

Dagstuhl Seminar
on
Time Service

Danny Dolev, Hebrew University
Rüdiger Reischuk, Med. Universität zu Lübeck
Fred B. Schneider, Cornell University
H. Raymond Strong, IBM Almaden Research

Schloß Dagstuhl, March 11. – March 15. 1996

Contents

Introduction	3
Final Seminar Programme	4
Abstracts of Presentation:	
David Mills: <i>Precision Network Time Synchronization</i>	5
Keith Marzullo: <i>Synchronizing Clocks and Reading Sensors</i>	5
John Rushby: <i>Formal Verification of Clock Synchronization Algorithms</i> ..	6
Ulrich Schmid: <i>Interval-based Clock Synchronization</i>	7
Klaus Schossmaier: <i>UTCSU - An ASIC for Supporting Clock Synchronization for Distributed Real-Time Systems</i>	8
Wolfgang Halang: <i>Time Services Based on Radio Transmitted Official UTC</i>	9
Christof Fetzer: <i>Fail-Aware Clock Synchronization</i>	9
Paulo Verissimo: <i>Cesium Spray: a Precise and Accurate Time Services on Large-Scaled Distributed Systems</i>	10
Shlomi Dolev: <i>Self-Stabilizing Clock Synchronization Algorithms</i>	10
Wolfgang Reisig: <i>A Temporal Logic for “As Soon As Possible”</i>	11
Augusto Ciuffoletti: <i>Self-Stabilization Issues in Clock Synchronization Algorithms</i>	11
Sergio Rajsbaum: <i>Extending Causal Order with Real-Time Specifications and Applications to Clock Synchronization</i>	12
Richard Hofmann: <i>A Fault-Tolerant Clock Synchronization Method with Correctness Proof and a Posteriori Correction</i>	12
Timothy Mann: <i>Marching to Many Distant Drummers</i>	13
Boaz Patt-Shamir: <i>Complexity of Clock Synchronization</i>	14
Ray Strong: <i>Reading Clocks by Eavesdropping</i>	14
M. Azadmanesh: <i>Exploiting Omissive Faults in Synchronous Approximate Agreement</i>	15
Roger Kieckhafer: <i>Hybrid Fault Models for Byzantine-safe Systems: Have they Reached the Limit?</i>	15
Rüdiger Reischuk: <i>Observable Clock Synchronization</i>	16
Participant List	17

Introduction

This seminar brought together a significant fraction of the world's researchers and engineers of network and distributed time services. Until now, there had been little communication between those who implemented time services for computer networks and those exploring new algorithms and analysis techniques for distributed time services. Our Dagstuhl seminar housed – for the first time – representatives from each of the major approaches to clock synchronization and each of the major providers of real-time synchronization services.

The formal presentations covered new clock synchronization algorithms, methods and actual analysis of these algorithms, and the details of fielded implementations. Participants learned not only the details of the current state of the Internet time protocol (NTP), but also about the history and origins of the algorithms in that implementation.

Some participants could see the practical impacts their fundamental research was having. Others learned about what problems are real and what problems could be safely put aside as academic curiosities. And, implementors learned of new algorithms ripe for trial.

The informal discussions and interactions were at least as important. The meeting forged a community of scientists and engineers with like interests where none had been. Workers in time services now realize the importance of their efforts and have a set of peers with whom to discuss their work. Clock synchronization and time services are a critical part of the infrastructure for any distributed network. By creating a community of workers in this area, the Dagstuhl seminar has itself created an infrastructure to facilitate further advances in the subject.

Schloss Dagstuhl provided an excellent atmosphere for this meeting. We like to thank the staff and the Dagstuhl foundation who made this exchange of research possible.

The organizers:

Danny Dolev, Fred Schneider, Rüdiger Reischuk, Ray Strong

Compiled by Andreas Jakoby

Monday, 11. March 1996

- 9.15 - 10.15 David Mills: *Precision Network Time Synchronization*
10.30 - 11.45 David Mills: *Precision Network Time Synchronization II*
15.30 - 16.30 Keith Marzullo: *Synchronizing Clocks and Reading Sensors*
16.45 - 17.45 John Rushby: *Formal Verification of Clock Synchronization Algorithms*

Tuesday, 12. March 1996

- 9.00 - 10.00 Ulrich Schmid: *Interval-based Clock Synchronization*
10.20 - 11.10 Klaus Schossmaier: *UTCSU - An ASIC for Supporting Clock Synchronization for Distributed Real-Time Systems*
11.15 - 12.05 Wolfgang Halang: *Time Services Based on Radio Transmitted Official UTC*
15.30 - 16.30 Christof Fetzer: *Fail-Aware Clock Synchronization*
16.45 - 17.45 Paulo Verissimo: *Cesium Spray: a Precise and Accurate Time Services on Large-Scaled Distributed Systems*

Wednesday, 13. March 1996

- 9.00 - 10.00 Shlomi Dolev: *Self-Stabilizing Clock Synchronization Algorithms*
10.30 - 11.15 Wolfgang Reisig: *A Temporal Logic for "As Soon As Possible"*
11.25 - 12.00 Augusto Ciuffoletti: *Self-Stabilization Issues in Clock Synchronization Algorithms*

Thursday, 14. March 1996

- 9.00 - 10.00 Sergio Rajsbaum: *Extending Causal Order with Real-Time Specifications and Applications to Clock Synchronization*
10.20 - 11.20 Richard Hofmann: *A Fault-Tolerant Clock Synchronization Method with Correctness Proof and a Posteriori Correction*
11.25 - 12.10 Timothy Mann: *Marching to Many Distant Drummers*
16.00 - 17.00 Boaz Patt-Shamir: *Complexity of Clock Synchronization*
17.15 - 17.45 Ray Strong: *Reading Clocks by Eavesdropping*

Friday, 15. March 1996

- 9.15 - 10.00 M. Azadmanesh: *Exploiting Omissive Faults in Synchronous Approximate Agreement*
10.30 - 11.15 Roger Kieckhafer: *Hybrid Fault Models for Byzantine-safe Systems: Have they Reached the Limit?*
11.20 - 12.05 Rüdiger Reischuk: *Observable Clock Synchronization*
12.15 end of the conference

Abstracts of Presentation

Computer Network Time Synchronization

David L. Mills, University of Delaware

This presentation consists of two parts, the first on the architecture, protocol and algorithms of the Network Time Protocol (NTP), which is used to synchronize computer clocks in the Internet, and the second on planned enhancements of NTP to support an autonomous configuration and cryptographic authentication service. The talk covers the architecture model, consisting of the algorithms to select the best sample from a series of samples for each of several servers, then determine the best subset of servers using an intersection algorithm modified from one due to Marzullo. Addutuibak algorithms are used to delete statistical outliers and improve accuracy by a weighted average of the survivors.

A key to the performance of NTP is the local clock model, which is based on a type-II, adaptive parameter, hybrid phase/frequency-lock loop. The parameters of the model and the phase-frequency mode in which it operates are automatically determined by the prevailing network delays, dispersions and measured local clock stability.

The talk continues to a discussion of enhancements to the NTP model, both in progress and proposed. These include autonomous configuration features which survey the environment and construct a hierarchical spanning tree based on a metric including delay and dispersion, subject to constraints in node degree and total distance. Other enhancements include provision of cryptographic authentication schemes using hybrid public and private key cryptosystems. Factors such as state space required and cryptosystem computation burden are discussed, along with specific proposed systems that combine the best features of both.

Synchronizing Clocks and Reading Sensors

Keith Marzullo, University of California at San Diego

We describe an internal clock synchronization protocol for a broadcast-based local area network. This protocol is especially adapted for hard real-time distributed systems because it makes very small the schedulability impact on both the processor and the network. It combines aspects of Cristian's probabilistic clock synchronization protocol and Verissimo's a posteriori clock synchronization.

We reviewed the attractions and the problems with these two clock synchronization protocols and present a new protocol, which we call silent clock synchronization, that combines the attractions of both while avoiding their problems. We then present an experimental setup and discuss how we measured the failure rates and the system parameters needed to build our protocol, and discuss the tightness that we attain.

Formal Verification of Clock Synchronization Algorithms Under Byzantine and Hybrid Fault Models

John Rushby, Computer Science Laboratory
SRI International Menlo Park California

Their intricacy and criticality makes clock synchronization attractive targets for mechanized formal verification. I will describe my verification (undertaken jointly with Friedrich von Henke) of the Interactive Convergence Algorithm of Lamport and Melliar-Smith. During this exercise, we found the proof of the main induction, and four of the five supporting lemmas, to be flawed in the original treatment. Our mechanized verification not only enabled us to detect and correct these flaws, but also to remove the approximations and to provide simpler and more uniform proofs than the originals. However, the main benefit of the formal specification and verification has been its ability to support exploration of variations on the original algorithm. The calculational character of mechanized deduction allows alternative designs and requirements to be examined in much the same way that computational fluid dynamics supports the exploration of different aerofoil designs. I will describe how, using these capabilities, we were able to eliminate one of the assumptions from the original treatment, to sharpen the statements of others, and to extend the algorithm to a more complex “hybrid” fault model that allows larger numbers of simple faults to be tolerated, without sacrificing Byzantine fault tolerance.

I will also describe similar lessons drawn from my colleague Shankar’s verification of Schneider’s general treatment of clock synchronization algorithms, and its subsequent refinement by Paul Miner of NASA.

Bibliography:

John Rushby and Friedrich von Henke. *Formal verification of algorithms for critical systems*. IEEE Transactions on Software Engineering, 19(1):13–23, January 1993.

John Rushby. *A formally verified algorithm for clock synchronization under a hybrid fault model*. In Thirteenth ACM Symposium on Principles of Distributed Computing, pages 304–313, Los Angeles, CA, August 1994. Association for Computing Machinery.

Natarajan Shankar. *Mechanical verification of a generalized protocol for Byzantine fault-tolerant clock synchronization*. In J. Vytopil, editor, Formal Techniques in Real-Time and Fault-Tolerant Systems, volume 571 of Lecture Notes in Computer Science, pages 217–236, Nijmegen, The Netherlands, January 1992. Springer-Verlag.

Paul S. Miner. *Verification of fault-tolerant clock synchronization systems*. NASA Technical Paper 3349, NASA Langley Research Center, Hampton, VA, November 1993.

Paul S. Miner, Shyamsundar Pullela, and Steven D. Johnson. *Interaction of formal design systems in the development of a fault-tolerant clock synchronization circuit*. In 13th Symposium on Reliable Distributed Systems, pages 128–137, Dana Point, CA, October 1994. IEEE Computer Society.

Interval-based Clock Synchronization

Ulrich Schmid, Klaus Schossmaier, Technical University of Vienna

We present¹ description and analysis of a novel algorithm suitable for fault-tolerant external clock synchronization. Unlike usual internal synchronization approaches, our interval-based *orthogonal accuracy algorithm* OA provides approximately synchronized clocks maintaining precision and accuracy w.r.t. external time (*universal time coordinated* UTC) simultaneously. This is accomplished by means of a time representation relying on intervals that capture UTC, providing accuracy information encoded in interval lengths. Our (generic) analysis, which considers important practical issues like non-zero clock granularity, utilizes a novel, interval-based framework for providing worst-case precision and accuracy bounds subject to a fairly realistic system model.

Keywords: external clock synchronization; fault-tolerant distributed real-time systems; universal time coordinated (UTC); convergence functions; precision analysis; accuracy intervals; clock granularity; Marzullo’s function.

¹This research is part of our project SynUTC, which is supported by the Austrian Science Foundation (FWF) under grant no. P10244-ÖMA.

UTCSU - An ASIC to Support Clock Synchronization for Distributed Real-time Systems

Klaus Schossmaier, TU Vienna

Targeting at a $1 \mu s$ precision/accuracy for a time-service in the real-time system area, we need to support the clock synchronization algorithm with adequate hardware. In our project SynUTC we are developing a costume VLSI chip with the following features:

- A rate and state adjustable local clock in the NTP-time format. We proposed a new kind of clock, composed of an oscillator and an adder, which adds a particular amount at each tick. This allows us to carry out continuous amortization, dealing with leap-seconds and the usage an a non-binary input frequency.
- Maintenance of an asymmetrical accuracy interval with proper deterioration dynamics in order to capture UTC. Also here, 2 adder-based clocks are in use with certain supplements.
- Times tamping events of packet transmission/arrival to reduce the clock reading error, 1PPS from GPS-receivers for injection of external time and for application purpose.
- Generating events with the help of duty-timers to support the protocol for a full message exchange (FME).
- For test and debugging purpose we implemented checksums, signatures, block-sums and snapshots.

The ASIC is manufactured by a $0.7\mu m$ standard cell CMOS technology and consumes about $60mm^2$.

Time Services Based on Radio Transmitted Official UTC

Wolfgang A. Halang, FernUniversität Hagen

ABSTRACT: Having severe deficiencies with respect to timing, predictable behavior, and delays between sensing and actuation, contemporary computer systems employed in real-time operation are inappropriate to a surprisingly high extent. It is pointed out how these deficiencies can be eliminated with available technology. The concept of alarm jobs and high-precision timers handling these alarm jobs is presented. They are to serve in embedded, distributed real-time systems to achieve precise time information, accurate time stamping of interrupt occurrences, precisely timed sensing and actuation operations, and comfortable time handling. No complicated clock synchronization by software is needed anymore, since the exact — and legal — Universal Time Co-ordinated (UTC) is received via satellite from GPS, the global navigation and positioning system. A hardware realization is described making use of an application specific integrated circuit. It is shown that the presented solution to the timing and clock synchronization problems is technically and economically feasible.

Fail-Aware Clock Synchronization

Christof Fetzer, University of California at San Diego

We address the problem of the impossibility of implementing internal clock synchronization in asynchronous systems. Fail-aware clock synchronization provides a “synchronization indicator” for each server and requires that

1. the deviation between the clocks of two servers is bounded by an a priori given constant whenever their indicators are true, and
2. a timely server which can communicate with a majority of servers in a timely manner has a indicator which is true.

We show how fail-aware clock synchronization can be implemented in timed asynchronous systems. These systems are characterized by

1. unbounded message transmission and process scheduling delays, and
2. that each process has access to a local hardware clock with a bounded drift rate.

CesiumSpray: a Precise and Accurate Global Clock Service for Large-scale Systems

P. Verissimo, Univ. Lisboa L. Rodrigues, Tech. Un. Lisboa
A. Casimiro, Univ. Lisboa

In large-scale systems, such as Internet-based distributed systems, classical clock-synchronization solutions become impractical or poorly performing, due to the number of nodes and/or the distance. We present a global time service for world-wide systems, based on an innovative clock synchronization scheme, named CesiumSpray. The service exhibits high precision and accuracy; it is virtually indefinitely scalable; and it is fault-tolerant. It is deterministic for real-time machinery in the local area, which makes it particularly well-suited for, though not limited to, large-scale real-time systems. The clock synchronization scheme is a hierarchical mix of external and internal synchronization. The root of the hierarchy are the GPS satellites, which “spray” their reference time over a set of nodes provided with GPS receivers, one per local network, where the second level of the hierarchy performs internal synchronization, further “spraying” the external time inside the local network. The algorithm of the second level is inspired on the high precision a posteriori agreement synchronization algorithm, modified to follow an external clock, and able to use simple group communication and membership facilities.

Self-Stabilizing Synchronization Algorithms

Shlomi Dolev, Ben-Gurion University

Three self-stabilizing clock synchronization algorithms will be presented. The first algorithm is self-stabilizing clock synchronization algorithm that uses bounded clock values. The algorithm guarantees convergence to a state where all clock values are identical, and are subsequently maintained to be identical. It is assumed that during the convergence period no new fault occurs.

The second algorithm is wait-free self-stabilizing clock synchronization algorithm. This algorithm is more robust than the first algorithm since it is able to cope with faults during the convergence period as well. The algorithm guarantees convergence even if processors repeatedly stop executing their program and restart.

The last algorithm consider the most sever faults. The algorithm copes with a more severe (and realistic) fault model than the traditional Byzantine fault model. This algorithm stabilizes to synchronize the system clocks in the presence of Byzantine faults — A processor that experience Byzantine fault

can exhibit arbitrary “malicious”, “two faced”, behavior as if it is controlled by an adversary.

The talk summarizes joint works with Anish Arora, Mohamed Gouda, and Jennifer L. Welch.

A Temporal Logic for “as soon as possible”

Wolfgang Reisig,

Time assumptions can be used to make causal assumptions (implemented by signals) redundant. Nevertheless, the untimed version of an algorithm, with causal dependencies made explicit, is frequently easier to understand and to prove.

A temporal logic is suggested that book-keeps the local resources needed for an action. A local progress operator describes an action’s occurrence “as soon as possible”, avoiding any interleaved clock synchronization. Its transitive composition describes progress “as soon as possible” with given resources. Powerful composition rules allow to embed local arguments into parallel context.

Self-Stabilization Issues in Clock Synchronization

Augusto Ciuffoletti, Universit  di Pisa

We present a new probabilistic clock synchronization algorithm, its prototype implementation and experimental results. The algorithm follows the client-server programming paradigm and is designed to work in a departmental environment with few servers and a number of clients connected through an arbitrary network topology.

At the core of the algorithm is a remote clock reading method that mitigates the negative effects of message delay uncertainty. The implementation proves the effectiveness of this approach and corroborates the theoretical speculations.

Extending Lamport's Partial Order with Real-time Specifications and Applications to Optimal Clock Synchronization ²

Sergio Rajsbaum, Instituto de Matematicas-UNAM, Mexico
Joint work with Boaz Patt-Shamir.

Abstract: We consider the problem of clock synchronization in a system with crash faults, with uncertain message delays and bounded clock drifts. To study what is the best possible precision, per execution rather than only in worst cases, we propose a paradigm that extends Lamport's partial order with real-time specifications. Typically, these are specifications related to bounds on message transmission delays and on clock drifts, but can model also other real-time specifications.

We present a characterization theorem for the tightest achievable estimate of the readings of a remote clock in any given execution of the system. Using this theorem, we obtain the first optimal on-line distributed algorithms for external clock synchronization. The general algorithm has unbounded space overhead, which is unavoidable, as we show. For systems with drift free clocks, we present a simple and efficient algorithm.

A Fault-tolerant Clock Synchronization Method with Correctness Proof and a Posteriori Correction

Richard Hofmann, University of Erlangen — IMMD VII

A clock synchronization method is presented that augments a universal distributed monitor system with an appropriate facility for time stamping. In order to find out what properties a clock for this purpose must have, the basic problems of monitoring and evaluation of parallel and distributed computer systems are figured out. It turns out that causality of events and time are important properties when analyzing such systems.

There are events that are causally related and other ones that are independent from each other. For the latter ones it is not necessary to know their sequence because it is irrelevant. On the other hand, causally related events must be reflected in the correct sequence. As there is a minimum time between the

² Appeared in Proc. 26th Symp. on Theory of Computing, May 1994.

causing event and the dependent event in a causal relationship, this minimum time interval is the measure for constructing clocks for monitor systems. Currently this time distance lies in the vicinity of half a microsecond.

In order to correctly reflect the sequence of causally related events all clocks used for time stamping in a parallel and distributed system must not deviate more than 1/4 of a microsecond. This can only be achieved by a hardware synchronization method. The method presented in the talk operates on two levels of information that are transferred over the same physical channel, RS 485 serial ports on the master and the slaves and suitable cabling between them.

The lower level of the synchronization mechanism consists of a PLL scheme that guarantees the same rate for all slave clocks. More precisely the PLL scheme was optimized for lowest time difference between any two slave clocks with respect to oscillator thermal noise and noise induced by the cabling. Mathematical worst-case analysis has shown error margins of about 5 nanoseconds; our measurements showed an error of less than 2 nanoseconds. Besides providing this precision the basic synchronization level also contains provisions for fault-tolerance by monitoring the status of the PLL. This status is documented as an attribute to each event.

Based on the services of the basic synchronization level, the upper level guarantees a synchronous start of all clocks, a supervising facility for the synchronization mechanism, and a synchronous stop of all clocks and measurement activities. During a measurement session, so-called sync-tokens are issued by the master of the synchronization scheme after regular time intervals (resynchronization intervals). Every time such a sync-token is received, an event is generated and stored in the same way as an event from the monitored system.

The regular time intervals allow to check the time stamps of those events generated by the sync-token — these events must have integer multiples of the length of the resynchronization interval. If this is not the case, corrections can be made according to the difference encountered at such events.

Marching to Many Distant Drummers

Timothy Mann, DEC Systems Research Center

I haven't worked on time services recently, but I've prepared a talk on some unpublished joint work I did with Leslie Lamport in 1990. My background also includes a time service implementation for SRC's Topaz distributed operating system, based on an earlier algorithm of Lamport's.

In the 1990 work, we address the problem of determining the time in a network where a node may obtain information indirectly from primary time

sources via intermediate nodes. Our key idea is to transmit and store each time datum as a pair, consisting of a time interval and a “failure predicate”, a boolean expression that indicates precisely which combinations of node failures could invalidate the interval. We describe some techniques based on this idea, but not a complete system design or implementation.

The Complexity of Clock Synchronization

Boaz Patt-Shamir, Northeastern University, Boston

Abstract: Based on our theory of clock synchronization, we prove that any clock synchronization algorithm which is both

- General (i.e., works for all systems), and
- Optimal (i.e., outputs the tightest bounds at any given instant)

is inherently inefficient. Specifically, we define a computational model (a variant of the branching program model), and show that in that model, any optimal synchronization protocol cannot have bounded space complexity. The result is demonstrated by considering a set of executions of a system consisting of four processors, two of which have drift-free clocks.

Joint work with Sergio Rajsbaum.

Reading Clocks by Eavesdropping

Ray Strong, IBM Almaden Research

(joint work with Danny Dolev, Rüdiger Reischuk, and Ed Wimmers)

The most straightforward way for a client to read a server’s clock involves round trip communication between client and server. If there is a requirement that each client read its server’s clock at least r times per hour, then the straightforward solution requires at least r message transmissions per hour from each client. We describe a simple protocol that allows one client to read its server’s clock by eavesdropping on communication between the server and another client. To be sure that the protocol works on a given network, we do not have to measure accurately such quantities as maximum transmission delay, minimum transmission delay, or tightness of broadcast receipt. We have only to establish that, relative to our method of reading clocks, there is some upper bound on the rate of observable clock drift.

Exploiting Omissive Faults in Approximate Agreement

M.H. Azadmanesh, University of Nebraska at Omaha

Abstract - The existing voting algorithms, including the Mean of Subsequenced Reduced (MSR) voting algorithms, based on the Thambidurai and Park fault model assume the worst case behavior of malicious faults. Hence, omissive faults are either assumed not to occur, or a predefined value is substituted for the missing values which has the effect of transforming omissive errors into more severe fault modes such as symmetric or asymmetric. As a result, existing voting algorithms can not exploit omissive faults.

This research has further partitioned the symmetric and asymmetric faults into disjoint transmissive and omissive faults. Thus, a new fault model is formed which consists of five modes of failure: benign, transmissive symmetric, omissive symmetric, transmissive asymmetric, and strictly omissive asymmetric. Based on this model, a new family of voting algorithms, called Omission-MSR (OMSR), is introduced. By exploiting omissive faults, OMSR algorithms allow non-faulty processes to locally discard self-evident errors. Furthermore, OMSR voting algorithms are shown to be significantly more fault tolerant than existing voting algorithms.

Hybrid Fault Models for Byzantine-Safe Systems: Have we Reached the Limit ?

Roger M. Kieckhafer, University of Nebraska - Lincoln

In recent years considerable attention has been paid to the development of Hybrid Fault models for Byzantine-safe systems. In particular, Thambidurai and Park's 3-mode fault model (TPH-3) has been used to show that the fault-tolerance of both Byzantine Agreement and Approximate Agreement algorithms are considerably better than predicted by the single-mode Byzantine fault model (Byz-1). We have recently devised a 5-mode Omissive/Transmissive fault model (OTH-5) and a new family of Approximate Agreement algorithms which are more fault-tolerant than any previous algorithms when analyzed under the OTH-5 fault model.

This presentation asks whether the improved fault-tolerance of the new algorithms actually translates into improved Dependability, or whether the new fault-modes have little impact. The motivation is to determine whether the development of new hybrid fault models is worth the effort relative to practical applications. A Generalized Stochastic Petri-Net (GSPN) Reliability

model is presented for a representative distributed fault-tolerant real-time system. GSPN Modeling results over a wide range of parameter values show that the new algorithms and fault model can indeed improve system reliability significantly, given certain values for system parameters. However, results also show that incorrect “tuning” of an algorithm can actually reduce reliability. In some cases the relationships between parameter values and Reliability are quite counter-intuitive. Hence, extreme care must be taken when trying to apply the OTH-5 fault model to real systems.

Observable Clock Synchronization

Rüdiger Reischuk, Med. Universität zu Lübeck

We develop a relativistic theory of observable clock synchronization that does not use or depend on a Newtonian framework or real time. Within the context of this theory, we focus on the problem of estimating the time on a remote clock. We generalize the concept of *rapport* to capture the situation when such an estimate is sufficient for clock synchronization purposes.

With a single property, called the *Observable Drift Property*, we characterize the information flow required for obtaining rapport. Finally, we compare our relativized and observable concepts with analogs based on the notion of real time in order to show that we are studying the right quantities.

Joint work with Danny Dolev and Ray Strong.

Participant List

M.H. Azadmanesh
University of Nebraska at Omaha
Center for Management of
Information Technology
308 CBA Building
Omaha NE 68182-0459 / USA
email: azad@ahvaz.unomaha.edu
tel: +1-402-554-39 76
<http://ahvaz.unomaha.edu>

Augusto Ciuffoletti
Universit  di Pisa
Dipartimento di Informatica
Corso Italia 40
I-56125 Pisa / Italy
email: augusto@dipisa.di.unipi.it
tel: +39-50-887-265

Shlomi Dolev
Ben Gurion University
Department of
Computer Science and Mathematics
Beer Sheva 84105 / Israel
email: dolev@CS.bgu.ac.il
<http://www.cs.bgu.sc.il/~dolev>

Christof Fetzer
University of California at San Diego
Department of Computer Science
9500 Gilman Drive
La Jolla CA 92093-0114 / USA
email: cfetzer@cs.ucsd.edu
tel: +1-619-597-20 18
www.cse.ucsd.edu/users/cfetzer

Wolfgang A. Halang
FernUniversit t Hagen
FB Elektrotechnik
Informationstechnik / Realzeitsysteme
D-58084 Hagen / Germany
wolfgang.halang@fernuni-hagen.de
tel: +49-2331-987 - 372
<http://www.fernuni-hagen.de/IT>

Jaap-Henk Hoepman
Centre for
Mathematics and Computer Science
Kruislaan 413
NL-1098 SJ Amsterdam
The Netherlands
email: jhh@cwi.nl
tel: +31-20-592-40 76
<http://www.cwi.nl/~jhh>

Richard Hofmann
Friedrich Alexander Universit t
Lehrstuhl Informatik VII
Martensstr. 3
91058 Erlangen / Germany
rhofmann@informatik.uni-erlangen.de
tel: +49-9131-85-70 26

Andreas Jakoby
Med. Universit t zu L beck
Naturwissenschaftliche Fakult t
Institut f r Theoretische Informatik
Wallstra e 40
D-23560 L beck / Germany
email: jakoby@informatik.mu-luebeck.de
tel: +49-451/7030-417
<http://www.informatik.mu-luebeck.de>

Roger Kieckhafer
University of Nebraska - Lincoln
Dept. of Computer Science
Lincoln NE 68588-0115 / USA
email: rogerk@cse.unl.edu
tel: +1-402-472-24 01

Timothy Mann
DEC SRC Systems Research Center
130 Lytton Ave.
Palo Alto CA 94301-1044 / USA
email: mann@src.dec.com
tel: +1-415-853-22 24

Keith Marzullo
University of California at San Diego
Department of
Computer Science and Engineering 0114
9500 Gilman Drive
La Jolla CA 92093-0114 / USA
email: marzullo@cs.ucsd.edu
tel: +1-619-534-37 29

Wolfgang Reisig
Humboldt Universität Berlin
Institut für Informatik
Sitz: Lindenstr. 54
Unter den Linden 6
D-10099 Berlin / Germany
reisig@informatik.hu-berlin.de
tel: +49-30-20 18 12 20

David L. Mills
University of Delaware
Dept. of Electrical Engineering
Newark DE 19716 / USA
email: mills@udel.edu
tel: +1-302-831-8247
<http://www/eecis.edel.edu/~mills>

John Rushby
SRI International
Computer Science Laboratory
333 Ravenswood Ave.
Menlo Park CA 94025 / USA
email: rushby@csl.sri.com
tel: +1-415-859-54 56
<http://www/csl/sri/com/fm.html>

Boaz Patt-Shamir
Northeastern University
College of Computer Science
161-CN
Boston MA 02115 / USA
email: boaz@ccs.neu.edu
tel: +1-617-373-20 78

Ulrich Schmid
Technische Universität Wien
Institut für Automation
Treitlstraße 3
A-1040 Wien / Austria
email: s@auto.tuwien.ac.at
tel: +43-1-58801-8189

Sergio Rajsbaum
UNAM - Instituto de Matematicas
Ciudad Universitaria
Mexico D.F. 04510 / Mexico
email: rajsbaum@servidor.unam.mx
tel: +52-5-622 45 20

Fred B. Schneider
Cornell University
Department of Computer Science
Upson Hall
Ithaca NY 14853-7510 / USA
email: fbs@cs.cornell.edu

Rüdiger Reischuk
Med. Universität zu Lübeck
Naturwissenschaftliche Fakultät
Institut für Theoretische Informatik
Wallstraße 40
D-23560 Lübeck / Germany
reischuk@informatik.mu-luebeck.de
tel: +49-451-7030-416
<http://www.informatik.mu-luebeck.de>

Klaus Schossmaier
Technische Universität Wien
Institut für Automation
Treitlstraße 3
A-1040 Wien / Austria
email: kmschoss@auto.tuwien.ac.at
tel: +43-1-58801-8189
www.auto.tuwien.ac.at/~kmschoss

H. Raymond Strong
IBM Almaden Research
Dept. K53/802
650 Harry Road
San Jose CA 95120-6099 / USA
email: strong@almaden.ibm.com
tel: +1-408-927 - 1758

Paulo Verissimo
University of Lisboa
Faculty of Sciences
Campo Grande C5
P-1700 Lisboa / Portugal
email: pjv@di.fc.ul.pt
tel: +351-1-750 01 03
<http://pandora.inesc.pt/>

Paul Vitanyi
CWI - Mathematisch Centrum
Kruislaan 413
NL-1098 SJ Amsterdam
The Netherlands
email: paulv@cwi.nl
tel: +31-205-924124
<http://www.cwi.nl/~paulv>