Design of Autonomous, Distributed Systems*

Tunc Simsek and Pravin Varaiya

University of California, Berkeley, CA 94720, USA
{tunc,varaiya}@eecs.Berkeley.EDU

Abstract. The paper describes an approach to organizing a collection of spatially distributed entities into a system capable of carrying out different complex missions. The entities are devices (artifacts like vehicles, sensors, actuators, computing and communications equipment) or agents or controllers, which control devices or other agents. The system-level capabilities concern complicated activities such as reconnaissance and logistics, requiring optimization, planning, and real-time control. The approach is to construct a system-level capability by appropriately organizing (composing) agents that can execute a small set of tasks. Agents are organized in a hierarchy: those at upper layers have authority over those at lower layers. The capability to conduct different missions is achieved by changing the agent architecture. The hybrid system specification and simulation language SHIFT is convenient for describing devices and agents, and their interaction. The paper gives two illustrative designs, one specified in SHIFT. To deploy systems organized in this way poses several research challenges, one of which is indicated here—the need for an exception-handling facility that provides a 'fallback' mechanism if an agent is unable to complete the task assigned to it.

1 Introduction

Figure 1 is an artist's conception of a futuristic battlefield. In our imagination, we can animate this static picture with a narrative describing several missions of surveillance, logistics and target tracking—conducted through the control and coordination of a variety of equipment, including vehicles, weapons, sensors and communications gear. We can also imagine that these entities could be reorganized and dedicated to different missions; and further, that these systems are autonomous, requiring little human intervention in the real-time control of equipment.

The paper describes an approach to organizing a collection of spatially distributed entities into a *system* capable of carrying out a variety of complex missions. Conceptually, we distinguish between entities that are *devices*, and those that are *controllers* or *agents*. Devices are artifacts such as vehicles, sensors, actuators, and communications equipment, which interact with the the physical world. These interactions are continuous-time signals. Controllers are artifacts

* Research supported by ONR Contract N00014-98-1-0585 and Darpa Contract F33615-00-C-1698

T.A. Henzinger and C.M. Kirsch (Eds.): EMSOFT 2001, LNCS 2211, pp. 435–450, 2001.
© Springer-Verlag Berlin Heidelberg 2001

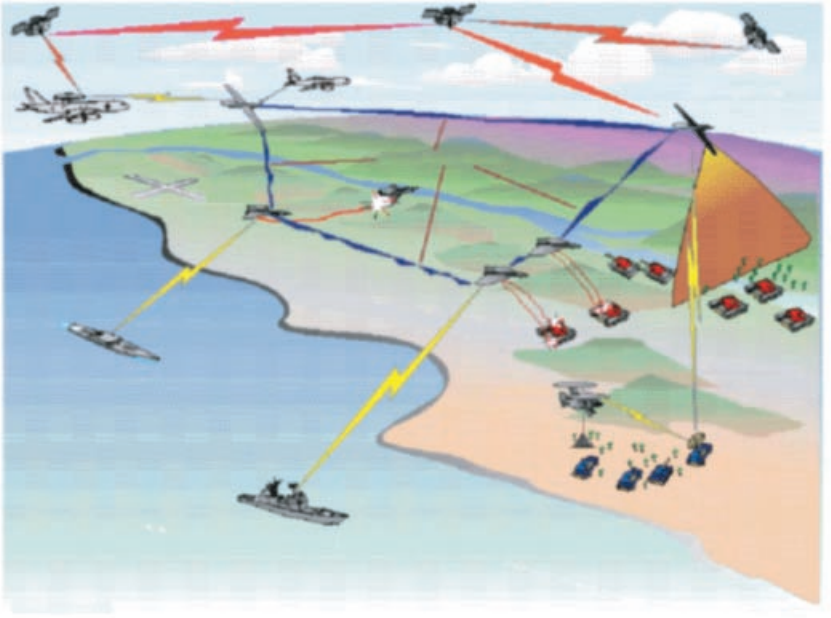

Fig. 1. Artist's conception of a future battlefield

such as algorithms, programs, and feedback control policies, which control the devices or other agents by exchanging symbolic messages. Interfaces "translate" between messages and signals. (The distinction between devices and agents is a matter of context: a vehicle may be viewed as a device or as a system comprising several devices and agents.)

To summarize, our interest is in organizing a system capable of complex missions involving optimization, planning, coordination, and control of entities over space and time. We suppose that each device has its own characteristic capability, which we cannot change. For example, a vehicle can move with a certain speed, a sensor can measure temperature, and a transmitter can broadcast a signal. On the other hand, we are free to design controllers, and to combine entities—devices and agents—into a system with a capability far superior to the capabilities of the constitutive devices. For example, the sensor, transmitter and vehicle could be organized into a system that autonomously measures, maps, and transmits the temperature profile over a designated land area. Evidently, the system-level capability to conduct a temperature survey emerges from an appropriate *organization* of the entities. The same entities can be organized in different ways to create different system-level capabilities.

Work in hierarchical system design forms two streams. The older stream of research, now called "mechanism design," began in the 1920s by asking if it is possible to organize a 'socialist' economy in which decisions $x_i \in X_i \subset R^n$ of