International Workshop on Processing Declarative Knowledge

Cristina Ribeiro

■ *The International Workshop on Processing Declarative Knowledge was held in Kaiserslautern, Germany, from 1 to 3 July 1991. The workshop was intended as a forum for the presentation of new approaches to processing declarative knowledge, the discussion of procedural versus alternative paradigms, and the issues concerned with efficient processing of realistic knowledge bases. Demonstrations of implemented systems were also announced.*

The International Workshop on Processing Declarative Knowledge was held in Kaiserslautern, Germany, from 1 to 3 July 1991. It was organized by the German Research Center for Artificial Intelligence (DFKI) in cooperation with the Association for Logic Programming and the Gesellschaft fuer Informatik. The workshop was hosted by DFKI at the pleasant campus of Kaiserslautern University. The Program Committee, chaired by M. Richter from DFKI, gathered a significant sample of well-known researchers from both the AI and the logic programming communities.

The workshop was intended as a forum for the presentation of new approaches to processing declarative knowledge, the discussion of procedural versus alternative paradigms, and the issues concerned with efficient processing of realistic knowledge bases. Demonstrations of implemented systems were also announced.

The 25 selected papers were distributed among eight sessions in the following areas: concept languages, semantics, applications, Prolog extensions, transformation, and abstract machines. A strong component on extensions of the logic programming paradigm and language implementations was manifest in the 12 system demonstrations available. Most of these demonstrations corresponded to presented papers.

One hundred five people attended the workshop, the majority from several research centers in Germany.

There were also people from a large number of other countries, including France, The Netherlands, the United Kingdom, Belgium, Italy, Sweden, USSR, Switzerland, Portugal, Egypt, Finland, the United States, Romania, and Australia.

Three invited talks covered distinguished aspects of declarative knowledge processing. Carlo Zaniolo (MCC Corporation, Austin, Texas) gave a talk entitled "Efficient Processing of Declarative Rule-Based Languages for Databases." Pascal van Hentenryck (Brown University) presented the talk "Constraint Logic Programming," and Andrew Taylor (University of Sydney) presented "High-Performance Prolog Implementation through Global Analysis."

Moss argued that…an integration of Prolog and a procedural language can be done in an object-oriented manner.

"Declarative and Procedural Paradigms—Do They Really Compete?" was the panel organized by H. Boley (DFKI); it provided a lively discussion arena for H. Ait-Kaci (Digital Equipment Corporation, Paris), C. Moss (Imperial College), M. Meier (European Computer-Industry Research Center, Munich), Richter, and A. Voronkov (International Laboratory of Intelligent Systems, Novosibirsk, USSR). The main discussion concerned whether Prolog, as a widespread declarative language, is expressive and efficient enough.

There seemed to be the general feeling that to obtain efficiency, one needs to use the nondeclarative features of the language. Moss argued that to be both declarative and efficient, an integration of Prolog and a procedural language can be done in an object-oriented manner. Ait-Kaci maintained that declarative and procedural notions are relative insofar as one language can only be said to be more or less declarative than another. He also argued that because conventional mathematics cannot elegantly account for simple intuitive and natural procedural constructs, their usefulness or even their correctness should by no means be discounted. He concluded that a healthier and more productive attitude must be to work toward elaborating a more appropriate mathematics to explain programming in terms that coincide with the lay person's perception and real-life needs.

The polemic initial position of Meier was that Prolog is not efficient and that most Prolog programs are also not declarative. The solution he proposed was to give up low-level programming and write programs at the highest possible level of abstraction, leaving the choice of appropriate processing to the system. It was pointed out that increasing efficiency at such a high level of program description is a complex matter. Y. Caseau (Bell Communications, Morristown, New Jersey) pointed out that in the development of large applications, declarative solutions can effectively be better because little effort is put into code optimization when people are developing software under pressure, even when using a procedural approach.

A special talk by Moss entitled "Commercial Applications of Large Prolog Knowledge Bases" concluded the workshop. He reported on a surprisingly large number of currently used database-sized applications developed in Prolog.

The workshop was a successful event, bringing together researchers with a broad spectrum of interests and allowing plenty of time for discussion and system demonstrations. The local arrangements were fine, thanks to H. Boley and the DFKI team. There was time for a very pleasant hike through the Palatinate Forest, starting at the university buildings and ending at a fine

integrating AI and software engineering, a main goal of the KBSE community. After describing a success story from the Gulf War (the use of the dynamic analytic replanning tool [DART]), Boehm concluded with a description of the KBSE challenges. These challenges include knowledge capture (acquisition) for adding real value to intelligent systems, knowledge updating, knowledge representation choice guidelines to help spread the technology, and knowledge base–project database interoperability. The biggest challenge of all, according to Boehm, is the scalability challenge: making sure our techniques and technology can truly scale up to large, real-world problems.

During the rest of the conference, 25 papers were presented in nine sessions: General Design Issues, Requirements, Automatic Programming, Reuse, KBSA, Debugging, Reverse Engineering, Learning, and Process Support. These papers were selected from 56 that were submitted to the conference and reflected the state of the field as follows: First, there is a lot of activity in the support of requirement and specification activity. Second, domain modeling and domain knowledge are increasingly viewed as important (this fact was revealed in a panel presentation as well). Third, although the lack of maturity of the field is reflected in the "toyness" of many of the problems, there are now some significant success stories. Rather than report on some or all of these presentations in detail, the reader is invited to examine the proceedings (Proceedings 1992).

Panels and demonstrations were also part of the conference. The three panels were "Domain Modeling," moderated by Neil Iscoe of EDS, Incorporated; "Encouraging Adaptation of KBSA and KBSE Technology," moderated by Bill Sasso of Andersen Consulting; and "Knowledge-Based Design," moderated by Michael Lowry of the Kestrel Institute. The panels were a combination of presentations by the panel members and questions from the audience. Demonstrations included both academic prototypes and commercially available systems.

The final day of the conference featured a plenary address by Thomas Cheatham of Software Options and Harvard University. Cheatham, one of the authors of the original KBSA report, described computing in the 1990s as being characterized by networks of workstations, with an occa-

sional supercomputer; multiple simultaneous users; life-cycle support; and, finally, activity coordination or process support. He then described an artifact-based software engineering environment being developed at Software Options called the E-L (environment language) system. The E-L system is based on the management of small typed artifacts, which, along with a coarse structure of relations of reference, predecessor, and successor, can describe both a software system and the process and environment used to generate it. The E-L system coordinates activity by multiple users over a large, interrelated base of artifacts and projects to provide activity distribution, communication, persistence, modularity, viewing, and extensibility.

KBSE-91 ended with an open discussion of the conference and the goals and status of the KBSE community. A number of suggestions were raised, from generating better publicity and getting wider European participation to having more formal tutorials and thinking hard about the education issues in generating and using KBSE technology. The primary technical concern was the maturing of KBSE work into more realistic domains so that knowledge-based techniques and technology could be evaluated properly.

Planning, organization, and implementation of future KBSE conferences will be facilitated by the formation of a permanent KBSE steering committee. Lewis Johnson of USC Information Sciences Institute is both next year's conference chair and program chair. Donald Yu of UNISYS will serve as local arrangement chair. The conference will be held in the Washington, D.C., area in mid- to late September 1992. For more information about this conference, send electronic mail to kbse7-request@cs.rpi.edu.

IEEE Expert started a special track on KBSE. This special track should appear early in 1992 with a number of the best papers from KBSE-91; subsequently, the track will be open to general submissions. For more information, contact Peter G. Selfridge, AT&T Bell Laboratories, Murray Hill NJ 07974, pgs@research.att.com.

References

Green, C.; Luckham, D.; Balzer, R.; Cheatham, T.; and Rich, C. 1986. Report on a Knowledge-Based Software Assistant.

In *Readings on Artificial Intelligence and Software Engineering*, eds. C. Rich and R. Waters, 377–428. San Mateo, Calif.: Morgan Kaufmann.

Proceedings of the Sixth Knowledge-Based Software Engineering Conference (KBSE-91). 1992. Washington, D.C.: IEEE Computer Society.

Peter G. Selfridge received his Ph.D. in computer science from the University of Rochester in 1982. Since then, he has been at AT&T Bell Laboratories and is currently in the AI Principles Research Department. Selfridge's main interest is in the application of knowledge representation technology to problems in large-scale software development, maintenance, and understanding.

Advances in Interfacing...

Milind Tambe received his M.S. in computer science from the Birla Institute of Technology and Science, Pilani, India, in 1986 and his Ph.D. in computer science from Carnegie Mellon University (CMU) in May 1991. He is currently a research associate at CMU's School of Computer Science. His interests are in the areas of integrated AI architectures and efficiency and parallelism of AI programs. He is a member of the American Association for Artificial Intelligence.

Processing Declarative Knowledge

restaurant where the conference banquet took place.

Proceeding preprints were available at the workshop, and book proceedings will be published as part of Springer-Verlag's Lecture Notes on Artificial Intelligence (Boley, H., and Richter, M. M., eds. 1992. *Processing Declarative Knowledge International Workshop (PDK '91)*. Berlin: Springer-Verlag. Forthcoming.

Cristina Ribeiro has a B.S. in electrical engineering from Porto University and an M.S. in electrical engineering from IST, Lisbon Technical University. She is currently working on a Ph.D. in AI at Lisbon New University. Her current research interests are temporal reasoning, logic programming, and constraint solving.