

Real-time and Performance Aspects of ANSI/EIA-709.1 Networks

Alexander Bauer, Peter Rössler
{bauera, roessler}@ict.tuwien.ac.at
Institute of Computer Technology, TU Wien, Austria

Abstract - This paper discusses aspects of performance and real-time capabilities of the fieldbus protocol ANSI/EIA-709.1 which is spread out all over the world. The VENUS (Vienna Embedded Networking Utility Suite) framework based on a dedicated programmable multi-protocol controller, currently developed by the Institute of Computer Technology and an Austrian company is presented, which implements new functionality to open new areas of applications for this widely used and well known network protocol.

I. TERMS AND DEFINITIONS

Fieldbusses are used to connect intelligent nodes in distributed control systems. They can replace conventional centralized solutions to save costs in cabling, system integration and maintenance. The applications of field area networks often require real-time capabilities to ensure that crucial deadlines can be met. For some applications high network throughput is required additionally.

Systems with these high performance features can be divided into soft- and hard real-time systems [1]. Hard real-time systems must meet their deadlines in order to avoid catastrophic results like e.g. in aircraft control. In soft real-time systems it is tolerated that occasionally deadlines can not be met. For the transmission of audio data e.g. it is not critical to lose a certain part of the information due to the high redundancy of spoken language, see also Section IV.3. However, in this case a (relatively) high data throughput is demanded.

A hard real-time communication system ensures that every transmitted message is delivered within a guaranteed delivery time. For soft real-time systems a certain (high) probability for adhering to these temporal requirements would be sufficient. The physical communication medium in combination with the network protocol is responsible for meeting the demanded deadlines and providing the needed communication performance.

II. THE ANSI/EIA-709.1 STANDARD

The ANSI/EIA-709.1 standard [2], developed by the American company Echelon in the late 80s under the name LonTalk, defines a network protocol which was originally suited to the demands of applications for

building- and home automation. Especially in the area of building automation we are talking about large networks, that is a total network expansion of several kilometers and thousands of nodes in one network. This stands in contrast to controller networks for industrial automation or drive-by-wire systems which are usually limited to local distances with a much fewer number of nodes. Hence, in most of these network protocols only the functionality of layers 1, 2 and 7 of the OSI (Open Systems Interconnection) -model is implemented which fits the needs of their application areas [3].

Due to the tremendous number of nodes in networks for building automation the resulting network traffic has to be limited which can be done by separating the overall traffic to local channels. For this purpose OSI-layer 3 (network layer) was implemented in the EIA-709.1 standard. Moreover the functionality in layers 4 and 5 (transport and session layers) provide a reliable point-to-point connection by supporting several types of services like acknowledged services for example. For this service type the sender of a message always gets information whether the sent message was received by all addressed nodes or not.

All in all EIA-709.1 nodes implement functionality of all the seven OSI-layers for a wide range of applications in the building- and home automation area like light control, energy management, access control or HVAC-systems [4].

III. COMMON EIA-709.1 IMPLEMENTATIONS

Today's common implementation of the EIA-709.1 protocol is the so called Neuron-Chip, a semiconductor device which is manufactured by the companies Cypress, Motorola and Toshiba [5]. Three 8-Bit-CPU's contained in the Neuron-Chip are processing layer 2-7 of the protocol and the application program (see Fig. 1, left side).

For a maximum of flexibility concerning the communication media, layer 1 (physical layer) can be easily exchanged by connecting different types of external network transceivers (RS-485, IR, RF, power line, fiber optic, ...). Hence, the Neuron-Chip is a cost effective solution for building general purpose low-cost network nodes.

A. Network Arbitration

Access to the physical media is controlled by a predictive p-persistent CSMA (Carrier Sense Multiple Access) algorithm which is processed in layer 2 of all network nodes. This algorithm is depicted in Fig. 2.

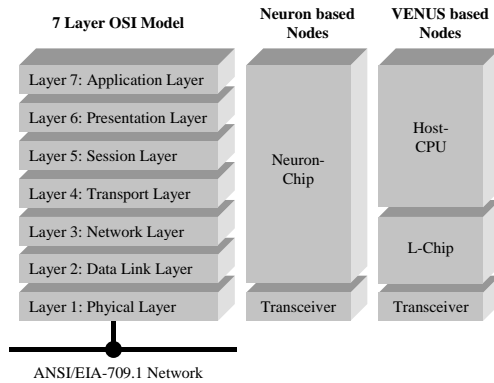


Fig. 1 - ANSI/EIA-709.1 in the OSI-model.




Fig. 2 - ANSI/EIA-709.1 Bus Arbitration.

In general an EIA-709.1 frame consist of a preamble for bit synchronization purposes, the address information containing source and destination address and the data field containing the user data and other information like the service type of the message. A CRC (Cyclic Redundancy Check) is transmitted with every frame for the purpose of detecting corrupt frames.

After a bus idle time which is also called $\beta 1$ -time in the EIA-709.1 terminology, bus arbitration starts with a number of $0 \leq p \leq 128$ priority time slots. Since the priority time slots are not used for most typical applications it is assumed that $p = 0$.

After the priority slots the actual p-persistent CSMA phase starts. During this time every node who wants to transmit a data packet on the network picks one of $r = 16 * BL$ random time slots and tries to send the packet by starting to transmit the preamble. If another node has picked an earlier time slot, the node has to wait for the next arbitration phase after the packet from the other node was transmitted and the network becomes idle again.

However, collisions of packets transmitted at the same time by two or several nodes may occur. To minimize the number of collisions the constant BL is changed dynamically between 1 and 63 depending on the estimated network traffic.

B. Performance on the Network Side and on the Upper Protocol Layers

The two columns on the left side of Table 1 show the packet throughput and the response time of Neuron based EIA-709.1 nodes (running @ 10 MHz and 20 MHz clock

frequency) depending on different bit rates on the network.

The first value in a cell gives the theoretical maximum packet throughput on the network in packets per second based on the p-persistent CSMA arbitration scheme shown in Fig. 2 (BL = 1, 2 bytes user data).

Bit rate	Packets/s on the network and packets/s and response time over the whole protocol stack		
	10 MHz Neuron	20 MHz Neuron	VENUS
1.25 Mbps	1500 P/s net 800 P/s total 7 ms	2500 P/s net 1500 P/s total 4 ms	10000 P/s 100 μ s + Host
2.5 Mbps		2900 P/s net 1600 P/s total 3.5 ms	20000 P/s 50 μ s + Host
5 Mbps			40000 P/s 25 μ s + Host

Tab. 1 - Several Parameters of Performance.

The second value gives the maximum packet throughput over the whole 7 layers of the protocol stack while the third value gives the measured response time. The application program in the receiving node for this test does nothing more than counting incoming packets.

IV. THE VENUS APPROACH

As described in Chapter III the Neuron-Chip is a low-cost solution which is best suited for applications in the area of building- and home automation. But for some applications where high volumes of data have to be transmitted (file transfer, audio and video transfer, master nodes, ...) there is a need for more performance [6]. Moreover by using a CSMA mechanism no hard deadline (as defined in Chapter I) for messages which are transmitted on the network can be guaranteed. This may be a problem for time critical applications for example in the area of industrial automation.

The VENUS (Vienna Embedded Network and Utility Suite) framework increases the performance of common EIA-709.1 nodes by using a dedicated programmable multi protocol controller which is currently developed by the Institute of Computer Technology and the Austrian company LOYTEC electronics GmbH. As shown on the right side of Fig. 1, the time critical layers 2 and 3 of the protocol are implemented in the L-Chip while the upper layers 3-7 and the application program are processed by a Host-CPU. The advantages and gains in performance of this approach are discussed in the following sections.

A. Performance on the Lower Layers

By looking at Fig. 2 one may ask the purpose of the $\beta 1$ -time. The reason for this bus idle time lies in the computing time which is required to process the lower protocol layers after a packet was received from the network. Because these functionality is not handled by a CPU as in the Neuron-Chip but a dedicated controller (L-

Chip) the β 1-time can be reduced from former 360 bits (@ 1.25 Mbps bit rate) to 20 bit times. For the same reason the preamble length can be reduced from 160 bits to 4 bits and the width of the random slots can be reduced from former 30 bits per slot to less than 1 bit per slot.

Moreover the L-chip runs at a maximum clock frequency of 40 MHz which allows bit rates on the network up to 5 Mbps. Finally layer 3 address comparison is handled by the L-chip on the fly (see next section). All in all this results in a packet throughput up to 40000 packets (see Tab. 1, right column).

B. Performance on the Upper Layers

Apart from the performance issues in lower layers of the EIA-709.1 protocol also layers 3-7 must be suited for high performance and real-time capabilities. The network layer implementation (layer 3) within the L-Chip performs the address filtering for incoming packets. This filter function is performed during the corresponding packet is received and stored in the packet receive buffer of the L-Chip. Hence, there is no performance degradation through layer 3 and the real-time parameters of the system are not affected. Additionally, the address filter disburdens the Host-CPU from processing every single packet on the network.

Since all higher protocol layers (4-7) are processed in the host, see Fig. 1, both the packet throughput and the response time of the node are scalable to a great extend by choosing an appropriate CPU. Further, for real-time applications a real-time operating system can be used to make possible the adherence to deadlines both for the higher protocol layers and the applications running on top of the protocol stack. Within the VENUS framework a portable protocol stack, called ORION, has been developed to suit the needs of high speed applications. Several advanced features have been implemented to overcome the limitations of the Neuron-Chip concerning packet throughput. A good example for the new functionality is the possibility of using multiple transaction spaces in the ORION protocol stack. This feature enables the node (in contrast to the Neuron-Chip) to process a high number of transactions with other nodes in parallel. Especially for real-time applications it cannot be accepted that a node is blocked by a single transaction, possibly for several seconds, because of communication problems with just one node. In the ORION implementation, even if there is a high response delay of one or several nodes, the protocol stack is not blocked and can process other transactions in the meantime.

However, for hard real-time systems, where strict adherence to deadlines and a deterministic behavior of the entire system are mandatory, the EIA-709.1 protocol implementation has to receive major modifications. There are plans to develop a special version of the ORION stack in order to achieve such capabilities. More information on this subject is presented in the following section.

C. Real-time Aspects

Due to the structure of conventional Neuron-Chip based nodes the EIA-709.1 protocol is not optimized for high performance and real-time applications. One important step in this context is the development of the network controller L-Chip. For soft real-time systems the improvements concerning packet throughput and response time are sufficient as long as the system has enough bandwidth for high-load situations. For hard real-time systems some problems with the EIA-709.1 protocol must be resolved in order to guarantee a deterministic temporal behavior. The following sections describe these problems and discuss possible solutions using the L-Chip.

Bus Arbitration

Since the CSMA protocol of EIA-709.1 still allows the occurrence of collisions on the network it is not possible to guarantee maximum transmission times. Furthermore starvation of nodes can not be avoided. Therefore a deterministic arbitration protocol must be used, which allows every node to send its messages within a certain deadline. The Medium Access Control (MAC) sublayer of the L-Chip could be modified to implement such a deterministic bus arbitration procedure. Starting from the p-persistent CSMA protocol of EIA-709.1 only the priority slots could be utilized. Additionally a mechanism for starvation avoidance would have to be established.

Corporation with Conventional Systems

In complex field area networks usually only a small part of the system requires high network performance and real-time capabilities. Simple control nodes e.g. for switches and lamps can be implemented using conventional inexpensive hardware (Neuron-Chips). Therefore considerations about the interconnection between the high-performance nodes and the rest of the network have to take place. It should be possible to connect nodes of all kind to the same bus but usually repeaters, routers or gateways will be used to create a firewall between critical and uncritical sections of the network. Since the L-Chip can be configured for both high- and low-speed channels, gateways can easily be implemented using one L-Chip for each channel.

Fault Tolerance

In safety-critical systems a single point of failure can not be tolerated. For a network this means that a failure in one node must not lead to a failure of the entire system. It is not acceptable that the real-time abilities of the network are dependent on single nodes. Even permanent physical failures like defective wires shall be considered. A cut wire could e.g. lead to two separate sub-systems which can keep the maximum response times amongst themselves. Redundancy in the form of several channels working in parallel can easily be implemented using

several L-Chips and possibly several Host-CPU's in each node.

Quality of Service

For reliable packet transmission EIA-709.1 uses different acknowledged and repeated services. On one hand these services provide the possibility of error detection and error correction within layers 4 and 5 of the protocol. On the other hand additional packets (acknowledgements, retries) are generated which can lead to unpredictably high network loads. Therefore in the implementation of the higher layers each service must be examined and possible adapted to retain the adherence to deadlines for all nodes.

Clock Synchronisation

One basic demand of real-time network applications is a clock synchronisation mechanism that allows exact time stamping of events as well as exact temporal control of devices in a distributed system. The L-Chip offers several functions to timestamp incoming and outgoing packets. It is also possible to connect an external clock synchronisation unit like the UTCSU (Universal Time Coordinated Clock Synchronisation Unit) [7].

Membership Service

Membership services [8] are used to distribute the functional status of nodes in a network. If nodes are e.g. not responding to requests they should be considered inactive and the required actions must be performed like raising an alarm or redirecting messages to other (still working) nodes. The required functionality would have to be implemented within the network management and diagnostic layers of the protocol.

D. Real-time Applications

Like mentioned in Chapter I real time systems can be classified as soft- or hard real-time systems according to their profile of requirements. A typical soft real-time communication application would be the transmission of video data over an EIA-709.1 channel using L-Chips, see Fig. 3.




Fig. 3 - Soft Real-Time Application.

A hard real-time application for fieldbuses would be e.g. a smoke detector system. In this case it is important, that alarm messages are guaranteed to be delivered, lost packets or packets that arrive after their deadline can lead to catastrophic results.

V. SUMMARY

Despite the rather complex structure of the network protocol as well as the drawbacks concerning real-time behaviour, EIA-709.1 is an interesting subject for research towards high-speed and real-time capabilities because of its popularity and the growing demand for more system performance and predictability. The development of the L-Chip shows the possibility of such improvement of a well established fieldbus concept. Both the high-performance functionality already implemented as well as new approaches for further developments towards real-time capabilities were presented in this paper.

REFERENCES

- [1] H. Kopetz, Real-Time Systems, Design Principles for Distributed Embedded Applications, Kluwer Academic Publishers, 1997.
- [2] ANSI/EIA-709.1-A-1999, ANSI/EIA Standard, Control Network Protocol Specification, 1999.
- [3] VDI/VDE 3687 Draft, Selection of Fieldbus-Systems by valuation of their performances for different applications, 1997.
- [4] D. Dietrich, D. Loy, H. Schweinzer, LON-Technologie, Verteilte Systeme in der Anwendung, Hüthig Verlag Heidelberg, 1998.
- [5] LonWorks, Technology Device Data, Rev.2, Motorola data book, 1996.
- [6] T. Rauscher, N. Reiter, H. Schweinzer, S. Soucek, Adressierungs- und Leistungsgrenzen großer Neuron-Chip-basierter LonWorks-Installationen, e&i 117 Jg 2000, H. 5, Springer Verlag, 2000.
- [7] D. Loy, GPS-Linked High Accuracy NTP Time Processor for Distributed Fault-Tolerant Real-Time Systems, Österr. Kunst- und Kulturverlag, Wien, 1997.
- [8] F. Cristian, Agreeing on Who is Present and Who is Absent in a Synchronous Distributed System, Proc. 18th Int. Symposium on Fault-Tolerant Computing, p. 206-211, Japan, 1988.