real-time communication services

Examples of real-time user services that can be offered, to the application or the programmer, by the network are: RTVC (Real Time Virtual Channel) [Ferrari and Verma 1990], guarantee-seeking messages [Arvind et al. 1991], and different kinds of best-effort and non-real-time services. If the service offered is deterministic, it is normally offering both a guaranteed minimum throughput and a bounded end-to-end delay. If the service offered is probabilistic, it can still be the case that a guarantee is offered but only as a guarantee to meet the specified QoS level at a certain probability. Some networks have inherent heterogeneous real-time support, i.e., having support for several traffic classes with rather different characteristics and real-time demands [Bergenhem and Jonsson 2003]. Using active networking, the QoS support can even be adapted during run-time by having service modules dynamically loaded into the network equipment [Metzler et al. 1999]. A related topic is about approaches to support real-time traffic over heterogeneous networks, i.e., networks composed of rather different sub-networks. One example is to support connection-oriented real-time communication in an FDDI-ATM-FDDI heterogeneous network [Chen et al. 1997]. Another example is the proposal of a general flexible TDMA (Time Division Multiple Access) approach that can be used to support real-time communication over a heterogeneous network with both WLAN (Wireless Local Area Network) and fieldbus technology (CAN; Controller Area Network) [Mock and Nett 1999], while a third example investigates the use of wireless communication to connect several PROFIBUS segments by forwarding and insertion of extra idle-time between messages [Alves et al. 2002]. It is not only pure data delivery services for communication between two nodes that can be enhanced with real-time support. One example is the support of real-time services for special communication patterns like many-to-many [Fan et al. 2004].

14.4 Local area networks

Traditionally, a LAN (Local Area Network) implies a shared-medium network where a MAC (Medium Access Control) protocol is used to control which node that should have the possibility to send in each instance. With shared-medium networks, we mainly refer to bus networks and ring networks. A MAC protocol with a deterministic behavior is needed if guaranteed real-time services shall be supported. For example, CSMA/CD (Carrier Sense Multiple Access, Collision Detect) used in Ethernet is not deterministic. Several of the deterministic MAC protocols can be used together with a schedability analysis to calculate the on-line performance of predefined communication patterns off-line.

Several MAC protocols and extensions to basic MAC protocols to support real-time communication over bus networks have been proposed. TDMA (Time Division Multiple Access) is one example where the access to the medium is divided into time-slots. In static TDMA, each node normally has the access to the bus in one time-slot per cycle but many variants of TDMA exist where, e.g., time-slots are dynamically booked. As long as not only relying on self-synchronization on transmitted frames (packets), clock-synchronization is very important in bus networks. Mars, proposed by Kopetz et al., is one example of a distributed real-time system where a TDMA based network is included [Kopetz et al. 1989].

Window protocols can also be used to obtain real-time functionality over bus networks, which the following example shows. Assume that the access to the medium is divided into windows (time intervals) and that a message has a priority/node/process specific value. Transmission is then only allowed if the value is in the window (and the medium is free). On collision, the window is split into several windows to reduce the risk of having two nodes in the same window. When having a deadline controlled window protocol, the basic idea is to have a value set according to the relative deadline and to let windows for earlier deadlines come earlier [Zhao et al. 1990]. Using VTCSMA (Virtual Time CSMA) [Molle and Kleinrock 1985] instead, all nodes are clock synchronized and a (virtual) time to start transmission is calculated depending on, e.g., deadline [Zhao and Ramamritham 1987]. Collisions can still appear and are treated with random delays etc. In other words, no deadline guarantees can be given and no schedability analysis can be made.

The IEEE 802.3 network (daily called Ethernet) has no inherent support for real-time communication but several kinds of extensions and modifications have been proposed. In RETHER [Venkatramani and Chiueh 1994], a software implemented token-based protocol is added on top of the normal MAC protocol. To reduce overhead, the token-based protocol is activated only when there is real-time traffic to be sent. Other researchers have proposed to control the amount of outgoing traffic from a node in an adaptive way to increase the overall possibilities in the network to meet the real-time demands [Kweon et al. 2000] [Carpenzano et al. 2002].

Token Ring (IEEE 802.5) is a ring network with a priority-based arbitration supporting eight priority levels. A single token is circulated in the ring and, as long as it is free, a node can grab the token to get permission to send. If a frame passes without the token (i.e., a normal data frame), a node can try to reserve the token for the next round by setting the special reservation field to the priority value of its message. This is only allowed if the priority field is not already having the same or a higher value. This reservation mechanism gives the node with the highest-priority message the possibility to send in the next round.

FDDI (Fiber Distributed Data Interface) is an optical fiber ring network, which relies on the same principles as in Token Bus (IEEE 802.4; see [Montuschi et al. 1992] for a discussion on timing), namely the Timed-Token protocol (see [Malcolm and Zhao 1994] for real-time aspects of the Timed-Token protocol). All nodes know the value of the constant $TTRT$ (Target Token Rotation Time), while the TRT (Token Rotation Time) is measured by each node for each turn the token travels. A node is allowed to transmit asynchronous traffic if $TTRT - TRT > 0$. In addition to asynchronous traffic, a predefined part of the bandwidth can be allocated for synchronous data for which a node is guaranteed to be allowed to send when it gets the token.

Switched Ethernet is a technology that is frequently used in LANs today. An additional Ethernet standard supported by many switches does even support priority queuing with between two and eight priority levels. An extra 3-bit header field is added to the frames, while the switches have priority queues. EtheReal was an early work on switched real-time Ethernet but was only throughput oriented (i.e., no explicit treatment of real-time demands) [Varadarajan and Chiueh 1998]. Later work on switched real-time Ethernet has been much concentrated on the fact that the collision probability can be totally eliminated by only including point-to-point links and network interface cards and switches that support full-duplex transmission [Alves et al. 2000] [Hoang et al. 2002].

14.5 Fieldbus networks

Fieldbus networks have their main origin in the demand to reduce cabling cost in, e.g., vehicles [Leen et al.] and automation industry. As many of the targeted applications can be classified as real-time systems, a lot of attention has been paid to real-time communication in fieldbus networks. Since typical data consists of sensor values or other short data, the frame length is normally rather short in fieldbus networks.

CAN (Controller Area Network) is one of the fieldbuses that have got a lot of attention. Each frame transports eight bytes of data at the most. A node can listen to arbitrary “addresses” and also to many “addresses” simultaneously. Several nodes can even listen to the same address (multicast). A special detail is that the ”address” also acts as a priority level. The MAC protocol works in the following way. When sending over the bus, a dominant bit overwrites a recessive bit. The sender, starting with the most significant bit of the “address”, checks if the received bit is identical with the transmitted bit. It stops transmission if the recessive bit is overwritten, i.e., because a collision with a higher-priority message has occurred. An already begun transmission (more than just “address”) is never interrupted. This MAC protocol resolves collisions in a way so the highest priority message in the whole network is ensured to be transmitted. However, the bit-by-bit arbitration mechanism limits the maximum distance depending on the bit rate used. For a 1 Mbit/s network, the maximum farthest distance between two nodes is set to 100 m.

Analyses of the real-time performance of CAN have been published in several papers (see, e.g., [Tindell et al. 1994]). Moreover, several extensions have been proposed, e.g., to support deadline scheduling by dynamically updating the priority value according to the deadline [Zuberi and Shin 1995] or by using server-based scheduling [Nolte et al.

2003]. TTCAN (Time-Triggered CAN) is a network in which the original CAN protocol is extended with a session layer on top of the CAN link layer to support time-triggered communication [Leen and Heffernan 2002]. One node (although several can exist for fault-tolerance) is the time master in the network and triggers a new cycle regularly by sending a reference message. FTT-CAN (Flexible Time-Triggered CAN) is similar to TTCAN but allows for traffic changes during runtime [Ferreira et al. 2002].

The master-slave arrangement where the slaves only transmit upon request by the master is an approach to make it easier to analyze the network and achieve a time-deterministic behavior. This approach is used in the MIL-STD-1553, developed for avionic applications in the early 1970's [Schuh 1988]. The network has been used in, e.g., JAS, F-15, and space shuttles. One or several buses (for redundancy) can be used where only the bus controller is allowed to initiate communication. The bus controller can, e.g., periodically ask remote terminals for information.

TTP [Kopetz and Grünsteidl 1994] and FlexRay [FlexRay] are two fieldbuses developed for safety critical applications like steer-by-wire. The TTP protocol (latterly called TTP/C) relies on static TDMA schedules where different nodes can have different amounts of scheduled transmission per TDMA round. Even though TTP/C is based on static scheduling, rapid mode-changes between different predefined schedules are supported. A bigger picture of how to use a TTP/C network, combined with segments using the simpler TTP/A protocol [Kopetz et al. 2000], in the Time Triggered Architecture (TTA) is presented in [Kopetz and Bauer 2003]. Both FlexRay and TTP/C support network replication for redundancy and a choice between a bus topology and a star topology [Rushby 2001]. FlexRay, however, supports both time-triggered and event-triggered communication by, at design time, defining two separate portions of the time cycle.

A lot of other work considering real-time aspects of fieldbus networks has been published, of which a few examples follows. The Timed-Token protocol variant used in Profibus is analyzed in [Tovar and Vasques 1999], while offline scheduling of sporadic traffic in FIP networks is analyzed in [Pedro and Burns 1997].

14.6 Packet-switched networks

In a typical packet-switched network, each packet traverses a number of physical links towards the final destination. After each hop, the packet is stored (queued) in a switch (or router). The choice of queuing architecture, traffic handling etc is essential for the QoS characteristics. Service disciplines for real-time communication in packet-switched networks can be classified into the following three classes according to the service they offer [Zhang 1995]. Guaranteed deterministic services offer a hundred percent guarantee of the stated QoS level by having an admission control mechanism to verify that the specified requirements can be met. A schedability analysis is run, when a new logical channel/RTVC is requested, for both existing and the new logical connections. Each switch on the path must guarantee the specified QoS, not violating QoS for existing connections. For this to be able, each source of a logical connection must obey to the predefined characteristics of its traffic generation. Guaranteed statistical services are similar to guaranteed deterministic services but have a non-zero loss probability. In other words, there is a small probability that delay and throughput bounds are violated or that buffer-overflow occurs. Having statistical bounds instead of counting on worst-case scenarios gives higher possible utilization for QoS traffic. For the third class, predicted services, guarantees are not given at all. Instead, a measurement-based admission control is used to predict the performance. For example, the current network load can be measured.

As stated, it is important that a source obeys to its specified maximum traffic generation. Because jitter can be accumulated for each hop, and therefore make a worst-case analysis very pessimistic, it is good if the switches also implement some policing mechanism to regulate the injection rate. The Leaky Bucket policing algorithm restricts traffic to follow the Linear Bounded Arrival Process Model (see above). Leaky bucket allows for an average rate of r and burst sizes of b by having a "bucket" filled with r tokens per second (see Figure 14.1). The bucket holds up to b tokens. Packets are only sent as long as there are tokens

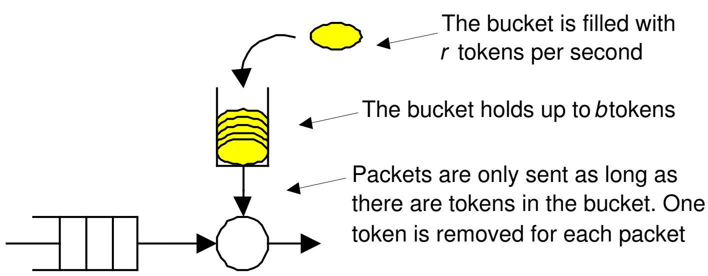

Figure 14.1: The leaky bucket algorithm.

in the bucket and one token is removed for each packet. If new tokens are generated when there are b tokens in the bucket, the tokens are discarded.

Below, we will give a few examples of guaranteed deterministic services, which can be divided into work-conserving service disciplines and non-work-conserving service disciplines. Using a work-conserving service discipline, transmission will occur as long as there are packets eligible for transmission. This maintains a good utilization. Non-work-conserving service disciplines instead, can be good with regard to jitter and hence the size of jitter eliminating buffers. WFQ (Weighted Fair Queuing) is one of the most well-known work-conserving service disciplines [Demers et al. 1989]. For each physical channel, each traffic class (or logical connection) has a FIFO (First In First Out) queue i, $1 \leq i \leq N$, with weight w_i, assuming N traffic classes (see Figure 14.2). Assuming non-empty queues when served, WFQ emulates an ideal system where packets are taken from all queues simultaneously in a bit-by-bit fashion, each queue being served at a rate equal to the fraction

$$\frac{w_i}{\sum_{j=1}^{N} w_j}$$

of the line rate. The earliest calculated finish times for a packet to be served in the ideal system decides the next queue to serve. A delay bound analysis for WFQ together with the leaky bucket algorithm used in each switch in the network is presented in [Parekh and Gallager 1994].

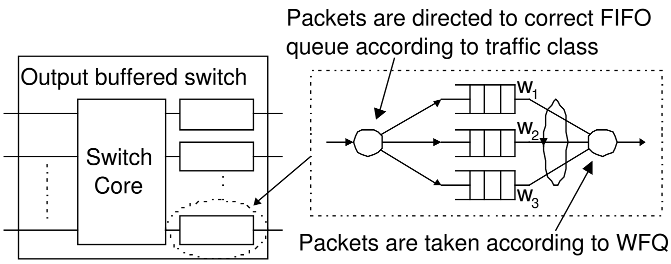

Figure 14.2: The WFQ service discipline exemplified in an output buffered switch.

Delay-EDD (Delay-Earliest-Due-Date) is another well-known work-conserving service discipline [Ferrari and Verma 1990]. EDF is used to sort packets in the switches and a deadline is set to each packet arriving at a switch. The deadline is set to the expected arrival time plus the delay bound for the current hop. By adding all delay bounds together for the path, an end-to-end delay bound is obtained. Important improvements of Delay-EDD to make it more generally applicable have been published [Kandlur et al. 1991] [Zheng and Shin 1994].

Jitter-EDD (Jitter-Earliest-Due-Date) [Verma et al. 1991] is an example of a non-work-conserving service discipline and is similar to Delay-EDD but with the following functional differences. Each switch on the route stamps each outgoing packet with the difference between its deadline and the actual transmission time, T_{diff}. The next switch on the route holds the packet for a period of T_{diff} before it is made eligible to be scheduled for further transmission.

14.7 Internet

As the Internet consists of a very wide variety of different networks coupled together, and with different operators, it is quite obvious that it is hard to give any real-time guarantees. If not just considering well-defined parts of the Internet were, for example, all routers have sophisticated QoS support and no shared-medium Ethernet networks are used, it is even impossible to give any guarantees. The fact that the IP network protocol is connectionless makes the situation even harder since routing decisions might change, and hence the possibilities for QoS, during a session. Anyhow, there are mechanisms for QoS over the Internet [Xiao and Ni 1999] [El-Gendy et al. 2003].

Using the RSVP signaling protocol in combination with IntServ (Integrated Services) gives the opportunity of having admission control for separate flows. Both a guaranteed service, with delay bound for specified traffic agreement, and a controlled load service, that simulates a lightly loaded network (a kind of statistical non-bounded service), are supported. The problem is, however, scalability. Hugh amounts of IP datagrams (packets) can pass the core Internet routers and each datagram should be inspected to discover which flow it belongs to and then handled according to specified QoS. Variable-sized headers make life even worse because simple offsets cannot be used. DiffServ (Differentiated Services) was developed with the scalability as the main driving force. With DiffServ, there is no treatment of separate flows. Instead, aggregation is introduced and the eight-bit ToS field (Type of Service) is used to distinguish among different traffic classes. All flows for a specific service class are treated in the same way (can vary with routing decision).

14.8 Networks for high-performance computing systems

High-performance computing systems are normally using a number of processors coupled together to form a parallel or distributed computer. The interconnection network is the heart of such a system. Since both high throughput and low delay are crucial for these systems, switched short-distance networks have been widely investigated and used (see, e.g., [Toda et al. 1994] [Rexford et al. 1998] for research considering real-time aspects), while wormhole routing is an important method to reduce the delay.

Wormhole routing is a kind of cut-through (i.e., only the header of the packet with the destination address must arrive before the switch/router can start to forward the packet to an output port [Kermani and Kleinrock 1979]) where it is not needed to store the whole packet in a router if an output port is busy [Dally and Seitz 1987]. Instead, a flow control signal is propagated back to the transmitter to order it, and intermediate routers, to stop sending. The transmission of the packet is resumed when the busy port becomes free.

Several methods to obtain real-time services over wormhole-routed networks have been proposed. One way is to divide the physical channels into multiple virtual channels, each with dedicated buffers in the routers [Dally 1992]. Packet deadlines can then be used in combination with different priorities for the virtual channels to schedule real-time traffic [Li and Mutka 1994]. Using TDMA to obtain real-time communication over wormhole networks has also been investigated [Garcia et al. 2000], while other examples are found in [Song et al. 1999] [Sundaresan and Bettati 1997].

RACEway from Mercury Computer Systems [Kuszmaul 1995] is an example of a network specially developed for parallel, embedded real-time systems. Circuit switching with source routing is used. Support for real-time traffic is obtained by using priorities, where a higher priority transmission preempts a lower priority transmission. The link bandwidth is 160 MByte/s. Infiniband is a more general emerging standard but likely to become important for both cluster computing and embedded systems. Support for different real-time traffic classes in Infiniband has been investigated [Pelissier 2000] [Alfaro et al. 2002].

14.9 Fiber-optic networks

To mention a few words about fiber-optic networks, networks taking advantage of the Wavelength Division Multiplexing (WDM) technique are chosen. The advantage with those networks compared to, e.g., FDDI, is that a number of channels (wavelengths of the light) can coexist in a totally passive optical network. The WDM star network is the most well known passive optical network.

A real-time protocol for packet switched communication in WDM star networks is described in [Yan et al. 1996]. The QoS associated with a real-time packet in the network is related to the probability of missing its deadline. The protocol tries to globally minimize the number of packets not managing to keep their QoS by adaptively changing the priority of queued packets. Although dynamic real-time properties are supported, the matter of the success or failure of a packet transmission depends on the global state of the network, and transmission success cannot be guaranteed in advance. A network that really supports hard real-time traffic is presented in [Jonsson et al. 1997]. The protocol is based on TDMA with semi-dynamic slot-allocation.

14.10 Wireless networks

Wireless networking is a field where, e.g., the high bit-error rate often must be specially treated. Moreover, the error characteristics can be rather bursty and location dependent. New scheduling algorithms must therefore be developed [Cao and Li 2001]. Other problems to cope with include the highly dynamic properties apparent in ad hoc (multi-hop) networks. Real-time protocols and scheduling algorithms have been proposed for both WLANs (Wireless LANs) [Gu and Zhang 2003] and for Wireless ATM [Frigon et al. 2001]. Routing issues for real-time communication in ad hoc networks have also been investigated [Chlamtac 1989] [Chakrabarti and Mishra 2001] [Bilstrup and Wiberg 1999], while another field is that of considering information theory to improve the chances to meet deadlines [Uhlemann et al. 2000].

14.11 Introduction to included papers

The four following appended papers gives a good coverage of the real-time communication research performed in ARTES. Also, they together provide good examples for several of the fields surveyed above.

The first paper presents novel work for the otherwise well explored CAN bus. Nolte, Nolin, and Hansson investigate the use of server-based scheduling with EDF for CAN, where a master node gives the other nodes credits to send in the next cycle. During the cycle, the native CAN protocol is used for medium access control. The master node keeps track of the other nodes' activities to, for example, be able to take care of unused resources when scheduling.

In the second paper, by Bergenhem and Jonsson, a ring network built up by fiber-ribbon links over which data and medium access control packets travel in separate fibers is described. The network has support for spatial reuse of the bandwidth to allow for several transmissions in different segments of the network simultaneously. Together with the slotted real-time protocol, this is shown to be a good solution for, e.g., high-performance signal processing applications.

In the third paper, Uhlemann, Rasmussen, and Wiberg present a framework for deadline dependent coding. This is a novel area of combining the fields of information theory and real-time communication, and is of special value for wireless communication systems because of the usually high bit-error rate. Both error correcting codes and retransmission schemes are discussed in the paper from a real-time point-of-view.

The fourth paper, by Hoang and Jonsson, presents work on using switched Ethernet for industrial real-time applications. It is described how to incorporate deadline scheduling in the switch and the end-nodes, still supporting the widespread TCP/IP suite. Moreover, deadline-partitioning schemes to divide the deadline (delay bound) for the different physical links to be traversed by logical connections are discussed. This can be used to remove bottlenecks, for example, between master nodes and the switch in a master-slave system.

Bibliography

[Alfaro et al. 2002] F. J. Alfaro, J. L. Sanchez, and J. Duato, "A Strategy to Manage Time Sensitive Traffic in InfiniBand," *Proc. Workshop on Communication Architecture for Clusters (CAC'02) in conjunction with International Parallel and Distributed Processing Symposium (IPDPS'02)*, Fort Lauderdale, FL, USA, Apr. 15, 2002.

[Alves et al. 2000] M. Alves, E. Tovar, G. Fohler, and G. Buttazzo, "CIDER - envisaging a COTS communication infrastructure for evolutionary dependable real-time systems," *Proc. for the Work in Progress and Industrial Experience Sessions, 12th Euromicro Conference on Real-Time Systems (ECRTS'2000)*, Stockholm, Sweden, June 2000, pp. 19-22.

[Alves et al. 2002] M. Alves, E. Tovar, F. Vasques, G. Hammer, and K. Rother, "Real-time communications over hybrid wired/wireless PROFIBUS-based networks," *Proc. 14th Euromicro Conference on Real-Time Systems*, pp. 142-151, 19-21 June, 2002.

[Arvind et al. 1991] K. Arvind, K. Ramamritham, and J. A. Stankovic, "A local area network architecture for communication in distributed real-time systems," *Journal of Real-Time Systems*, vol. 3, no. 2, pp. 115-147, May 1991.

[Bergenhem and Jonsson 2003] Bergenhem, C. and M. Jonsson, "A pipelined ring network – heterogeneous real-time in radar signal processing," *Proc. IASTED International Conference on Parallel and Distributed Computing and Networks (PDCN'2003)*, Innsbruck, Austria, Feb. 10-13, 2003, pp. 825-832.

[Bilstrup and Wiberg 1999] U. Bilstrup and P.-A. Wiberg, "Routing protocol for wireless real-time multihop network," *Proc. 11th Euromicro Conference on Real-Time Systems*, York, UK, June 9-11, 1999, pp. 5-7.

[Cao and Li 2001] Y. Cao and V. O. K. Li, "Scheduling algorithms in broad-band wireless networks," *Proceedings of the IEEE*, vol. 89, no. 1, pp. 76-87, Jan. 2001.

[Carpenzano et al. 2002] A. Carpenzano, R. Caponetto, L. Lo Bello, and O. Mirabella, "Fuzzy traffic smoothing: an approach for real-time communication over Ethernet networks," *Proc. 4th IEEE International Workshop on Factory Communication Systems*, pp. 241-248, Aug. 28-30, 2002.

[Chakrabarti and Mishra 2001] S. Chakrabarti and A. Mishra, "QoS issues in ad hoc wireless networks," *IEEE Communications Magazine*, vol. 39, no. 2, pp. 142-148, Feb. 2001.

[Chen et al. 1997] B. Chen, A. Sahoo, W. Zhao, and A. Raha, "Connection-oriented communications for real-time applications in FDDI-ATM-FDDI heterogeneous networks," *Proc. 17th International Conference on Distributed Computing Systems*, pp. 225-234, 27-30 May 1997.

[Chlamtac 1989] I. Chlamtac, "Bounded delay routing for real time communication in tactical radio networks," *Proc. IEEE MILCOM*, 1989.

[Dally 1992] W. J. Dally, "Virtual-channel flow control," *IEEE Transactions on Parallel and Distributed Systems*, vol. 3, no. 2, pp. 194-205, Mar. 1992.

[Dally and Seitz 1987] W. J. Dally and C. L. Seitz, "Deadlock-free message routing in multiprocessor interconnection networks," *IEEE Transactions on Computers*, vol. 36, no. 5, pp. 547-553, May 1987.

[Demers et al. 1989] A. Demers, S. Keshav, and S. Shenker, "Analysis and simulation of a fair queueing algorithm," *Proc. ACM SIGCOMM'89*, Austin, Texas, USA, Sept. 25-27, 1989, pp. 1-12.

[El-Gendy et al. 2003] M. A. El-Gendy, A. Bose, K. G. Shin, "Evolution of the Internet QoS and support for soft real-time applications," *Proceedings of the IEEE*, vol. 91, no. 7, pp. 1086-1104, July 2003.

[Fan et al. 2004] Fan, X., M. Jonsson, and H. Hoang, "Efficient many-to-many real-time communication using an intelligent Ethernet switch," *Proc. International Symposium on Parallel Architectures, Algorithms, and Networks (I-SPAN'2004)*, Hong Kong, China, May 10-12, 2004.

[Ferrari and Verma 1990] D. Ferrari and D. C. Verma, "A scheme for real-time channel establishment in wide-area networks," *IEEE Journal of Selected Areas in Communications*, vol. 8, no. 3, pp. 368-379, Apr. 1990.

[Ferreira et al. 2002] J. Ferreira, P. Pedreiras, L. Almeida, and J. A. Fonseca, "The FTT-CAN protocol for flexibility in safety-critical systems," *IEEE Micro*, vol. 22, no. 4 , pp. 46-55, July-Aug. 2002.

[FlexRay] The FlexRay Consortium: http://www.flexray.com/.

[Frigon et al. 2001] J.-F. Frigon, V. C. M. Leung, and H. C. B. Chan, "Dynamic reservation TDMA protocol for wireless ATM networks," *IEEE Journal on Selected Areas in Communications*, vol. 19, no. 2, pp. 370-383, Feb. 2001.

[Garcia et al. 2000] A. Garcia, L. Johansson, M. Jonsson, and M. Weckstén, "Guaranteed periodic real-time communication over wormhole switched networks," *Proc. ISCA 13th International Conference on Parallel and Distributed Computing Systems (PDCS-00)*, Las Vegas, NV, USA, Aug. 8-10, 2000.

[Gu and Zhang 2003] D. Gu and J. Zhang, "QoS enhancement in IEEE 802.11 wireless local area networks," *IEEE Communications Magazine*, vol. 41, no. 6, pp. 120-124, June 2003.

[Hassan et al. 2000] M. Hassan, A. Nayandoro, and M. Atiquzzaman, "Internet telephony: services, technical challanges, and products", *IEEE Communications Magazine*, vol. 38, no. 4, pp. 96-103, Apr. 2000.

[Hoang et al. 2002] Hoang, H., M. Jonsson, U. Hagström, and A. Kallerdahl, "Switched real-time Ethernet with earliest deadline first scheduling - protocols and traffic handling," *Proc. Workshop on Parallel and Distributed Real-Time Systems (WPDRTS'2002) in conjunction with International Parallel and Distributed Processing Symposium (IPDPS'02)*, Fort Lauderdale, FL, USA, April 15-16, 2002.

[Jonsson et al. 1997] M. Jonsson, K. Börjesson, and M. Legardt, "Dynamic time-deterministic traffic in a fiber-optic WDM star network," *Proc. 9th Euromicro Workshop on Real Time Systems (EWRTS'97)*, Toledo, Spain, June 11-13, 1997, pp. 25-33.

[Kandlur et al. 1991] D. D. Kandlur, K. G. Shin, and D. Ferrari, "Real-time communication in multi-hop networks," *Proc. 11th International Conference on Distributed Computing Systems*, pp. 300-307, May 20-24, 1991.

[Kermani and Kleinrock 1979] P. Kermani and L. Kleinrock, "Virtual cut-through: a new computer communication switching technique," *Computer Networks*, vol. 3, pp. 267-286, 1979.

[Kopetz and Grünsteidl 1994] H. Kopetz and G. Grünsteidl, "TTP - a protocol for fault-tolerant real-time systems," *Computer*, vol. 27, no. 1, pp. 14-23, Jan 1994.

[Kopetz and Bauer 2003] H. Kopetz and G. Bauer, "The time-triggered architecture," *Proceedings of the IEEE*, vol. 91, no. 1, pp. 112-126, Jan. 2003.

[Kopetz et al. 1989] H. Kopetz, A. Damm, C. Koza, M. Mulazzani, W. Schwabl, C. Senft, and R. Zainlinger, "Distributed fault-tolerant real-time systems: the Mars approach," *IEEE Micro*, vol. 9, no. 1, pp. 25-40, Feb. 1989.

[Kopetz et al. 2000] H. Kopetz, M. Holzmann, and W. Elmenreich, "A universal smart transducer interface: TTP/A," *Proc. Third IEEE*

International Symposium on Object-Oriented Real-Time Distributed Computing (ISORC'2000), pp. 16-23, Mar. 15-17, 2000.

[Kuszmaul 1995] B. C. Kuszmaul, "The RACE network architecture," *Proc. 9th International Parallel Processing Symposium (IPPS'95)*, Santa Barbara, CA, USA, Apr. 25-28, 1995, pp. 508-513.

[Kweon et al. 2000] S.-K. Kweon, K. G. Shin, and G. Workman, "Achieving real-time communication over Ethernet with adaptive traffic smoothing," *Proc. 6th IEEE Real-Time Technology and Applications Symposium (RTAS'2000)*, Washington, D.C., USA, 31 May - 2 June 2000, pp. 90-100.

[Leen and Heffernan 2002] G. Leen and D. Heffernan, "TTCAN: a new time-triggered controller area network," *Microprocessors and Microsystems*, vol. 26, no. 2, pp. 77-94, Mar. 17, 2002.

[Leen et al.] G. Leen, D. Heffernan, and A. Dunne, "Digital networks in the automotive vehicle," *Computing & Control Engineering Journal*, vol. 10, no. 6, pp. 257-266, Dec. 1999.

[Li and Mutka 1994] J.-P. Li and M. W. Mutka, "Priority based real-time communication for large scale wormhole networks," *Proc. Eighth International Parallel Processing Symposium*, Cancun, Mexico, Apr. 26-29, 1994, pp. 433-438.

[Malcolm and Zhao 1994] N. Malcolm and W. Zhao, "The timed-token protocol for real-time communications," *Computer*, vol. 27, no. 1, pp. 35-41, Jan. 1994.

[Metzler et al. 1999] B. Metzler, T. Harbaum, R. Wittmann, and M. Zitterbart, "AMnet: heterogeneous multicast services based on active networking," *Proc. IEEE Second Conference on Open Architectures and Network Programming (OPENARCH'99)*, pp. 98-105, 26-27 Mar. 1999.

[Mock and Nett 1999] M. Mock and E. Nett, "Real-time communication in autonomous robot systems," *Proc. Fourth International Symposium on Autonomous Decentralized Systems (ISADS'1999)*, Tokyo, Japan, 21-23 March 1999, pp. 34-41.

[Molle and Kleinrock 1985] M. Molle L. Kleinrock, "Virtual time CSMA: why two clocks are better than one," *IEEE Transactions on Communications*, vol. 33, no. 9, pp. 919-933, Sept. 1985.

[Montuschi et al. 1992] P. Montuschi, L. Ciminiera, and A. Valenzano, "Time characteristics of IEEE 802.4 token bus protocol," *IEE Proceedings-E (Computers and Digital Techniques)*, vol. 139, no. 1, pp. 81-87, Jan. 1992

[Nolte et al. 2003] T. Nolte, M. Sjödin, H. Hansson, "Server-based scheduling of the CAN bus," *Proc. IEEE International Conference on Emerging Technologies and Factory Automation (ETFA'03)*, vol. 1, pp. 169-176, Sept. 16-19, 2003.

[Parekh and Gallager 1994] A. K. Parekh and R. G. Gallager, "A generalized processor sharing approach to flow control in integrated services networks: the multiple node case," *IEEE/ACM Transactions on Networking*, vol. 2, no. 2, pp. 137-150, Apr. 1994.

[Pedro and Burns 1997] P. Pedro and A. Burns, "Worst case response time analysis of hard real-time sporadic traffic in FIP networks," *Proc. Ninth Euromicro Workshop on Real-Time Systems*, Toledo, Spain, June 11-13, 1997, pp. 3-10.

[Pelissier 2000] J. Pelissier, "Providing quality of service over Infiniband architecture fabrics," *Proc. Hot Interconnects*, 2000.

[Rexford et al. 1998] J. Rexford, J. Hall, and K. G. Shin, "A router architecture for real-time communication in multicomputer networks," *IEEE Transactions on Computers*, vol. 47, no. 10, pp. 1088-1101, Oct. 1998.

[Rushby 2001] J. Rushby, "Bus architectures for safety-critical embedded systems,"
First Workshop on Embedded Software (EMSOFT'2001), Oct. 8-10, 2001, Lake Tahoe, CA, USA. Published in Springer-Verlag *Lecture Notes in Computer Science*, vol. 2211, pp. 306–323.

[Schuh 1988] R. A. Schuh, "An overview of the 1553 bus with testing and simulation considerations," *Proc. 5th IEEE Instrumentation and Measurement Technology Conference (IMTC-88)*, pp. 20-25, Apr. 20-22, 1988.

[Song et al. 1999] H. Song, B. Kwon, and H. Yoon, "Throttle and preempt: a flow control policy for real-time traffic in wormhole networks," *Journal of Systems Architecture*, vol. 45, no. 8, pp. 633-649, Feb. 1999.

[Sundaresan and Bettati 1997] S. Sundaresan and R. Bettati, "Distributed connection management for real-time communication over wormhole-routed networks," *Proc. 17th International Conference on Distributed Computing Systems*, Baltimore, MD, USA, May 27-30, 1997, pp. 209-216.

[Stankovic et al. 2003] J. A. Stankovic, T. E. Abdelzaher, C. Lu, L. Sha, and J. C. Hou, "Real-time communication and coordination in embedded sensor networks," *Proceedings of the IEEE*, vol. 91, no. 7, pp. 1002-1022, July 2003.

[Tindell et al. 1994] Tindell, K. W., H. Hansson, and A. J. Wellings, "Analysing real-time communications: controller area network (CAN)," *Proc. 15th IEEE Real-Time Systems Symposium*, pp. 259-263, 1994.

[Toda et al. 1994] K. Toda, K. Nishida, E. Takahashi, and Y. Yamaguchi, "A priority forwarding router chip for real-time interconnection networks," *Proc. 15th IEEE Real-Time Systems Symposium*, pp. 63-73, 1994.

[Tovar and Vasques 1999] E. Tovar and F. Vasques, "Real-time fieldbus communications using Profibus networks," *IEEE Transactions on Industrial Electronics*, vol. 46, no. 6, pp. 1241-1251, Dec. 1999.

[Uhlemann et al. 2000] E. Uhlemann, P.-A. Wiberg, T. M. Aulin, and L. K. Rasmussen, “Deadline dependent coding - a framework for wireless real-time communication,” *Proc. 7th International Conference on Real-Time Computing Systems and Applications (RTCSA’00)*, Cheju Island, South Korea, Dec. 12-14, 2000, pp. 135-142.

[Varadarajan and Chiueh 1998] S. Varadarajan and T. S. Chiueh, ”EtheReal: a host-transparent real-time Fast Ethernet switch,” *Proc. 6th International Conference on Network Protocols*, 1998, pp. 12-21.

[Venkatramani and Chiueh 1994] C. Venkatramani and T. Chiueh, “Supporting real-time traffic on Ethernet,” *Proc. Real-Time Systems Symposium (RTSS’1994)*, pp. 282-286, Dec. 7-9, 1994.

[Verma et al. 1991] D. Verma, H. Zhang and D. Ferrari, “Guaranteeing delay jitter bounds in packet switching networks,” *Proc. Tri-Comm’91*, Chapel Hill, NC, USA, Apr. 1991, pp. 35-46.

[Wolf and R. Steinmetz 1997] L. C. Wolf and R. Steinmetz, ”Multimedia communication,” *Proceedings of the IEEE*, vol. 85, no. 12, Dec. 1997.

[Xiao and Ni 1999] X. Xiao and L. M. Ni, “Internet QoS: a big picture,” *IEEE Network*, vol. 13, no. 2, pp. 8-18, Mar./Apr. 1999.

[Yan et al. 1996] A. Yan, A. Ganz, and C. M. Krishna, “A distributed adaptive protocol providing real-time services on WDM-based LAN’s,” *Journal of Lightwave Technology*, vol. 14, no. 6, pp. 1245-1254, June 1996.

[Zhang 1995] H. Zhang, “Service disciplines for guaranteed performance in packet-switching networks,” *Proceedings of the IEEE*, vol. 83, no. 10, pp. 1374-1396, Oct. 1995.

[Zhao and Ramamritham 1987] W. Zhao and K. Ramamritham, “Virtual time CSMA protocols for hard real-time communications,” *IEEE Transactions on Software Engineering*, vol. 13, no. 8, pp. 938-952, Aug. 1987.

[Zhao et al. 1990] W. Zhao, J. A. Stankovic, K. Ramamritham, “A window protocol for transmission of time-constrained messages,” *IEEE Transactions on Computers*, vol. 39, no. 9, pp. 1186-1203, Sept. 1990.

[Zheng and Shin 1994] Q. Zheng and K. G. Shin, “On the ability of establishing real-time channels in point-to-point packet-switched networks”, *IEEE Transactions on Communications*, vol. 42, no. 2/3/4, pp. 1096- 1105, Feb./Mar./Apr. 1994.

[Zuberi and Shin 1995] K. M. Zuberi and K. G. Shin, “Non-preemptive scheduling of messages on controller area network for real-time control applications,” *Proc. Real-Time Technology and Applications Symposium (RTAS’1995)*, pp. 240-249, May 15-17, 1995.