Summary of the discussion sessions: FLP and real-time

Two extra evening discussion sessions during the workshop explored the implications of the Fischer-Lynch-Paterson Theorem (FLP) for practical system design, and the various notions of "real time" used in distributed settings.

The FLP Theorem states that consensus is unsolvable in a distributed system with asynchronous communications and one single faulty processors. This has serious consequences, for example when data is replicated to make a system fault tolerant. In order to keep replicated data globally consistent, the same updates have to be applied in the same order to each replica. Agreement on order, however, is a consensus problem.

Several ways of working around the consequences of the FLP Theorem were suggested in the discussions. According to Sam Toueg, consensus can be solved in an asynchronous system by adding a failure detector to the environment. This work is described in the PODC-91 paper written by Chandra and Toueg. However, not everyone understands the result. One objection is that, obviously, the weakest failure detector to solve consensus still cannot guarantee progress in an asynchronous environment. Moreover, it is not clear how the weakest failure detector could actually be "implemented" in a real system. The approach advocated by systems like Horus, Isis and Transis consists in weakening the notion of consensus, and only try to achieve consensus among groups of "responsive members". Each process in such group agrees on the membership of the group, and any process outside the group is ignored by the members. Processes outside a group cannot do any damage, but they may block until communication to the group is re-established. In the Horus system, the application programmer is also allowed to plug in any arbitrary failure detector, so that an application-dependent solution can be obtained. In these systems, however, network partitioning can prevent progress.

It is hoped that, within a few years, a rigorous approach to consensus based on a formally defined failure detector (for example, as advocated by Chandra, Toueg and Hadzilacos), will be reduced to practice by the community building fault-tolerant distributed systems. Such a development would lead to to systems with well defined liveness properties and a more explicit characterization of limitations stemming from FLP.

The second discussion section focused on notions of "real-time" in distributed computing systems. The correctness of many distributed real-time systems relies on the fact that fault-free processors can deterministically reach a consensus within a given upper time bound. The FLP Theorem states that in this case one must make the so-called "synchronous communication" assumption, i.e., that the underlying communication system delivers inter-processor messages within a bounded delay. Of course, this is a very strong assumption for any real communication system. Even if one fixes an upper bound on communication delay that is several orders of magnitude higher than the mean delay, there is still a non-zero probability of this assumption being violated due to repetitive message

corruption and retransmission. When can such a system be claimed to be "100% correct"?

During the discussions, David Powell argued that the validity of any assumption should be quantified by an associated probability value (such as the probability of messages being delivered within the assumed maximum delay). He sketched a framework for so-called quantified proofs, in which the truth of properties is always expressed in probabilistic terms, and he explored the notion of "probabilistic validations". The relevance of this concept was disputed by several other participants, arguing that a proof is either correct or incorrect and that no probability should be involved in a demonstration.

Ken Birman argued that protocols have properties and corresponding probability distributions, parameterized by the time scale on which the system is examined. If one is concerned with real-time deadlines of minutes or tens of seconds, that is if the temporal scale is sufficiently coarse, asynchronous protocols are almost synchronous, and actually become real-time protocols whose performance is often better than synchronous protocols for the same model. Gerard LeLann disputed this view. According to him, one should not confuse performance of protocols for a particular implementation with the algorithmic nature of the protocols. He insisted on the fact that real-time issues arise whenever a specification contains direct or indirect references to physical time, and that real-time should not be defined with respect to the magnitude of the timeliness constraints (i.e., requirements that refer to physical time).

According to Gerard LeLann, the main problem that arises in a real-time environment is to schedule actions consistently with time constraints. Hard real-time raises proof obligations, which mainly consist in giving an upper bound on response times for all tasks under all possible situations. These proof obligations must be fulfilled under knowledge (about arrival laws of external events etc.) that is still limited by the information provided in the problem specification. In distributed systems, knowledge about the behavior of subcomponents and the arrival patterns of events is in general rather limited. Hence, correct scheduling algorithms for real-time distributed systems can only be pure on-line algorithms.

There was also a discussion on real-time communications. David Powell, Farnam Jahanian, and others had the view that polling or token-passing offer correct solutions to general real-time communications problems. Gerard LeLann disputed that view: Polling or token-passing inevitably rests on off-line decisions regarding scheduling of messages. Such decisions being known, it is easy to generate arrival scenarios whereby messages miss their deadlines, despite the fact that the problem might have solutions that can be found by a pure on-line protocol.

Finally, another issue was briefly addressed in one of the discussion sessions, raised by a contribution from Paulo Verissimo about clock-less protocols in real-time systems. Some participants insisted on the fact that either clocks (and some notion of global time) are used and relied upon at some system level or not. If this is not the case, then it is hard to see how a clock-less system can be shown to exhibit real-time properties.

List of Authors

Lecture Notes in Computer Science

For information about Vols. 1–865
please contact your bookseller or Springer-Verlag

Vol. 866: Y. Davidor, H.-P. Schwefel, R. Männer (Eds.), Parallel Problem Solving from Nature - PPSN III. Proceedings, 1994. XV, 642 pages. 1994.

Vol 867: L. Steels, G. Schreiber, W. Van de Velde (Eds.), A Future for Knowledge Acquisition. Proceedings, 1994. XII, 414 pages. 1994. (Subseries LNAI).

Vol. 868: R. Steinmetz (Ed.), Multimedia: Advanced Teleservices and High-Speed Communication Architectures. Proceedings, 1994. IX, 451 pages. 1994.

Vol. 869: Z. W. Raś, Zemankova (Eds.), Methodologies for Intelligent Systems. Proceedings, 1994. X, 613 pages. 1994. (Subseries LNAI).

Vol. 870: J. S. Greenfield, Distributed Programming Paradigms with Cryptography Applications. XI, 182 pages. 1994.

Vol. 871: J. P. Lee, G. G. Grinstein (Eds.), Database Issues for Data Visualization. Proceedings, 1993. XIV, 229 pages. 1994.

Vol. 872: S Arikawa, K. P. Jantke (Eds.), Algorithmic Learning Theory. Proceedings, 1994. XIV, 575 pages. 1994.

Vol. 873: M. Naftalin, T. Denvir, M. Bertran (Eds.), FME '94: Industrial Benefit of Formal Methods. Proceedings, 1994. XI, 723 pages. 1994.

Vol. 874: A. Borning (Ed.), Principles and Practice of Constraint Programming. Proceedings, 1994. IX, 361 pages. 1994.

Vol. 875: D. Gollmann (Ed.), Computer Security – ESORICS 94. Proceedings, 1994. XI, 469 pages. 1994.

Vol. 876: B. Blumenthal, J. Gornostaev, C. Unger (Eds.), Human-Computer Interaction. Proceedings, 1994. IX, 239 pages. 1994.

Vol. 877: L. M. Adleman, M.-D. Huang (Eds.), Algorithmic Number Theory. Proceedings, 1994. IX, 323 pages. 1994.

Vol. 878: T. Ishida; Parallel, Distributed and Multiagent Production Systems. XVII, 166 pages. 1994. (Subseries LNAI).

Vol. 879: J. Dongarra, J. Waśniewski (Eds.), Parallel Scientific Computing. Proceedings, 1994. XI, 566 pages. 1994.

Vol. 880: P. S. Thiagarajan (Ed.), Foundations of Software Technology and Theoretical Computer Science. Proceedings, 1994. XI, 451 pages. 1994.

Vol. 881: P. Loucopoulos (Ed.), Entity-Relationship Approach – ER'94. Proceedings, 1994. XIII, 579 pages. 1994.

Vol. 882: D. Hutchison, A. Danthine, H. Leopold, G. Coulson (Eds.), Multimedia Transport and Teleservices. Proceedings, 1994. XI, 380 pages. 1994.

Vol. 883: L. Fribourg, F. Turini (Eds.), Logic Program Synthesis and Transformation – Meta-Programming in Logic. Proceedings, 1994. IX, 451 pages. 1994.

Vol. 884: J. Nievergelt, T. Roos, H.-J. Schek, P. Widmayer (Eds.), IGIS '94: Geographic Information Systems. Proceedings, 1994. VIII, 292 pages. 19944.

Vol. 885: R. C. Veltkamp, Closed Objects Boundaries from Scattered Points. VIII, 144 pages. 1994.

Vol. 886: M. M. Veloso, Planning and Learning by Analogical Reasoning. XIII, 181 pages. 1994. (Subseries LNAI).

Vol. 887: M. Toussaint (Ed.), Ada in Europe. Proceedings, 1994. XII, 521 pages. 1994.

Vol. 888: S. A. Andersson (Ed.), Analysis of Dynamical and Cognitive Systems. Proceedings, 1993. VII, 260 pages. 1995.

Vol. 889: H. P. Lubich, Towards a CSCW Framework for Scientific Cooperation in Europe. X, 268 pages. 1995.

Vol. 890: M. J. Wooldridge, N. R. Jennings (Eds.), Intelligent Agents. Proceedings, 1994. VIII, 407 pages. 1995. (Subseries LNAI).

Vol. 891: C. Lewerentz, T. Lindner (Eds.), Formal Development of Reactive Systems. XI, 394 pages. 1995.

Vol. 892: K. Pingali, U. Banerjee, D. Gelernter, A. Nicolau, D. Padua (Eds.), Languages and Compilers for Parallel Computing. Proceedings, 1994. XI, 496 pages. 1995.

Vol. 893: G. Gottlob, M. Y. Vardi (Eds.), Database Theory – ICDT '95. Proceedings, 1995. XI, 454 pages. 1995.

Vol. 894: R. Tamassia, I. G. Tollis (Eds.), Graph Drawing. Proceedings, 1994. X, 471 pages. 1995.

Vol. 895: R. L. Ibrahim (Ed.), Software Engineering Education. Proceedings, 1995. XII, 449 pages. 1995.

Vol. 896: R. N. Taylor, J. Coutaz (Eds.), Software Engineering and Human-Computer Interaction. Proceedings, 1994. X, 281 pages. 1995.

Vol. 897: M. Fisher, R. Owens (Eds.), Executable Modal and Temporal Logics. Proceedings, 1993. VII, 180 pages. 1995. (Subseries LNAI).

Vol. 898: P. Steffens (Ed.), Machine Translation and the Lexicon. Proceedings, 1993. X, 251 pages. 1995. (Subseries LNAI).

Vol. 899: W. Banzhaf, F. H. Eeckman (Eds.), Evolution and Biocomputation. VII, 277 pages. 1995.

Vol. 900: E. W. Mayr, C. Puech (Eds.), STACS 95. Proceedings, 1995. XIII, 654 pages. 1995.

Vol. 901: R. Kumar, T. Kropf (Eds.), Theorem Provers in Circuit Design. Proceedings, 1994. VIII, 303 pages. 1995.

Vol. 902: M. Dezani-Ciancaglini, G. Plotkin (eds.), Typed Lambda Calculi and Applications. Proceedings, 1995. VIII, 443 pages. 1995

Vol. 903: E. W. Mayr, G. Schmidt, G. Tinhofer (Eds.), Graph-Theoretic Concepts in Computer Science. Proceedings, 1994. IX, 414 pages. 1995.

Vol. 904: P. Vitányi (Ed.), Computational Learning Theory. EuroCOLT'95. Proceedings, 1995. XVII, 415 pages. 1995. (Subseries LNAI).

Vol. 905: N. Ayache (Ed.), Computer Vision, Virtual Reality and Robotics in Medicine. Proceedings, 1995. XIV, 567 pages. 1995.

Vol. 906: E. Astesiano, G. Reggio, A. Tarlecki (Eds.), Recent Trends in Data Type Specification. Proceedings, 1995. VIII, 523 pages. 1995.

Vol. 907: T. Ito, A. Yonezawa (Eds.), Theory and Practice of Parallel Programming. Proceedings, 1995. VIII, 485 pages. 1995.

Vol. 908: J. R. Rao Extensions of the UNITY Methodology: Compositionality, Fairness and Probability in Parallelism. XI, 178 pages. 1995.

Vol. 909: H. Comon, J.-P. Jouannaud (Eds.), Term Rewriting. Proceedings, 1993. VIII, 221 pages. 1995.

Vol. 910: A. Podelski (Ed.), Constraint Programming: Basics and Trends. Proceedings, 1995. XI, 315 pages. 1995.

Vol. 911: R. Baeza-Yates, E. Goles, P. V. Poblete (Eds.), LATIN '95: Theoretical Informatics. Proceedings, 1995. IX, 525 pages. 1995.

Vol. 912: N. Lavrac, S. Wrobel (Eds.), Machine Learning: ECML – 95. Proceedings, 1995. XI, 370 pages. 1995. (Subseries LNAI).

Vol. 913: W. Schäfer (Ed.), Software Process Technology. Proceedings, 1995. IX, 261 pages. 1995.

Vol. 914: J. Hsiang (Ed.), Rewriting Techniques and Applications. Proceedings, 1995. XII, 473 pages. 1995.

Vol. 915: P. D. Mosses, M. Nielsen, M. I. Schwartzbach (Eds.), TAPSOFT '95: Theory and Practice of Software Development. Proceedings, 1995. XV, 810 pages. 1995.

Vol. 916: N. R. Adam, B. K. Bhargava, Y. Yesha (Eds.), Digital Libraries. Proceedings, 1994. XIII, 321 pages. 1995.

Vol. 917: J. Pieprzyk, R. Safavi-Naini (Eds.), Advances in Cryptology - ASIACRYPT '94. Proceedings, 1994. XII, 431 pages. 1995.

Vol. 918: P. Baumgartner, R. Hähnle, J. Posegga (Eds.), Theorem Proving with Analytic Tableaux and Related Methods. Proceedings, 1995. X, 352 pages. 1995. (Subseries LNAI).

Vol. 919: B. Hertzberger, G. Serazzi (Eds.), High-Performance Computing and Networking. Proceedings, 1995. XXIV, 957 pages. 1995.

Vol. 920: E. Balas, J. Clausen (Eds.), Integer Programming and Combinatorial Optimization. Proceedings, 1995. IX, 436 pages. 1995.

Vol. 921: L. C. Guillou, J.-J. Quisquater (Eds.), Advances in Cryptology – EUROCRYPT '95. Proceedings, 1995. XIV, 417 pages. 1995.

Vol. 922: H. Dörr, Efficient Grapg Rewriting and Its Implementation. IX, 266 pages. 1995. (Subseries LNAI).

Vol. 923: M. Meyer (Ed.), Constraint Processing. IV, 289 pages. 1995.

Vol. 924: P. Ciancarini, O. Nierstrasz, A. Yonezawa (Eds.), Object-Based Models and Languages for Concurrent Systems. Proceedings, 1994. VII, 193 pages. 1995.

Vol. 925: J. Jeuring, E. Meijer (Eds.), Advanced Functional Programming. Proceedings, 1995. VII, 331 pages. 1995.

Vol. 926: P. Nesi (Ed.), Objective Software Quality. Proceedings, 1995. VIII, 249 pages. 1995.

Vol. 927: J. Dix, L. Moniz Pereira, T. C. Przymusinski (Eds.), Non-Monotonic Extensions of Logic Programming. Proceedings, 1994. IX, 229 pages. 1995. (Subseries LNAI).

Vol. 928: V.W. Marek, A. Nerode, M. Truszczyński (Eds.), Logic Programming and Nonmonotonic Reasoning. Proceedings, 1995. VIII, 417 pages. 1995. (Subseries LNAI).

Vol. 929: F. Morán, A. Moreno, J.J. Merelo, P. Chacón (Eds.), Advances in Artificial Life. Proceedings, 1995. XIII, 960 pages. 1995 (Subseries LNAI).

Vol. 930: J. Mira, F. Sandoval (Eds.), From Natural to Artificial Neural Computation. Proceedings, 1995. XVIII, 1150 pages. 1995.

Vol. 931: P.J. Braspenning, F. Thuijsman, A.J.M.M. Weijters (Eds.), Artificial Neural Networks. IX, 295 pages. 1995.

Vol. 932: J. Iivari, K. Lyytinen, M. Rossi (Eds.), Advanced Information Systems Engineering. Proceedings, 1995. XI, 388 pages. 1995.

Vol. 933: L. Pacholski, J. Tiuryn (Eds.), Computer Science Logic. Proceedings, 1994. IX, 543 pages. 1995.

Vol. 934: P. Barahona, M. Stefanelli, J. Wyatt (Eds.), Artificial Intelligence in Medicine. Proceedings, 1995. XI, 449 pages. 1995. (Subseries LNAI).

Vol. 935: G. De Michelis, M. Diaz (Eds.), Application and Theory of Petri Nets 1995. Proceedings, 1995. VIII, 511 pages. 1995.

Vol. 936: V.S. Alagar, M. Nivat (Eds.), Algebraic Methodology and Software Technology. Proceedings, 1995. XIV, 591 pages. 1995.

Vol. 937: Z. Galil, E. Ukkonen (Eds.), Combinatorial Pattern Matching. Proceedings, 1995. VIII, 409 pages. 1995.

Vol. 938: K.P. Birman, F. Mattern, A. Schiper (Eds.), Theory and Practice in Distributed Systems. Proceedings,1994. X, 263 pages. 1995.

Vol. 939: P. Wolper (Ed.), Computer Aided Verification. X, 451 pages. 1995.

Vol. 941: M. Cadoli, Tractable Reasoning in Artificial Intelligence. XVII, 247 pages. 1995. (Subseries LNAI).

Vol. 942: G. Böckle, Exploitation of Fine-Grain Parallelism. IX, 188 pages. 1995.

Vol. 943: W. Klas, M. Schrefl, Metaclasses and Their Application. IX, 201 pages. 1995.