

Preface

In 1992 we initiated a research project on large scale distributed computing systems (LSDCS). It was a collaborative project involving research institutes and universities in Bologna, Grenoble, Lausanne, Lisbon, Rennes, Rocquencourt, Newcastle, and Twente. The World Wide Web had recently been developed at CERN, but its use was not yet as common place as it is today and graphical browsers had yet to be developed. It was clear to us (and to just about everyone else) that LSDCS comprising several thousands to millions of individual computer systems (nodes) would be coming into existence as a consequence both of technological advances and the demands placed by applications. We were excited about the problems of building large distributed systems, and felt that serious rethinking of many of the existing computational paradigms, algorithms, and structuring principles for distributed computing was called for. In our research proposal, we summarized the problem domain as follows: “*We expect LSDCS to exhibit great diversity of node and communications capability. Nodes will range from (mobile) laptop computers, workstations to supercomputers. Whereas mobile computers may well have unreliable, low bandwidth communications to the rest of the system, other parts of the system may well possess high bandwidth communications capability. To appreciate the problems posed by the sheer scale of a system comprising thousands of nodes, we observe that such systems will be rarely functioning in their entirety. The potential for unintentional unavailability of system services will result not only from failed hardware and software modules, but also from congested gateways and communications channels. The system is expected to change continually: nodes will be mobile, new components and services will be added, existing services will be modified, some components will be taken out of service and so forth.*”

The project, named BROADCAST (standing for: Basic Research On Advanced Distributed Computing: from Algorithms to SysTems), was funded by the European Union from 1992 to 1995 under the ESPRIT Basic Research Scheme. The project continued as a working group funded by the Union from 1996 to 1999, and incorporating two industrial partners (ANSA/APM, now part of Citrix, and Chorus systems, now part of Sun Microsystems).

The papers in this volume represent some of the key research results produced by the BROADCAST partners. We have structured this volume in four parts. Part one, *Distributed Algorithms*, chapters 1-6, contains papers on topics that address fundamental issues, such as time, clocks, and ordering of events, consensus in distributed systems, group communication, and garbage collection. Part two, *Systems Architecture*, chapters 7-12, addresses system structuring issues that range from treatment of network partitions to dependable distributed objects and Web services. A number of application support environments are described in the third part, *Application Support*, chapters 13-17; these include support for Web caching, transactional workflow system, component based programming environment and support for mobile Java applications. The fourth and final part of the volume, *Case Studies*, chapters 18-21, describes a number of working systems built by the partners.

We are confident that the papers in this volume will prove rewarding to all computer scientists working in the area of distributed systems. As already stated, the work reported here, spanning from 1992 to 1999, has been funded in part by the ESPRIT Project BROADCAST (project No. 6366) and the ESPRIT Working Group BROADCAST (project no. 22455). We thank Nick Cook, who acted as the administrative coordinator during this period, for overseeing the smooth running of our research activities.

November 1999

Sacha Krakowiak
Santosh Shrivastava