Institut für Automation
Abt. für Automatisierungssysteme

Technische Universität Wien

Projektbericht Nr. 183/1-95
August 1999

Basic Features of the Wireline/Wireless Factory/Facility Fieldbus

Ulrich Schmid

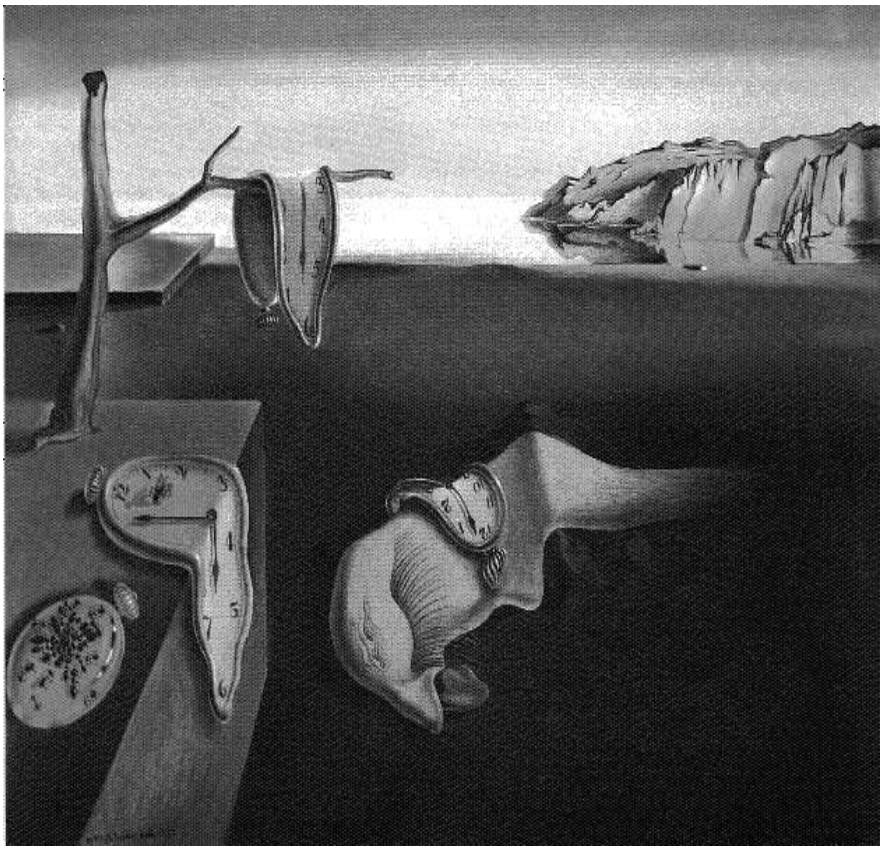

Salvador Dali, "Die Beständigkeit der Erinnerung"

Basic Features of the Wireline/Wireless Factory/Facility Fieldbus

Ulrich Schmid
Technische Universität Wien
Department of Automation
Treitlstraße 1, A-1040 Vienna
Email: s@auto.tuwien.ac.at

August 1999

Abstract

The project W2F (Wireline/Wireless Factory/Facility Fieldbus) aims at the development of a next-generation LAN/fieldbus for distributed automation. Targeted to the major areas factory automation and home/facility automation, W2F will employ spread-spectrum CDMA technology both on wireline and wireless media. Major design goals are full distribution, security, fault-tolerance, and real-time issues as well as flexibility with respect to wireline/wireless interconnections. In sharp contrast to existing fieldbusses, W2F will be proven to provide its fault-tolerance, security and real-time properties —or degrade gracefully— even in case of failures, attacks, and overload. Therefore, it should be a suitable basis for building up dependable and safety-critical distributed real-time applications.

Keywords: real-time computer communications, fieldbusses, wireless LANs, spread spectrum CDMA, fault-tolerant distributed real-time systems, dependability, security and integrity.

1 Motivation

It is well-known in practice that existing fieldbus technology cannot cope with the fault-tolerance and real-time requirements of future applications. Yet, in view of the existing standards' dominance, we argue that there is not much hope for a new fieldbus to get acceptance without either

(1) being compatible with some existing fieldbus, or
(2) providing "revolutionary" additional features.

In our project W2F (Wireline/Wireless Factory/Facility Fieldbus)[1], we will explore how far one can get with (2) by developing a fieldbus that employs spread-spectrum CDMA technology both on wireline and wireless media.

[1] Supported by the Austrian START-programme Y41-MAT.

Apart from treating wireline and wireless interconnections uniformly, this approach entirely avoids a shared channel and its multiple access problem, provides superior noise immunity, and allows to exploit all the communications bandwidth available in (existing) cabling infrastructure. The various protocols implementing W2F will be proven to provide its fault-tolerance, security and real-time properties —or degrade gracefully— even in case of failures, attacks, and overload. Therefore, it should be a suitable basis for building up dependable and safety-critical distributed real-time applications.

2 Topology

We consider an arbitrary topology of *subnets* (SNs) interconnected by an arbitrary number of *routing nodes*, which are (not necessarily dedicated) nodes participating in two or more SNs. A non-redundant communications architecture, as used in most existing applications, results in a tree-like topology of SNs interconnected by a single routing node. Figure 1 shows an example of a slightly more redundant communications architecture.

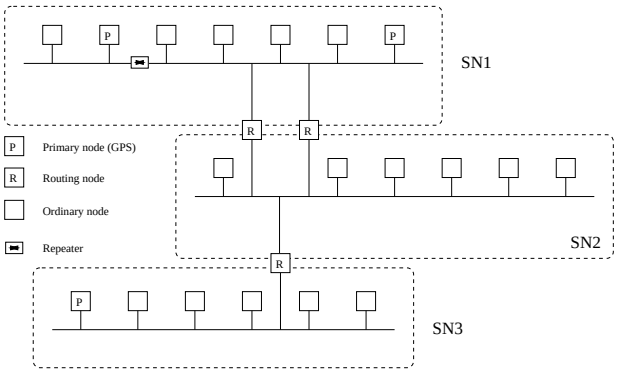

Figure 1: *Example of a system topology with three subnets interconnected by several routing nodes*

Each SN is made up of one or more *segments*, which are interconnected by *repeaters* that simply pass through any traffic. Any node in a single SN uses the same frequency band and the same set of CDMA codes for communication.

We impose the following rules for drawing SN borders:

(1) Non-routing nodes can be in at most one SN at any time.

(2) All nodes in an SN must normally be reachable[2] from each other.

(3) A single SN should contain all nodes that are normally directly reachable from each other (in particular, an SN should be closed w.r.t. physical segments),

(4) A single SN should contain all nodes being at the same "level of abstraction" (i.e., an SN should be closed w.r.t. logical domains).

Clearly, those restrictions do not prevent an SN to span several unrelated logical domains (a collection of nodes that collaborate to provide some specific functionality) and/or multiple physical segments. Moreover, SNs can be overlapping[3] in that routing nodes are contained in multiple SNs.

We further assume that

- there are up to 1000+ nodes in an SN,
- the interconnection medium within a single segment of the SN may either be a single wire or wireless, employing a specific frequency band and CDMA code set for communication,
- repeaters can connect wireline/wireline or wireline/wireless,
- most nodes are statically assigned to their SN(s),
- (mobile) nodes may join and leave (wireless) SNs.
- routing nodes are ordinary or dedicated nodes with multiple transmit and receive units,
- redundant paths are normally exploited only for communication between routing nodes, not for interconnected SNs as a whole,[4]

[2]Two nodes are normally reachable from each other if they can communicate with each other in absence of total partitioning. In case of wireless media, this does not mean directly reachable from each other but might involve a path of intermediate (non-routing) nodes.

[3]Multiple cells within an single wireless SN must be realized via the group membership protocol (see Section 4).

[4]The philosophy behind this is: "If you need a redundant path from node p to q for reliability reasons, provide it by attaching both to another SN." Non-direct redundant paths between p and q are only exploited in case of partitioning. Without this restriction, all adjacent SNs had always to cope with the whole traffic of each other, apart from the fact that a non-direct redundant path is not as good as a direct path.

The communication patterns expected in the system are

- a few (often a single) request/reply message exchanges between a client node and usually many different server nodes, with end-to-end message delivery times in the few ms-range (e.g. for alarm messages),
- periodic messages and message exchanges with periods 100 ms ... 100 s and ms-range delivery times (e.g. for sensor polling),
- streamed data up to 1 Mb/s (e.g. for CCD-camera data).

Both point-to-point and multicasting communications will be required. One can assume, however, that the overall traffic is bursty, with relatively low average load.

3 Spread-Spectrum CDMA

Unlike all other existing fieldbusses, we will not employ simple baseband transmission on wireline interconnections. Instead, we will use spread-spectrum CDMA for transmission both on wires and wireless, which offers a number of interesting prospects:

- Simultaneous real-time data transmission between all peers
- Potentially infinite channels between two peers
- Cryptography and inherent security at the lowest system level
- Receiver-controlled peer selection
- Superior noise-immunity
- Exploitation of all the available bandwidth of (existing) cables

The most obvious disadvantage of spread-spectrum transmission is much higher node interfacing costs. However, we think that the advances in digital and mixed mode ASIC technology will eventually drop it to a tolerable level.

Still, there are several other difficulties that remain to be solved. General CDMA-related problems are

- How to generate potentially infinitely many cryptographically secure codes?
- How to rule out (timing-)disturbance due to multipath reception?
- How to accomplish multi-user interference reduction (MUIR) etc. without knowing all codes?
- How to cope with low-powered units?
- How much bandwidth can we utilize e.g. on dedicated twisted-pair or coaxial cables?

- How to build wireline/wireline and, in particular, wireline/wireless repeaters without disturbing the channel characteristics too much?

A major source of problems, however, is the required many-to-many communication, which raises questions like

- How to accomplish many-to-many communication with a limited number (10 ...100) of receiver channels?
- How to achieve a bit error rate comparable to baseband networks?
- How to achieve very small and deterministic code acquisition time required for real-time guarantees of sporadic messages?

Another important issue that requires support at this level is interference detection and location of an interferer's position. In fact, aiming at a fault-tolerant and secure system, effective counter measures against jamming and interference have to be provided by W$_2$F:

- Location of short circuits and cable breaks for wireline channels
- Selective filters for increasing immunity to brute-force CW jamming
- Advanced detection of pulsed CW and wideband interference

Our research will focus on novel CDMA schemes for solving those problems, which will exploit techniques/particularities like the following:

- Configurable multichannel CDMA-receivers
- Advanced MUIR techniques
- Lowest-level fault-tolerant clock synchronization for system-wide chip-level synchronization
- Exploit accurate transmission delay and position information for power control and interferer locationing
- Primarily static system structure
- "Public key code sequences"

Analytical and simulation studies will be used to assess our solutions. Moreover, a prototype implementation will eventually be built to facilitate experimental evaluation.

4 Basic Protocols

The major part of W$_2$F is a complete, proven-correct protocol suite that implements a variety of services at the SN level. A suite of higher-level protocols (see Section 5) will transparently extend those services across SN boundaries. The most important basic services are:

- Secure admission control
- Reliable datagrams
- Overload protection
- Fault-tolerant position acquisition
- Topology management (transient disconnections & hidden nodes within an SN)
- Fault-tolerant external clock synchronization
- Interference detection and locationing
- Connection management
- Group management
- Atomic multicasting

The protocols will provide fault-tolerance, security and real-time performance guarantees under normal conditions and degrade gracefully under exceptional conditions like partitioning and denial-of-service attacks, for example. Particular emphasis will also be put on (auto-)configuration, monitoring and maintenance, since ease of use is a key feature for any fieldbus.

The envisioned protocols will differ from existing ones by the fundamental role dedicated to some global knowledge of *position* and *time* at all nodes. More specifically, a fault-tolerant topology service will be responsible for continuously providing an accurate view of the positions (and hence mutual distances) of all nodes. Using this information, the quality of service of many other protocols can be considerably improved. This is particularly true for the clock synchronization service, which will use similar techniques as developed in our SynUTC project [SKM+00] to establish a highly accurate common notion of time at all nodes.

By making available synchronized interval clocks with very high accuracy at the lowest layer, we can eventually rely upon timestamped messages (which are usually required by the atop running applications anyway) in all our protocols. For example, contrasting usual approaches based on sequence numbers, our *Sequenced Synchronized Clock Message Protocol* SSCMP [SP95] provides sequenced at-most-once delivery of messages by means of a clock-based connection management protocol. Unlike handshake-based ones, this protocol is also well-suited for W$_2$F's sporadic request/reply communication patterns. Note that the setup overhead of usual handshake-based connection management protocols would reduce the throughput down to 50%.

Last but not least, synchronized clocks enable both synchronous atomic broadcasting and time-driven transmission scheduling algorithms. Therefore, our approach offers conceptual coherence and improved performance at the same time. Graceful degradation —at least fail-awareness [Fet97]— can be achieved

by using the accuracy information made available by interval-based clock synchronization.

Our research will focus upon proven-correct protocol design and analytical worst case performance analysis; simulation will also be used where appropriate. Moreover, we will build up a reasonably complete prototype implementation for experimental evaluation.

5 Higher-Level Protocols

According to Section 4, the basic protocols of W$_2$F do not support

- routing of messages across SN borders,
- multicast groups exceeding SN borders,
- transparent handover/roaming of mobile nodes.

Whereas this clearly restricts transparent communications, it is nevertheless true that our basic protocols are sufficient for certain applications: For example, if a monitoring node n_m in a higher-level SN1 needs to track the data from a sensor node n_s distributed via atomic multicasting in a lower-level SN2, it would suffice to add n_m to n_s's multicast group. Nevertheless, this solution not only requires a complex multicast/group membership protocol, but is also arguable at the application level: Usually, the message sent by n_s will have the level of abstraction of SN2, i.e., will usually carry information that is meaningless at the level of SN1. Thus, the monitoring node n_m will be forced to throw away most of the message content and, more problematic, employ data manipulations usually only known to SN2-nodes. For that reason, existing solutions usually rely upon some kind of proxy representing the sensor connected to n_s at some SN1/SN2 routing node, which in turn provides n_m with the required data. Communication protocols that work transparently across SN boundaries are hence not needed here.

Still, higher-level protocols implementing such features will eventually be needed by more advanced applications. In order to retain simplicity and performance of the basic protocols, W$_2$F will provide this functionality atop of the latter: If a multicast group spanning several SNs —violating SN-rule (4)— is needed for some reason, a (more costly) higher-level protocol must be used to achieve the multicast functionality usually provided by the basic protocols at the SN level.

The higher-level protocols of W$_2$F will hence offer the following services transparently across SN boundaries:

- Naming service
- Routing across SN borders
- Global topology management
- System configuration, monitoring and maintenance
- Global group management
- Global atomic multicasting

We will build those protocols atop the basic ones, i.e., pursue a layered approach. This allows the design of well-defined and high-performance basic protocols ("intra-SN"), which can be used as building blocks for almost any high-level protocol ("inter-SN"). Note that flexibility w.r.t. higher-level protocols is particularly important to ensure interoperability with existing standard protocols.

6 Hardware Architecture

The envisioned protocols need certain hardware support, which is to be integrated with the required communications facilities. At the lowest level, we have to deal with

- wireless/wireline interfacing,
- power supply over wire,
- low-level security/encryption mechanisms,
- very high-accuracy synchronized clocks.

Note that W$_2$F will exploit CDMA's superior noise immunity to use the data communications wire for low-voltage power supply of nodes without dedicated power supply. A recharge scheduling protocol will ensure regular recharging of all battery-powered nodes, despite of the limited power available through the network.

In addition, there are also basic protocols like SSCMP that must be implemented partly in hardware for performance reasons. Another important issue affecting the hardware architecture is the idea of employing secure coprocessors for establishing a trusted and continuously operational computing base for the whole system.

Our research will focus on how to efficiently provide the required mechanisms on top of state-of-the-art hardware/software technology. Compatibility with COTS components and ease of maintenance will also be a major issue here.

7 Research Areas

The project involves several important research areas, which are very active fields of their own. Therefore, our particular work can rely upon a substantial body of existing work:

- *CDMA spread-spectrum communications* [Goi98]
- *High-accuracy fault-tolerant clock synchronization* [SKM+00], [FC97]
- *Cryptography, security protocols & proof logics* [SW99], [GSG99]
- *Basic protocols & formal verification* [Fet97], [SP95], [Mul93]

Acknowledgements

I am indebted to Christof Fetzer, Alois Goiser and Thilo Sauter for their valuable suggestions and comments on earlier versions of this paper.

References

[FC97] Christof Fetzer and Flaviu Cristian. Integrating external and internal clock synchronization. *J. Real-Time Systems*, 12(2):123–172, March 1997.

[Fet97] Christof Fetzer. *Fail-Awareness in Timed Asynchronous Systems.* Dissertation, University of California, San Diego, Computer Science, 1997.

[Goi98] Alois M.J. Goiser. *Handbuch der Spread-Spectrum Technik.* Springer, Wien, New York, 1998.

[GSG99] S. Gritzalis, D. Spinellis, and P. Georgiadis. Security protocols over open networks and distributed systems: formal methods for their analysis, design, and verification. *Computer Communications*, 22:697–709, 1999.

[Mul93] Sape Mullender. *Distributed Systems.* ACM Press/Addison Wesley, New York, 2nd ed. edition, 1993.

[SKM+00] Ulrich Schmid, Johann Klasek, Thomas Mandl, Herbert Nachtnebel, Gerhard R. Cadek, and Nikolaus Kerö. A Network Time Interface M-Module for distributing GPS-time over LANs. *J. Real-Time Systems*, 18(1), 2000. (to appear).

[SP95] Ulrich Schmid and Alfred Pusterhofer. SSCMP: The sequenced synchronized clock message protocol. *Computer Networks and ISDN Systems*, 27:1615–1632, 1995.

[SW99] Sean W. Smith and Steve Weingart. Building a high-performance, programmable secure coprocessor. *Computer Networks*, 31:831–860, 1999.