Coordination of Safety-Critical Mobile Real-Time Embedded Systems

Aline Senart, Mélanie Bouroche, Barbara Hughes and Vinny Cahill
Distributed Systems Group
Department of Computer Science
Trinity College Dublin
{first.last}@cs.tcd.ie

Abstract

Safety-critical mobile applications running on resource-constrained embedded systems will play an increasingly important role in domains such as automotive systems, space, robotics and avionics. Such applications are composed of mobile autonomous components interacting spontaneously without any infrastructure. Therefore, to undertake safety-critical missions, the applications will have to coordinate the behaviour of their components in real-time, while overcoming the limitations of (ad hoc) wireless communication. In this paper, we outline the key research challenges to be addressed in order to achieve such real-time coordination.

1 Introduction

Mobile computing has received much interest in recent years due to the proliferation and diversification of mobile computing devices and significant advances in wireless networking capabilities [MCE04]. Following this trend, we can now envisage the widespread development and deployment of mobile safety-critical applications in daily life becoming a reality. Early examples include automated guided vehicles or distributed process control in wastewater treatment plans [KS03] and encompass a variety of domains such as automotive systems, space, robotics and avionics.

Mobile safety-critical applications enable a collection of autonomous components equipped with wireless communication facilities to interact in a possibly ad hoc wireless network. These components are typically deployed within embedded systems as control systems for critical infrastructure or as mission-critical systems exhibiting autonomous and proactive behaviour. Their failure or malfunction may result in serious injury to people and loss or severe damage to equipment. Among them, hard real-time safety-critical components are those that suffer a critical failure if time constraints are violated [Kop97]. They are traditionally found interacting with each other in a distributed shared environment where they have to complete cooperative tasks in a specified time-period. Therefore, they need to coordinate their behaviour and adapt it to the state of their environment within a time-bound interval. For example, robots used in urban search and rescue cooperate together and with humans in overlapping workspaces. For this working environment to remain safe and secure, not only must internal computations of robots meet their deadlines, but timely coordination of robots behaviour is also required.

In this paper, we explain why real-time coordination of mobile embedded systems is a key requirement to be addressed in order to ensure application-wide safety constraints. We show that there are several impediments to achieving such coordination in (ad hoc) wireless networks. First, timely communication is hard to provide in a network where the topology changes dynamically and the rate of link failures is high. Secondly, the global behaviour of coordinated real-time mobile embedded systems has to guarantee that the safety constraints of the application are always respected.

The structure of the paper is as follows. In the next section, we discuss the absence of adequate real-time (ad hoc) wireless communication models. Following this, we demonstrate why the unavailability of real-time coordination models is an obstacle to achieving coordination.

2 Research challenges

Ensuring that safety constraints are never violated by mobile real-time embedded systems that coordinate their behaviour with each other is a particularly hard problem. In this section, we present the main challenges that are raised.

2.1 Absence of adequate real-time (ad hoc) wireless communication models

In order to achieve coordination of safety-critical mobile real-time embedded systems, there is a strong need for real-time (ad hoc) wireless communication. This is very difficult to provide in ad hoc wireless networks since nodes communicate directly with one another in a peer-to-peer fashion. Furthermore, dynamic node mobility, limited resources of embedded systems in terms of power and transmission range, and the varying and unpredictable latency of wireless links may potentially result in critical situations where a hard real-time event cannot be delivered to all relevant nodes [PH02]. Consequently, predictive techniques are required in order to find new routes prior to the failure of existing ones and proactively allocate the required resources [PH99]. In this context, some of the questions that arise are: Can real-time wireless communication be improved by previous knowledge of the periodicity of transmissions? Would partition anticipation help to achieve proactive and preemptive real-time communication? Is mobility awareness useful? Can mobility be predicted? If mobility can be predicted, would mobility patterns be of any use? Embedded systems are resource-constrained and cannot tolerate excessive resource reservation. Then, how to limit the overhead of resource reservation? May the reserved resources be available for use by other components until required by the original one?

In the safety-critical mobile embedded systems that we are considering, providing guarantees about timely communication without any given assumption is clearly not possible [HC03]. It would be interesting to investigate if a real-time (ad hoc) wireless communication model can be designed to provide such guarantees by imposing some additional constraints on connectivity, topology, number of nodes of such networks or characteristics that applications have to expose, for example. Furthermore, many guarantees can be ensured but the most crucial one are very difficult to provide. For example, when there is an alteration in the underlying physical infrastructure, like a network partition or a significant decrease in achievable latency, communicating components have to be informed of the change in order to adapt their behaviour to respect safety and timeliness constraints. This is a challenging requirement, particularly in ad hoc wireless networks, where dynamic changes in the connectivity and network topology are frequent. Many issues are raised: What form of feedback can we give to safety-critical mobile applications? How can this feedback be provided in real-time? Can we guarantee that this feedback will always be returned to applications?

2.2 Unavailability of real-time coordination models

Actions of autonomous mobile real-time components need to be tightly-coupled in order to ensure that the safety constraints of the emerging system are satisfied. As they operate in a environment characterised by extensive resource sharing (e.g., sensors/actuators, processors and communication channels), their behaviour must be coordinated to accommodate system-wide safety constraints. These safety constraints are typically application-specific and place stringent constraints on the state and behaviour of individual components. Determining constraints on individual components from high-level system-wide constraints is a hard task. To overcome this difficult problem, we need a real-time coordination model that would automatically translate system-wide constraints into constraints on the behaviour of individual components to ensure that no catastrophe can occur, i.e., that specified safety constraints are never violated. Regarding this coordination model, several questions can be raised: How can we specify real-time safety constraints? Can we provide a formalism powerful enough to fully capture any existing an future constraint? Is it possible to fully automate the translation of application constraints into constraints on the behaviour of individual components? If not, can we provide a framework that would make the development of mission-critical application easy? How can we verify the behaviour of a set of communicating components?

Furthermore, in large-scale and complex safety-critical systems composed of collections of embedded systems, each autonomous component can change its own state asynchronously by reaction to events that are relevant, arising from its local environment. Keeping distributed views and actions timely and consistent is a hard problem, if not impossible. Consensus-based techniques assume continuous connectivity and use of atomic operations that wait for every component to be ready. However, this may not be viable for mobile real-time embedded systems, where the link failure rate is significant and tasks have to be performed in a timely manner. This leads to the following questions: Is hard consistency in such environments possible? Can we reduce the complexity of the problem by adding constraints on the components? How to resynchronise outdated views and states of components in a bounded interval of time? How to ensure system-wide safety constraints without a distributed consistent and accurate state of the safety-critical application?

3 Conclusion

In this paper, we have described the need for real-time coordination, a major requirement of safety-critical mobile real-time embedded systems. Then, we have identified the challenges that are still open to achieve it. We believe that the definition of appropriate (ad hoc) wireless communication models and real-time coordination models is a key requirement to allow such coordination in (ad hoc) wireless networks. Our current research is addressing this issue in the context of the Space-Elastic Model [BHM$^+$05]. This communication model gives feedback to safety-critical applications in the form of a description of a varying geographical area in which real-time communication is guaranteed for specified Quality of Service requirements. We are currently investigating how this feedback can be used to provide timeliness and reliability guarantees and to define a coordination model which is able to ensure system-wide safety constraints [BHC06].

References

[BHC06] Mélanie Bouroche, Barbara Hughes, and Vinny Cahill. Building reliable mobile applications with space-elastic adaptation. In *4th International Workshop on Mobile Distributed Computing (MDC 2006)*, June 2006. to appear.

[BHM$^+$05] M. Boulkenafed, B. Hughes, R. Meier, G. Biegel, and V. Cahill. Providing hard real-time guarantees in context-aware applications: Challenges and requirements. In *4th IEEE International Symposium on Network Computing and Applications*, pages 119–127, Cambridge, MA, 2005.

[HC03] B. Hughes and V. Cahill. Achieving real-time guarantees in mobile wireless ad hoc networks. In *Real-Time Systems Symposium (RTSS'03) - Work In Progress Session*, pages 37–40, Cancun, Mexico, December 2003.

[Kop97] Hermann Kopetz. *Real-Time Systems: Design Principles for Distributed Embedded Applications.* Kluwer Academic Publisher, 1997.

[KS03] J. Kjeldskov and J. Stage. The process of developing a mobile device for communication in a safety-critical domain. In *Ninth IFIP TC13 International Conference on Human-Computer Interaction*, Zürich, Switzerland, September 2003.

[MCE04] C. Mascolo, L. Capra, and W. Emmerich. *Principles of Mobile Computing Middleware*, chapter 11, pages 261–280. John Wiley, middleware for communications edition, 2004.

[PH99] M. R. Pearlman and Z. J. Haas. Determining the optimal configuration for the zone routing protocol. *IEEE Journal on Selected Areas in Communication*, 17(8):1395–1414, 1999.

[PH02] D. D. Perkins and H. D. Hughes. A survey on qos support for mobile ad hoc networks. *Wireless Communications and Mobile Computing*, 2:503–513, 2002.