Real-Time Systems, 12, 1–4 (1997)

Guest Editor's Preface

ULRICH SCHMID s@auto.tuwien.ac.at
Department of Automation, Technical University of Vienna, Treitlstraße 3, A-1040 Vienna, Austria

Et dicimus tempus et tempus, tempora et tempora: "Quandiu dixit hoc ille", "Quandiu fecit hoc ille" et: "Quam longo tempore illud non vidi" et: "Duplum temporis habet haec syllaba ad illam simplam brevem". Dicimus haec et audimus haec et intellegimur et intellegimus. Manifestissima et usitatissima sunt, et eadem rursus nimis latent et nova est inventio eorum.
Aurelius Augustinus[1], *Confessiones*, XI, 22 (28)

During the last decade, much research has been conducted on how to provide synchronized clocks in fault-tolerant distributed systems. Global time and approximately synchronized clocks are, therefore, well-established terms in computer science — phrases like *"... algorithm xx requires synchronized clocks with precision ε ..."* are commonly used in modern textbooks on very different subjects, like communication protocols and data base systems. The clock synchronization algorithms of Halpern/Simons/Strong/Dolev [3], Lundelius/Lynch [5], Lamport/Melliar Smith [4], Mahaney/Schneider [6], Cristian/Aghili/Strong [1], Srikant/Toueg [7] (and the many variants derived from them), but also related theoretical work form a well-established basis for applications that require synchronized clocks.

So why wasting time and money by publishing a 1997 special issue devoted to global time in distributed systems at all?

Because enabling technology —more specifically, the *Global Positioning System* (GPS)— has been put into operation, which inevitably creates new user demands.

The earlier research mentioned above is almost exclusively devoted to internal clock synchronization, which provides a common, *but arbitrary* notion of time among the nodes of a distributed system. However, today's applications —in particular, very large scale ones like flight control systems— require some global time that has a well-defined relation to the only official and *legal* world-wide standard time: *Universal Time Coordinated* (UTC). Providing a global time in distributed systems that satisfies this additional accuracy requirement is usually termed the external synchronization problem, since UTC is provided externally to the system. Up to now, it did not receive much attention in research and practice, primarily due to the fact that there was no affordable way of acquiring UTC with high accuracy and high reliability.

The introduction of GPS, which caused a revolution in fields like navigation and land surveying, has also changed this situation completely. Apart from opening up new possibilities of transferring time with ultra high accuracy (ns range), it also provides access to UTC with some 100 ns accuracy at very low cost. Therefore, users —and hence, application engineers— become increasingly reluctant to accept some internal system time that does not relate to UTC. It was eventually recognized by the research community that this requires new research efforts, opening up a new and promising field of both theoretical and applied research. In fact, several independent activities pushing research in this area were launched recently: Apart from this special issue, which was initiated by the European editor-in-chief of *J. Real-Time Systems*, Wolfgang A. Halang, one of the celebrated *Dagstuhl-Seminars* —organized by Danny Dolev, Rüdiger Reischuk, Fred Schneider, and Ray Strong— was held March, 10–15, 1996, which was entirely devoted to the subject of *Time Services*.

Both quality and quantity of the papers contained in this special issue verify the liveness and fruitfulness of recent research in that area: A whole volume, i.e., three issues, of J. Real-Time Systems was required to accomodate the final material.

Issue 12:1

The opening paper of this issue, *Global Positioning System (GPS) Time Dissemination for Real-Time Applications* provides a comprehensive yet reasonably thorough introduction to the capabilities of GPS. I am particularly pleased that PETER H. DANA, who is a well-known GPS expert, accepted my invitation to write this paper. To supplement the information on external time sources, an overview of the capabilities of traditional radio-based time distribution systems is provided in the paper *Terrestrial Time Signal Dissemination* by REINER LICHTENECKER.

Given the fact that GPS has matured rapidly, it is only natural that more and more applications come up that exploit its power. The paper *Mobile, Distributed Real-Time Computing in a Vehicle Tracking System* by BERNHARD STERZBACH surveys an example of a typical state-of-the-art industrial application. Apart from that, GPS also provides enabling technology for novel real-time multiprocessor architectures, like the one outlined in the paper *High Accuracy Concurrent Event Processing in Hard Real-Time Systems* by WOLFGANG A. HALANG AND MARKUS WANNEMACHER. The concluding paper *Possible and Impossible Self-Stabilizing Digital Clock Synchronization in General Graphs* by SHLOMI DOLEV deals with the problem of how to provide synchronized clocks in such shared memory multiprocessor systems.

— ⋄ —

Today's applications typically depend exclusively on (the reliability of) GPS, which is a complex system with many possible sources of errors. For safety-critical systems, such solutions are of course not feasible. In fact, is was early recognized by the research community that the problem of fault-tolerant external synchronization constitutes a research topic in its own right. Internal clock synchronization algorithms cannot easily be extended to provide accuracy as well as precision, since obeying "authoritative" information of a few UTC time sources is incompatible with fault-tolerance. Moreover, internal synchronization is usually limited to LAN distances, which are certainly exceeded by future generation distributed systems. The remaining papers of this special issue form a quite representative overview of current efforts to solve those problems.

Issue 12:2

This issue contains two comprehensive papers primarily devoted to the theoretical background of fault-tolerant external synchronization. The first contribution by Christof Fetzer and Flaviu Cristian on *Integrating External and Internal Clock Synchronization* provides an external clock synchronization algorithm that degrades to internal synchronization in case of excessively many faulty UTC sources, along with a number of important lower bound results. The paper *Interval-based Clock Synchronization* by Ulrich Schmid and Klaus Schossmaier provides a novel generic framework for analyzing worst case precision and accuracy of interval-based algorithms, which relies upon a very detailled system model, incorporating issues like clock granularities and broadcast latencies in full detail.

Issue 12:3

This last issue primarily deals with engineering aspects in fault-tolerant external synchronization. The opening contribution *CesiumSpray: a Precise and Accurate Global Clock Service for Large-scale Systems* by Paulo Veríssimo, Luís Rodrigues, and António Casimiro provides an overview of a fully engineered solution for a particular fault-tolerant distributed real-time system architecture, which provides high precision/accuracy without hardware support by using the *a posteriori* agreement technique. By contrast, an elaborate time processor supporting very high accuracy interval-based clock synchronization by hardware is described in the paper *Specification and Implementation of the Universal Time Coordinated Synchronization Unit (UTCSU)* by Klaus Schossmaier, Ulrich Schmid, Martin Horauer, and Dietmar Loy.

To complete the overall picture, in particular, with respect to the upcoming research on clock synchronization in partially synchronous systems, an appendix containing the abstracts of the presentations given at the Dagstuhl-Seminar on Time Services is provided. I am grateful to Danny Dolev, Rüdiger Reischuk, Fred B. Schneider, and H. Raymond Strong for their consent to reprint the official seminar report [2] in this special issue.

— ⋄ —

Selecting all the papers above, I tried to collect a special issue that may be considered reasonably representative for the current status of research on the subject. To maintain high scientific quality, thorough peer-reviewing was employed; I am particularily grateful to all reviewers for their expertise, and to the authors for their endurance and cooperation. I would be happy if the research community could benefit from the efforts spent on it. In particular, the special issue will have served its purpose if the interested reader recognizes the relevance of the problem — and, maybe, spends some time on scrutinizing the material provided therein.

Notes

1. For the reader's convenience, we also provide the translation of the quotation taken from St. Augustin's Confessions: *And we speak as to time and time, times and times,— "How long is the time since he said this?" "How long the time since he did this?" and, "How long the time since I saw that?" and, "This syllable hath double the time of that single short syllable." These words we speak, and these we hear; and we are understood, and we understand. They are most manifest and most usual, and the same things again lie hid too deeply, and the discovery of them is new.*

References

1. F. Cristian, H. Aghili, R. Strong. *Clock Synchronization on the Presence of Omission and Performance Faults and Processor Joins*, Proc. 16th International Symposium on Fault Tolerant Computing Systems, Vienna, Austria, July 1986, p. 218–223.
2. D. Dolev, R. Reischuk, F.B. Schneider, R. Strong (eds.) Seminar-Report ???? on Dagstuhl-Seminar *Time Services*, Schloß Dagstuhl, Germany, March 1996.
3. J. Halpern, B. Simons, R. Strong, D. Dolev. *Fault-Tolerant Clock Synchronization*, Proc. 3rd Annual ACM Symposium on Principles of Distributed Computing, August 1984, p. 89-102.
4. L. Lamport, P.M. Melliar-Smith. *Synchronizing Clocks in the Presence of Faults*, Journal of the ACM, 32(1), January 1985, p. 52–78.
5. J. Lundelius, N. Lynch. *A New Fault-Tolerant Algorithm for Clock Synchronization*, Proc. 3rd ACM Symposium on Principles of Distributed Computing, Vancouver, Canada, August 1984, p. 75–88.
6. S. Mahaney, F.B. Schneider. *Inexact Agreement: Accuracy, Precision and Graceful Degradation*, Proc. 4th Annual ACM Symposium on Principles of Distributed Computing, August 1985, p. 237–249.
7. T.K. Srikanth, S. Toueg. *Optimal Clock Synchronization*, Journal of the ACM, 34(3), July 1987, p. 626–645.