

Defining a CAN-based Infrastructure for Fault-Tolerant Real-Time Distributed Computing

José Rufino
ruf@digitais.ist.utl.pt
IST-UTL*

Paulo Veríssimo
pjuv@di.fc.ul.pt
FC/UL†

Guilherme Arroz
pcegsa@alfa.ist.utl.pt
IST-UTL

Abstract

Standard fieldbuses are today a cost-effective solution for distributed computer control systems. However, the efficient implementation of fault-tolerance and real-time mechanisms on the simple fieldbus environment presents non-negligible problems. This paper outlines our approach on the use of the Controller Area Network (CAN) as an off-the-shelf component in the design of fault-tolerant real-time distributed systems.

1 Introduction

The design and implementation of distributed computer control systems intended for real-world interfacing, i.e. integrating sensors and/or actuators, have increasingly been based on standard fieldbuses as an alternative to specialized and thus costly architectures [4]. The development of applications for such environments may greatly benefit from the availability of services such as clock synchronization, reliable group communication, membership and failure detection.

However, the migration of fault-tolerant communication systems to the realm of fieldbuses presents non-negligible problems, some of them addressed by our ongoing research in the context of CAN, the Controller Area Network [11, 8, 6, 5]. CAN is a fieldbus that has assumed increasing importance and widespread acceptance in application areas as diverse as shop-floor control, robotics or automotive [3].

This paper outlines our approach on how to use CAN as an off-the-shelf component in the design of fault-tolerant real-time distributed systems. The paper is organized as follows: Section 2 analyzes CAN dependability; Section 3 defines the system model; CAN non-stop operation in the presence of medium failures is discussed in Section 4; Section 5 addresses CAN inaccessibility and its impact on hard real-time operation. The failure detection, membership and reliable group communication services are discussed in

the remaining sections. A companion clock synchronization service can be found in [5].

2 Dependability of CAN

CAN uses a twisted pair cable as transmission medium. The CAN physical layer specified in [3] allows tolerance of some cabling faults (one wire open/short failures), by switching from differential to single-wire operation. However, this mechanism can not provide resilience to the simultaneous interruption of both wires.

CAN is a multi-master fieldbus. Bus signaling takes one out of two values: *recessive*, otherwise the state of an idle bus; *dominant*, which always overwrites a recessive value. This behavior, together with the uniqueness of frame identifiers, is exploited for bus arbitration. A *carrier sense multi-access with deterministic collision resolution* policy is used. The node transmitting the frame with the lowest identifier always goes through and gets the bus. Frames that have lost arbitration or have been destroyed by errors are automatically retransmitted. A *frame* is a piece of encapsulated information traveling on the network. It may contain a *message*, a user-level piece of information.

Though CAN fault-confinement and error detection mechanisms ensure that most failures are perceived consistently by all nodes [7], some subtle errors can lead to inconsistency. Inconsistent frame omissions occur when faults hit the last but one bit of a frame at some recipients¹, tagged \times set in Figure 1-B. This may lead to: the message to be accepted in duplicate by the recipients in the \bullet set of Figure 1-B, upon retransmission; inconsistent message omission, if the sender fails before retransmission. A thorough discussion of these failure scenarios can be found in [8]. However infrequent they may be, the probability of its occurrence is high enough to be taken into account, at least for highly fault-tolerant applications of CAN.

*Instituto Superior Técnico - Universidade Técnica de Lisboa, Avenida Rovisco Pais, 1049-001 Lisboa, Portugal. Tel: +351-1-8418397 - Fax: +351-1-8417499. NavIST Group CAN WWW Page - <http://pandora.ist.utl.pt/CAN>.

†Faculdade de Ciências da Universidade de Lisboa, Portugal. Navigators Home Page: <http://www.navigators.di.fc.ul.pt>.

¹The set may have only one element. Examples of causes for inconsistent detection are: electromagnetic interference or deficient receiver circuitry.

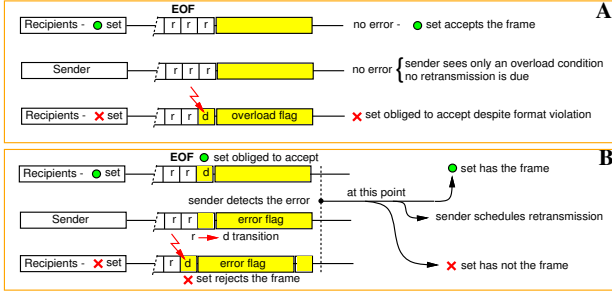


Figure 1: Inconsistency in CAN error handling

3 System Model

We enumerate our fault assumptions, formalizing the discussion of Section 2, and present the properties of our system model. The model addresses a set of processes communicating through CAN. Each process is attached to the network through a CAN interface. Together, they form a node. We assume that the processes are fail-silent and blame all temporary failures on the CAN network components. We say a component is **weak-fail-silent** if it behaves correctly or crashes if it does more than a given number of omission failures in an interval of reference, called the component's *omission degree*. The **CAN network components** are modeled by the failure semantics used in [8]:

- individual components are **weak-fail-silent** with *omission degree* f_o ;
- failure bursts never affect more than f_o transmissions in an interval of reference²;
- omission failures may be inconsistent (i.e., not observed by all recipients);
- there is no permanent failure of shared network components, as justified in Section 4.

CAN MAC-level properties

CAN has a medium access control (MAC) sub-layer that basically exhibits the same kind of properties identified in previous works on LANs [10]. A first formalization of CAN MAC-level properties in [8] proved quite effective. Figure 2 complements that definition with the time-related properties MCAN5-MCAN7.

MCAN4 maps the failure semantics introduced earlier onto the operational assumptions of CAN, being $k \geq f_o$. MCAN6 specifies a maximum frame transmission delay, which is T_{td} in the absence of faults. It depends on message latency classes and offered load bounds [9, 12]. The *bounded transmission delay* includes T_{ina} (MCAN5), the maximum duration of an

²For instance the duration of a message transaction round. Note that this assumption is concerned with the total number of failures of possibly different components.

MCAN1 - Broadcast: correct nodes receiving an uncorrupted frame transmission, receive the same frame.

MCAN2 - Error Detection: correct nodes detect any corruption done by the network in a locally received frame.

MCAN3 - Network Order: any two frames received at any two correct nodes, are received in the same order at both nodes.

MCAN4 - Bounded Omission Degree: in a known time interval T_{rd} , omission failures may occur in at most k transmissions.

MCAN5 - Bounded Inaccessibility: in a known time interval T_{rd} , the network may be inaccessible at most i times, with a total duration of at most T_{ina} .

MCAN6 - Bounded Transmission Delay: any frame queued for transmission is transmitted on the network within a bounded delay of $T_{td} + T_{ina}$.

MCAN7 - Tightness: correct nodes receiving an uncorrupted frame transmission, receive it at real time values that differ, at most, by a known small constant $\Delta\Gamma_{tight}$.

Figure 2: CAN MAC-level properties

inaccessibility fault [11]. MCAN7 is crucial for achieving high precision on synchronized clocks [5].

LCAN1 - Validity: if a correct node broadcasts a message, then the message is eventually delivered to a correct node.

LCAN2 - Best-effort Agreement: if a message is delivered to a correct node, then the message is eventually delivered to all correct nodes, if the sender remains correct.

LCAN3 - At-least-once Delivery: any message delivered to a correct node is delivered at least once.

LCAN4 - Non-triviality: any message delivered to a correct node was broadcast by a node.

LCAN5 - Total Order: *not ensured*.

LCAN6 - Bounded Inconsistent Omission Degree: in a known time interval T_{rd} , inconsistent omission failures may occur in at most j transmissions.

Figure 3: Native CAN LLC-level properties

CAN LLC-level properties

CAN has error-recovery mechanisms on top of the basic MAC sub-layer functionality, that yield interesting message properties. These mechanisms provide additional dependability guarantees, in some way with the flavor of the logical link control (LLC) sub-layer in LANs: the omission failures specified by MCAN4 are masked in general at the LLC level by the retry mech-

anism of CAN. However, the existence of inconsistent omissions, as discussed in Section 2, postulates:

- that there may be message duplicates when they are recovered;
- that some j of the k omissions ($j \ll k$) will show at the LLC interface as inconsistent omissions.

Figure 3 recalls from [8] the LLC-level properties of CAN. The first five properties characterize the reliability of CAN communication and its shortcomings. L_{CAN6} provides the grounds for the design of efficient dependability enforcement mechanisms [8, 6].

4 Network Availability

Our ideas to enhance CAN network availability are sketched in Figure 4. They rely on the replication of the physical path – cable medium and transceivers – used by the MAC entities to communicate (channel).

The strategy for channel media replication assumes: each cable replica is routed differently, being reasonable to consider failures in different media as independent; any bit issued from a MAC sub-layer is simultaneously transmitted on all the redundant media interfaces; media stuck-at-dominant failures are handled through a special-purpose *watchdog* timer.

The bare CAN wired-AND nature is exploited for handling channel replicated media: the signals from the different redundant media receivers are combined in an AND function, before interfacing the MAC sub-layer. This simple method ensures resilience against CAN physical partitions:

- nodes at the *in-partition*³ receive a correct signal on all redundant media interfaces;
- in the *out-partition*, the recessive signal from the (idle) failed media is combined with the redundant media signals to produce a correct channel output.

Each network cable can include an additional differential pair, as a way to support a Dual-MAC architecture with small overhead costs⁴. The second MAC can be used to improve: the overall system dependability; the available network bandwidth; hard real-time message delivery guarantees. The architecture of Figure 4 can be made compliant with the standard physical layer, by using spare pins in the CiA connector [1].

5 CAN Inaccessibility

Even in a continuously connected network, the occurrence of certain events (e.g.: bit errors; receiver glitches) in its operation produces a subtle form of

³I.e., the partition that includes the transmitter.

⁴For example, using the state-of-the-art Motorola MPC555 that integrates in a single chip a 32-bit PowerPC micro-controller and two advanced CAN 2.0B controllers.

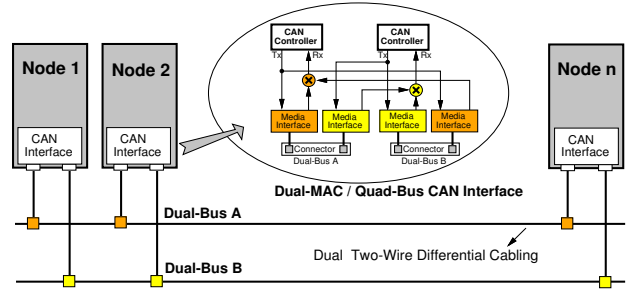


Figure 4: Highly-available CAN infrastructure

virtual partitioning. CAN has its own means of recovering from these situations, but this takes time [7, 11]. An **inaccessibility** fault occurs when a component *temporarily refrains* from providing service.

The effect of inaccessibility on real-time communication is the error it introduces in timing bounds, such as message latencies. Most message schedulability analyses consider the network as always functioning normally [9, 2]. Bounds are established that may be violated upon the occurrence of inaccessibility events. In consequence, the system (e.g. [2]) may exhibit an unpredictable behavior and ultimately fail.

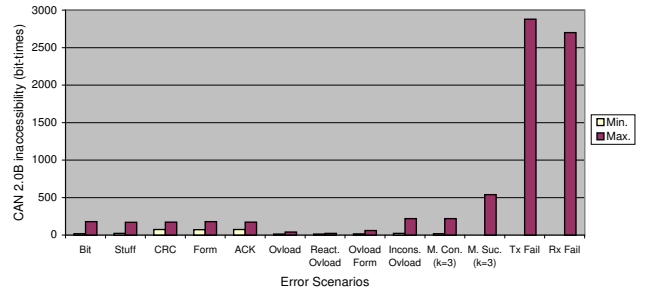


Figure 5: Normalized CAN inaccessibility

One solution to avoid failures due to inaccessibility events is: to compute the duration of all possible inaccessibility faults, as done in [7, 11] (Figure 5 summarizes the results); accommodate inaccessibility bounds in the timeliness model (e.g. M_{CAN6}).

6 Failure Detection and Membership

A membership service is intended to provide, at any given time, consistent information about failed/correct nodes. Our approach to this problem is based on the observation that many CAN applications [9, 2] exhibit a periodic traffic pattern. A failure detection/membership service matching strict application-level latency requirements can be designed with minimal costs in network bandwidth.

The periodic high-level messages are implicitly used

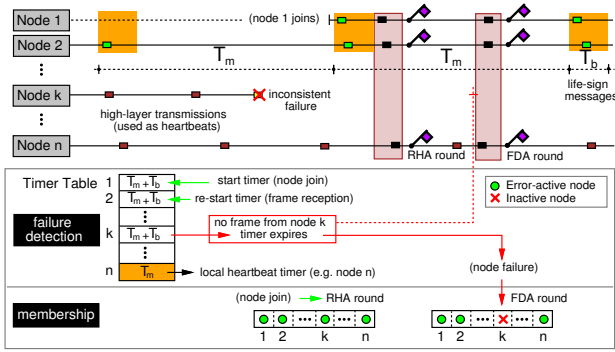


Figure 6: Failure detection and membership

as heartbeats (Figure 6). Specific *life-sign* messages need only to be issued by nodes with message periods higher than the failure detection latency or transmitting only sporadic/aperiodic traffic. If a node remains silent during a period longer than the detection latency, that will be a failure. Consistency of membership information is ensured through: a *reception history agreement* (RHA) protocol, in the presence of node join/leave operations; an optimized *failure detection agreement* (FDA) protocol, upon node failure.

Optimization tradeoffs exist with a protocol variant where all nodes explicitly issue *life-sign* messages.

7 Reliable Group Communication

The problem of reliable CAN communication was originally addressed in [8], dismissing the current belief that CAN supports an atomic broadcast service and providing a protocol suite that handles the problem effectively. Next, we use and extend the results from [8] to outline the current architecture of a CAN-based reliable group communication service.

Interfacing the standard CAN layer (Figure 7) we use the fundamental fault-tolerant broadcast protocols of [8, 6]: SDCAN enhances LKAN3, by ensuring that each message is delivered **at-most-once**, if no message ordering is required [6]; EDCAN enhances LKAN3 in the same way but also enhances LKAN2, removing the condition of the sender not failing (cf. Fig 3) and securing all the properties of a reliable broadcast service [8, 6].




Figure 7: CAN real-time fault-tolerant protocol suite

A versatile real-time group communication service, offering different *qualities of service*, is defined above this layer. The *G-Filtering* sub-layer restricts processing of higher layers to the traffic addressed to the node. The top sub-layer include (totally ordered) atomic (T-CAN) and reliable (R-CAN) group communication protocols, which are variants of the protocols in [8], and two new protocols: L-CAN, a reliable group communication protocol that trades a high message delivery bound with a low utilization of network bandwidth; F-CAN, a companion protocol that exploits MCAN3 and LKAN2 to support an efficient message fragmentation scheme that does not need to use sequence numbers for fragment ordering.

The failure detection and membership protocols are also included in Figure 7. A companion clock synchronization protocol is described in [5].

8 Concluding Remarks

There is an increasing demand for fault-tolerant and real-time distributed systems based on fieldbuses. In this paper, after formalizing the properties actually secured by CAN, we have shown that with the appropriate techniques one can draw a modular solution able to add fault-tolerance and hard real-time attributes to the basic functionality offered by CAN off-the-shelf components.

References

- [1] CiA - CAN in Automation. *CAN Physical Layer for Industrial Applications - CiA/DS102-1*, October 1992.
- [2] J. Gil, A. Pont, G. Benet, J. Blanes, and M. Martinez. A CAN architecture for an intelligent mobile robot. In *Proceedings of the IFAC International Symposium on Intelligent in Components and Instrumentation for Control Applications*, Anney, France, 1997. IFAC.
- [3] ISO. *International Standard 11898 - Road vehicles - Interchange of digital information - Controller Area Network for high-speed communication*, November 1993.
- [4] H. Kopetz and G. Grunsteidl. TTP - a protocol for fault-tolerant real-time systems. *IEEE Computer*, 27(1):14-23, January 1994.
- [5] L. Rodrigues, M. Guimarães, and J. Rufino. Fault-tolerant clock synchronization in CAN. In *Proceedings of the 19th Real-Time Systems Symposium*, Madrid, Spain, December 1998. IEEE.
- [6] J. Rufino, N. Pedrosa, J. Monteiro, P. Veríssimo, and G. Arroz. Hardware support for CAN fault-tolerant communication. In *Proceedings of the 5th International Conference on Electronics, Circuits and Systems*, pages 263-266, Lisbon, Portugal, September 1998. IEEE.
- [7] J. Rufino and P. Veríssimo. A study on the inaccessibility characteristics of the Controller Area Network. In *Proceedings of the 2nd International CAN Conference*, pages 7.12-7.21, London, England, October 1995. CiA.
- [8] J. Rufino, P. Veríssimo, G. Arroz, C. Almeida, and L. Rodrigues. Fault-tolerant broadcasts in CAN. In *Digest of Papers, The 28th International Symposium on Fault-Tolerant Computing Systems*, pages 150-159, Munich, Germany, June 1998. IEEE.
- [9] K. Tindell and A. Burns. Guaranteeing message latencies on Controller Area Network. In *Proceedings of the 1st International CAN Conference*, pages 1.2-1.11, Mainz, Germany, September 1994. CiA.
- [10] P. Veríssimo. Real-time Communication. In S.J. Mullender, editor, *Distributed Systems*, ACM-Press, chapter 17, pages 447-490. Addison-Wesley, 2nd edition, 1993.
- [11] P. Veríssimo, J. Rufino, and L. Ming. How hard is hard real-time communication on field-buses? In *Digest of Papers, The 27th International Symposium on Fault-Tolerant Computing Systems*, Washington - USA, June 1997. IEEE.
- [12] K. Zuberi and K. Shin. Scheduling messages on Controller Area Network for real-time CIM applications. *IEEE Transactions on Robotics and Automation*, 13(2):310-314, April 1997.