

An Event Model for Real-Time Systems in Mobile Environments

René Meier¹, Jörg Kaiser², Barbara Hughes¹, Cristiano Brudna², Vinny Cahill¹

¹Department of Computer Science, Trinity College Dublin, Ireland

²Department of Computer Structures, University of Ulm, Germany

{rene.meier, barbara.hughes, vinny.cahill}@cs.tcd.ie, {kaiser, cristiano.brudna}@informatik.uni-ulm.de

Abstract

This paper describes an event model that has been designed to address the predictability requirements of applications operating in mobile environments based on hierarchically structured WAN-of-CANs networks. The event model supports an event channel concept for modeling the guarantees provided by the underlying heterogeneous communication infrastructure. The networks that comprise such a WAN-of-CANs may provide fundamentally different degrees of quality of service and as a result can be viewed as zones within which certain guarantees can be enforced. Event channels operating in CAN-based subnetworks typically with strong timing behavior may support hard temporal and reliability attributes whereas channels interconnecting these subnetworks using wireless networks support weaker timing attributes.

1. Introduction

Advances in information technology are encouraging new classes of applications that are based on a large number of networked components acting autonomously in response to a myriad of sensors and actuators to assess and control aspects of the environment. Examples include telematics, traffic management and home automation to name a few. To a large extent, such systems operate proactively and independently of direct human control driven by their perception of the environment and the ability to organize computations dynamically. The challenging characteristics of these applications include sentience and autonomy of components, issues of responsiveness and safety criticality, geographical dispersion, mobility and evolution.

Such systems require a degree of co-operation, adaptability, extensibility and reliability that is not available today. The problem is to provide a communication and interaction scheme that supports the large-scale many-to-many communication relation typical of these applications. Additionally, it should be possible to seamlessly disseminate relevant information generated by deeply embedded controllers to all interested entities in the global network.

The contribution of the paper is an event model for real-time systems in heterogeneous environments. Such environments may be composed from networks with largely different quality characteristics ranging from highly predictable real-time buses to ad-hoc wireless networks in which mobility of nodes may result in frequent link failures. The event-based communication model represents a paradigm for middleware that asynchronously interconnects

the components that comprise distributed applications [1]. Event-based middleware has been recognized as being well suited to addressing the requirements of mobile application [2, 3] and to connecting the components in control applications with timeliness requirements [4-6]. However, to date event models have not been designed to offer sufficient levels of predictability for real-time systems while supporting mobile application components.

The event model presented in this paper addresses the predictability requirements of large-scale applications operating in mobile environments based on heterogeneous communication infrastructures. Such infrastructures may include Controller Area Networks (CAN), Local Area Networks (LAN), and Wide Area Networks (WAN), especially those based on wireless technology, and typically consists of a hierarchically structured WAN-of-CANs internetwork. A WAN-of-CANs internetwork typically comprises CAN subnetworks that are interconnected by means of (wireless) LANs and WANs. The networks that comprise such a WAN-of-CANs characteristically provide fundamentally different degrees of Quality of Service (QoS), ranging from CAN networks with strong timing and reliability behavior to wireless networks supporting mobile application components and as a result providing weaker guarantees. We regard individual networks as QoS containers, called zones, within which given degrees of predictability in terms of timing and reliability can be enforced.

Our event model provides a programming model based on the concept of event channels. A number of event channel classes with different temporal and reliability attributes are supported to integrate real-time aspects into the event channel model. Depending on the guarantees available from the underlying network, these event channels can be mapped to certain QoS zones. Generally, channels can be associated with zones providing equal or stronger guarantees.

The remainder of this paper is structured as follows: Section 2 presents the system architecture including the underlying network infrastructure and the classes of supported event channels and discusses issues arising when applying real-time to a WAN-of-CANs network, especially those including wireless technology. Section 3 describes the programming model supported by our event model. Section 4 outlines the communications architecture. Section 5 concludes this paper by summarising our work.

2. System Architecture

When striving for a real-time event system, the communications architecture substantially affects the temporal and reliability properties of event dissemination. Thus, in quality terms the underlying network is not transparent for the application using the event system. The goal for our event system is to express the quality requirements of event dissemination on the abstraction level of events rather than to resort to lower-level network details. The middleware then automatically maps these requirements to the underlying network. Hence, we will briefly introduce the general assumptions about the network architecture.

2.1 General Network Structure

From an application point of view, the event system should support cooperating autonomous entities, such as cars, robots, and people wearing computers and moving through a smart environment. As an example consider autonomous mobile robots which may be equipped with all kinds of smart sensors, such as lasers, radars, cameras, navigational sensors, and chemical sensors, to achieve adequate perception of their environment. The necessary highly reactive behaviour of such an individual robot has to be achieved by a tight cooperation of the sensor / actuator system over some special purpose network inside the robot, called a CAN (Controller Area Network in a broad sense). Additionally, the vehicles may communicate via a wireless link carrying out joint tasks. At an abstract layer, this is modelled as islands of tight control cooperating via a WAN-of-CANs structure [7]. We assume a certain guaranteed level of predictability as an intrinsic property of a CAN. A CAN therefore constitutes a zone of coherent QoS provision.

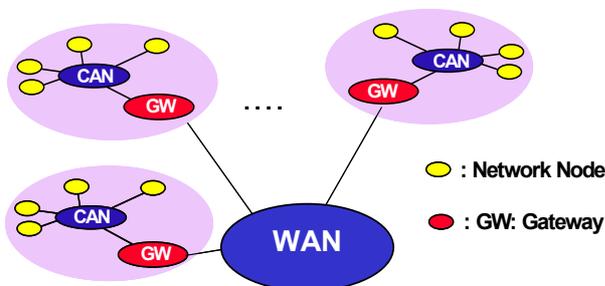


Figure 1. The WAN-of-CANs network architecture.

The general structure depicted in Figure 1 is recursive in a sense that a CAN may be composed from a hierarchy of networks itself. This issue is not further discussed in this paper. It should only be noted that each of the interconnected subnets constitutes a zone in which a certain QoS is defined. The gateways have multiple functions. Firstly, they route events over network boundaries by maintaining an event interface that is transparent with respect to functional properties throughout the network. For non-functional quality attributes they provide the needed awareness about the network properties. Thus, they have to marshal events depending on the underlying abstract network infrastructure,

which is described below. Secondly, they constitute filters that allow to scope events according to their context and their quality attributes. A realisation of this structure for cooperating mobile robots internally equipped with a CAN-Bus, which is extended via gateways to a wireless LAN for event dissemination between the robots is presented in [8]. Furthermore, a middleware service that has been designed for interconnecting mobile application components using ad hoc wireless local area networks has been described in [9].

2.2 Applying Timeliness to WAN-of-CANs

Previous research in real-time event-based communication [4, 6] has focused on wired networks. The assumptions made are not applicable in WAN-of-CAN networks comprising wireless ad hoc networks where dynamic connections, unpredictable latency and limited resource availability, impact the real-time communication guarantees.

2.2.1 Dynamic Connections

The absence of a fixed infrastructure means that mobile nodes constitute the communication infrastructure. As nodes move in and out of range of other nodes, the connectivity and network topology changes dynamically.

Communication between mobile nodes requires the Received Signal Strength (RSS) to be adequate to connect to another mobile node. The RSS is significantly affected by the movement of wireless nodes, terrain configuration and transmission power which may depend on battery life [10]. The changes in RSS lead to highly unpredictable connections between mobile nodes.

Unlike fixed infrastructure networks where link failures are comparatively rare events, the rate of link failure due to node mobility is the primary obstacle to routing in ad hoc networks [11]. Greater mobility increases the fluctuations in link connectivity, the volume of topological updates, the time spent processing the updates (e.g. for route discovery protocols), and congestion due to increased update transmissions and retransmissions. Link failures may also result in network partitions and the possibly critical situation that real-time communication is not possible with all relevant wireless nodes.

Topology changes introduced by node mobility and wireless link failures must somehow be communicated to other nodes. Since wireless and computation resources, for example bandwidth and battery power, are limited in wireless ad hoc networks, any overhead must be kept to a minimum and additional communication delays due to an increase in the volume of topological updates must be avoided.

Topology updates throughout an ad hoc network cannot happen instantaneously. The time-varying capacity of wireless links, limited resources and node mobility make maintaining accurate routing information very difficult in ad hoc wireless networks. Routing for real-time event-based communication must ensure resource availability (e.g. bandwidth) whilst maintaining minimum latency. Routing decisions may be compromised by inaccurate network information and time-bounded route determination, where

optimal routes may not be found within the time available that may be critical for real-time event-based communication.

2.2.2 Unpredictable Latency

Wireless transmissions in ad hoc networks are broadcast through a shared physical communication channel. Collisions in wireless communications can be caused by simultaneous transmissions by two or more wireless nodes sharing the same frequency band, or the hidden terminal problem [12]. Collisions cause unpredictable latency for medium access that is unacceptable in real-time event-based communication.

The lack of a fixed infrastructure and the limited power of wireless mobile nodes reducing the transmission range, means that wireless nodes are designed to serve as routers if needed. The result is a distributed multi-hop network with a time-varying topology with typically short-lived routes. The unpredictable latency for route determination and medium access (encountered at each hop) makes an estimate of end-to-end delivery latency that is critical in real-time event-based communication, very difficult and with a high probability of being inaccurate.

2.2.3 Limited Resource Availability

In mobile ad hoc wireless networks the available bandwidth is very limited and some wireless devices have severe energy constraints. Hence, communication is an expensive operation and any additional control packet overhead (e.g. resource reservation) must be kept to a minimum. In [10], Campbell et al. identify that a critical design issue for wireless ad hoc networks is the development of suitable communication architectures, protocols and services that efficiently reduce power consumption thereby increasing the operational lifetime of the wireless device. Increasing battery life will help reduce the number of link breaks caused by node failure, and thus the volume of topological updates, the competition for scarce resources, and communication unpredictability that impede real-time event-based communication.

3. The Programming Model

As a consequence of the analysis, a system incorporating wired and wireless networks has to cope with heterogeneity in many respects: dissemination mechanisms, quality of service, topology of networks etc. Thus, a programming model should strive to provide abstractions that allow to express the communication requirements on an adequate level. Moreover, as a consequence of the interaction with the physical world, the pace of interactions and of the computational progress in the distributed system is dictated by the progression of real time and the properties of the environment. To support coordination of actions and cooperation between mobile entities, not only the events must be disseminated over a communication network in a timely manner, but also the direct interaction, perceived by sensors has to be considered. Perception and fusion of elementary events like sensor data in the sentient object model [13] become part of the overall interaction. An architecture which describes this model of generic events in

the sentient object model has been presented by Verissimo and Casimiro [14] in the context of CORTEX project [7]. In this model, objects only interact via an event layer. The event layer hides the network as well as the transformation process of real-world events. The basic objective behind this architecture is to reason in terms of generic events. Because all events appear at the event layer, it is possible to order the events consistently whether they are generated in the environment or in the system. The event and event channel model introduced here constitute a specific way to realize such an architecture. An event is an instance of an event type, which is characterized by subject, attributes, and content:

event = <subject, attribute_list, content>

The subject is attached to an event as a tag to explicitly identify the event's type and enable basic filtering and some form of type checking on event dissemination. An event is further characterized by a set of functional and non-functional attributes. Functional attributes are related to the application relevant context in which an event has been generated and may include information like location, range of dissemination (relative location or proximity) and time of occurrence, or system related information like a mode of operation or a network zone in a WAN-of-CANs architecture. Non-functional attributes represent intrinsic properties and quality aspects of an event. These may be deadlines necessary to express temporal requirements of event dissemination or a validity interval according to an aging function. This is particularly important in a real time event system to prevent outdated events competing for resources with current events. Finally, the content of an event carries the data that is represented as a structured set of functional parameters. The content is accessible by specific methods.

Predictability incorporates the timely delivery of events under anticipated load and failure conditions. The nature of events may range from a safety-critical event signaling that a crash of a mobile vehicle with an obstacle is about to happen to the dissemination of non-critical events describing a room temperature or illumination intensity that have less demanding predictability requirements. As a matter of fact, there is a trade-off between the resources expended and the degree of predictability. Consequently, the respective requirements for the underlying abstract network may range from synchronous reliable broadcasts [8, 15] to best effort communication.

To facilitate the specification of the particular requirements for event dissemination the notion of event channels is introduced. An event channel encapsulates properties of the underlying communication system and allows quality attributes to be specified on an abstraction level where it is assessable to an application programmer. The benefits are twofold: Firstly, a check can be performed as to whether the non-functional attributes of an event match to the quality attributes defined for the event channel. This allows early timing failure detection. Secondly, predictability requires that resources are available when they are needed. An event channel can be established and the necessary resources can be assigned by the middleware before an event

has to be disseminated. The general form of an event channel representation is given by:

```
event_channel = <subject, attribute_list, handlers>
```

In contrast to the attributes of an event, which describes the properties of a single individual occurrence of an event, the attributes of the event channel abstract the properties of the underlying communication network and dissemination scheme. Therefore the attributes define quality properties and include e.g. latency, dissemination constraints and reliability parameters. The subject of the event channel must match the subject of the event that is disseminated through the channel. The "handlers"-field allows specifying notification and exception handlers for the event channel.

4. The Communications Architecture

The design of the communications architecture is motivated by our approach of providing a programming model based on the concept of event channels with different timeliness and reliability properties. The communications architecture employs a mechanism for enforcing hard real-time event channel properties in CAN networks and another mechanism based on bounding the propagation range of events for providing soft real-time guarantees with high probability in wireless networks.

4.1 Mapping Event Channels to Zones

Applications may comprise numerous components representing real-world objects that may be mobile and distributed over a large geographical area. Such components are typically location-aware and depending on their location may interact using different parts of the underlying WAN-of-CANs network architecture. Applications may define various event channels disseminating events with different temporal and reliability properties. Depending on the guarantees available from the underlying network, these event channels can be mapped to certain QoS zones. Generally, event channels can be associated with zones providing equal or stronger guarantees. For example, a traffic management application may include vehicles interacting through wireless networks in order to exchange information on the current traffic situation thereby contributing to better driver awareness and consequently to safer driving. Such information may include an accident notification disseminated by a broken-down car to approaching vehicles. Various components representing intra vehicle objects, such as breaks, accelerator, and speedometer, and lights, might communicate using a CAN-based network. Such inter-vehicle and intra-vehicle communication may use different event channel classes for interconnecting their respective components. Hard real-time event channels may be mapped to a zone incorporating intra vehicle components whereas event channels with weaker delivery guarantees may be mapped to a zone comprising components using wireless communication.

Multiple event channels may be mapped to a particular zone sharing the available resources. Such event channels typically disseminate events describing different information and may support different properties. For example, intra-

vehicle communication may include a hard real-time event channel disseminating events on behalf of breaks and accelerators and a non real-time event channel controlling the lights of the vehicle.

Event channels may connect components across multiple parts of the network architecture and as a result may be mapped to multiple zones. The properties that can be enforced by such an event channel depend on the level of QoS provided by each of the zones involved. An event channel can be associated with multiple zones if every zone involved provides equal or stronger guarantees. Moreover, an event channel may only operate across the boundaries of multiple zones if the underlying networks are connected through designated gateway components. Such gateways act as producers and consumers of events on either side of the networks they connect. Events received on one side are disseminated on the other side and vice versa. In addition, gateways allow applications to specify network specific mapping of event data and attributes and handle implicit attributes, such as location. A gateway may use a location service to retrieve its own location and may append this location information to events generated by nearby nodes lacking direct access to a location service. For example, a gateway located in a vehicle may connect a CAN-based intra-vehicle network and a wireless inter-vehicle network. It may attach its own location to events generated by nodes on the CAN network as these reside close to the gateway and maintain their location relative to the gateway.

4.2 Mapping Channels to CAN Networks

In zones, covering special CAN networks like an intra-vehicle network, better real-time support may be necessary because a highly reactive and predictable behaviour is needed. We exploited the specific hardware supported, lower level mechanisms of a CAN-Bus, which is popular in the automotive industry, to implement various event channel classes. Reliable hard real-time event channels are provided for time and safety critical events. They are based on a TDMA scheme in which channels are established in pre-reserved time slots. Hard real-time event channels can tolerate a specified number of omission failures and guarantee the delivery of events at the event deadline. Soft real-time event channels use the temporal validity of events to define deadlines for event transmission. However, in transient overload situations, these deadlines may be missed and temporal validity may expire. The event system will only provide awareness of such expiration to enable application dependent reactions. Finally, non real-time event channels have no temporal constraints. For a detailed description of the real-time channel classes the reader is referred to [16].

4.3 Mapping Channels to Wireless Networks

As previously discussed, enforcing hard real-time guarantees in wireless environments in general and in ad hoc networks in particular is problematical due to the dynamic nature of these networks. However, we argue that soft real-time event channels can be supported in zones that contain wireless networks. Significantly, we envisage using a set of

techniques in order to provide soft real-time guarantees with a high probability.

We propose to bound the propagation range of events in wireless environments [9] by defining event channel attributes that describe geographical areas within which events are valid. Such proximity-based event dissemination represents a natural way to limit the scope of an event channel, thereby allowing entities to interact based on their current location. An example scenario illustrating such behavior might include a broken-down car disseminating an accident notification to vehicles in its vicinity.

Moreover, we envisaged using a predictable medium access protocol, such as the light-weight, location-aware, atomic multicast protocol for Time-Bounded Medium Access Control (TBMAC) [17], when propagating event notifications. The TBMAC protocol is based on time-division multiple access with dynamic but predictable slot allocation and has been designed for use in multi-hop ad hoc networks. It provides, with high probability, time-bounded access to the wireless medium that can be exploited by event channels with guaranteed response time requirements.

5. Conclusions

We have presented an approach to incorporating the topology of a heterogeneous communication infrastructure into an event-based programming model. We have described how an event model may address the predictability requirements of applications operating in mobile environments based on hierarchically structured WAN-of-CANs networks. The subnetworks that comprise such a WAN-of-CANs define QoS containers that can be viewed as zones within which certain guarantees can be enforced. Our programming model is based on a concept of event channels supporting a number of event channel classes with different temporal and reliability attributes. Depending on the guarantees available from the underlying network, these event channels can be mapped to certain QoS zones. Generally, event channels can be associated with zones providing equal or stronger guarantees.

We have introduced two prototype implementations of our event model. The CAN version enforces hard-real time guarantees in CAN-based subnetworks using lower level mechanisms of the CAN-Bus. The LAN version uses techniques, including geographical bounding of the event propagation range and predictable medium access, to provide soft real-time guarantees with a high probability in wireless networks.

Acknowledgments

The work described in this paper was partly supported by the Irish Higher Education Authority's Programme for Research in Third Level Institutions cycle 0 (1998-2001) and by the FET programme of the Commission of the EU under research contract IST-2000-26031 (CORTEX).

References

[1] J. Bacon, K. Moody, J. Bates, R. Hayton, C. Ma, A. McNeil, O. Seidel, and M. Spiteri, "Generic

Support for Distributed Applications," *IEEE Computer*, vol. 33, pp. 68-76, 2000.

- [2] G. Cugola, E. D. Nitto, and A. Fuggetta, "The JEDI Event-Based Infrastructure and its Application to the Development of the OPSS WFMS," *IEEE Transactions on Software Engineering (TSE)*, vol. 27, pp. 827-850, 2001.
- [3] R. Meier, "Communication Paradigms for Mobile Computing," *ACM SIGMOBILE Mobile Computing and Communications Review (MC2R)*, vol. 6, pp. 56-58, 2002.
- [4] T. Harrison, D. Levine, and D. Schmidt, "The Design and Performance of a Real-Time CORBA Event Service," in *Proceedings of the 1997 Conference on Object-Oriented Programming Systems, Languages and Applications (OOPSLA'97)*. Atlanta, Georgia, USA: ACM Press, 1997, pp. 184-200.
- [5] R. Rajkumar, M. Gagliardi, and L. Sha, "The Real-Time Publisher/Subscriber Inter-Process Communication Model for Distributed Real-Time Systems: Design and Implementation," in *Proceedings of the IEEE Real-time Technology and Applications Symposium*, 1995.
- [6] J. Kaiser and M. Mock, "Implementing the Real-Time Publisher/Subscriber Model on the Controller Area Network (CAN)," in *Proceedings of the 2nd International Symposium on Object-oriented Real-time distributed Computing (ISORC99)*. Saint-Malo, France, 1999.
- [7] P. Verissimo, V. Cahill, A. Casimiro, K. Cheverst, A. Friday, and J. Kaiser, "CORTEX: Towards Supporting Autonomous and Cooperating Sentient Entities," in *Proceedings of the European Wireless Conference*. Florence, Italy, 2002.
- [8] J. Kaiser and C. Brudna, "A Publisher/Subscriber Architecture Supporting Interoperability of the CAN-bus and the Internet," in *Proceedings of the IEEE International Workshop on Factory Communication Systems (WFCS 2002)*. Västerås, Sweden, 2002.
- [9] R. Meier and V. Cahill, "Exploiting Proximity in Event-Based Middleware for Collaborative Mobile Applications," in *Proceedings of the 4th IFIP International Conference on Distributed Applications and Interoperable Systems (DAIS'03)*, LNCS 2893. Paris, France: Springer-Verlag Heidelberg, Germany, 2003, pp. 285-296.
- [10] J. Gomez, A. T. Campbell, M. Naghshineh, and C. Bisdikian, "Conserving Transmission Power in Wireless Ad Hoc Networks," in *Proceedings of the 9th International Conference on Network*

Protocols (ICNP 2001). Riverside, California, USA, 2001.

- [11] A. B. McDonald and T. Znati, "A Mobility Based Framework for Adaptive Clustering in Wireless Ad-Hoc Networks," *IEEE Journal in Selected Areas in Communications (JSAC)*, vol. 17, pp. 1466-1487, 1999.
- [12] F. A. Tobagi and L. Kleinrock, "Packet Switching in Radio Channels: Part II - The Hidden Terminal Problem in Carrier Sense Multiple-Access and the Busy-Tone Solution," *IEEE Transactions on Communications*, vol. 23, pp. 1417-1433, 1975.
- [13] A. Fitzpatrick, G. Biegel, S. Clarke, and V. Cahill, "Towards a Sentient Object Model," presented at Workshop on Engineering Context-Aware Object Oriented Systems and Environments (OOPSLA/ECOOSE'02), Seattle, Washington, USA, 2002.
- [14] P. Verissimo and A. Casimiro, "GEAR: The Generic-Events Architecture," presented at WORDS 2003, Cancun, Mexico, 2003.
- [15] F. Cristian, "Synchronous Atomic Broadcast for Redundant Broadcast Channels," *The Journal of Real-Time Systems*, vol. 2, pp. 195-212, 1990.
- [16] J. Kaiser, C. Brudna, and C. Mitidieri, "A Real-Time Event Channel Model for the CAN Bus," in *Proceedings of the Eleventh International Workshop on Parallel and Distributed Real-Time Systems (WPDRTS 2003)*. Nice, France, 2003.
- [17] R. Cunningham and V. Cahill, "Time Bounded Medium Access Control for Ad Hoc Networks," in *Proceedings of the Second ACM International Workshop on Principles of Mobile Computing (POMC'02)*. Toulouse, France: ACM Press, 2002, pp. 1-8.