New Programs at DARPA and NSF

Xiaolei Qian
Computer Science Laboratory, SRI International
333 Ravenswood Avenue, Menlo Park, CA 94025
qian@csl.sri.com

Abstract

We will share with readers some good news on NSF and Defense budget, and report on several interesting new programs at DARPA and NSF.

1 NSF and Defense Budget in Good Shape

In the FY97 budget that passed Congress in September, both NSF and Defense fared pretty well. NSF got $3.27B, compared with $3.22B last year. The funding for Research and Related Activities, including Computing and Information Science and Engineering, got a substantial boost, $118M more than last year. Defense research funding increased more than 7%, with the major beneficiaries being basic research (6.1 funds) and Balistic Missile Defense.

2 New Collaborative Research Initiatives from NSF

NSF's Computer and Information Science and Engineering (CISE) Directorate is introducing 3 new initiatives in collaborative team research. Teams can represent different subactivities of computing research as well as other scientific or engineering disciplines. CISE will also consider for joint funding multidisciplinary proposals that are submitted to programs in other Directorates and that include computing research. Award sizes are expected to range from $250,000 to 600,000$ over 2 to 4 years. Workshops and 12-month planning grants will also be considered. The three new initiatives are as follows (Further information will be available at www.cise.nsf.gov.):

Experimental Software Systems (ESS) "The ESS program will support projects with significant experimental content, conducted by small teams of investigators. Appropriate subjects for the research may include experiments in software engineering, prototyping and evaluation of novel software systems, experiments in distributed and parallel computing, or other topics. The program will build on earlier, successful research that may have been conceptual or theoretical in nature, and support efforts to evaluate the most promising results." The contact for this program is Dr. William Agresti, (703) 306-1911, wagresti@nsf.gov.

Special Projects in Networking and Communications "The projects, typically submitted by small groups, will include researchers in networking and/or communications and researchers from other computing areas such as operating systems, databases, software environments, architecture or from the social sciences, such as economics, psychology or sociology." For more details, contact Dr. Darleen Fisher, Division of Networking and Communications Research and Infrastructure (703) 306-1949, dlfisher@nsf.gov.

Challenges in Computer and Information Science and Engineering These challenges "will provide support for research that is focused on a single problem, the solution of which requires scientific advances across two or more subactivities of computer science and engineering. Although successful proposals submitted to this initiative are expected to have a primary focus on computing research, proposals including participation of disciplines other than computer science and engineering are also encouraged." More information can be obtained from Dr. Robert Voigt, (703) 306-1900, rvoigt@nsf.gov.

3 DARPA Interested in High-Performance Knowledge Bases

The Defense Advanced Research Projects Agency (DARPA) has announced a 4-year program for research and development in the area of High-Performance Knowledge Bases (HPKB). The goal of this new program is to "produce the technology needed to en-

able system developers to rapidly (within months) construct large (100K-1M axiom/rule/frame) knowledge-bases that provide comprehensive coverage of topics of interest, are reusable by multiple applications with diverse problem-solving strategies, and are maintainable in rapidly changing environments." The construction of large, comprehensive, reusable, and maintainable knowledge bases would involve three major steps:

1. "Building Foundation Knowledge: creating the foundation knowledge (e.g., selecting the knowledge representation scheme, assembling theories of common knowledge, defining domain-specific terms and concepts) to enable the construction and population of large, comprehensive knowledge bases for particular domains of interest – by selecting, composing, extending, specializing, and modifying components from a library of reusable ontologies, common domain theories, and generic problem-solving strategies."

2. "Acquiring Domain Knowledge: constructing and populating a complete knowledge base – by using the foundation knowledge to generate domain-specific knowledge acquisition, data mining, and information extraction tools – to enable collaborating teams of domain (non-computer) experts to easily extend the foundation theories, define additional domain theories and problem solving strategies, and acquire domain facts to populate a comprehensive knowledge base covering the domains of interest."

3. "Efficient Problem Solving: enabling efficient problem solving – either by providing efficient inference and reasoning procedures to operate on a complete knowledge base, or by providing tools and techniques to select and transform knowledge from a complete knowledge base into optimized problem-solving modules tailored to the unique requirements of an application."

DARPA intends to produce, by the end of FY00, integrated knowledge-base development environments to be used in multiple DARPA application projects, in particular, Dynamic Multi-user Information Fusion (DMIF), Joint Task Force Advanced Technology Demonstration (JTFATD), Technology Development for the Joint Force Air Component Commander (JFACC), Advanced Logistics Program (ALP), Battlefield Awareness Data Dissemination (BADD), and Information Gathering, Processing and analyses in Support of Crisis Management (Project Genoa).

DARPA is soliciting proposals in two categories, one being the development and integration of individual knowledge base technologies and environments, and the other being the development and evaluation of challenge problems to be used to evaluate and test the technologies and environments. In the first category, proposals are solicited both for individual technology efforts to create the foundation-building, knowledge-acquisition, and problem-solving technology, and also for larger integration efforts to produce integrated knowledge-base development environments in which the individual technologies have been combined into cohesive tool sets. The areas of interest for individual technology efforts are summarized below.

Foundation-Building Technology DARPA is interested in new techniques for "creating, maintaining, and manipulating large libraries of reusable ontologies, common domain theories, and generic problem-solving strategies to enable the creation of foundation knowledge for constructing and populating of large, comprehensive knowledge bases for particular domains of interest – by selecting, composing, extending, specializing, and modifying components from those libraries. This would include developing:

- Tools and techniques to define, test, verify, validate, and edit large libraries of ontologies, domain theories, and problem-solving strategies;
- Tools and techniques to extend and specialize those ontologies, domain theories, and problem-solving strategies to create new foundation knowledge;
- Tools and techniques to compare and compose independently developed ontologies, domain theories, and problem-solving, especially to compare independently developed or modified theories to detect conflicting information;
- Tools and techniques to translate ontologies, domain theories, and problem-solving strategies among different knowledge representation languages and frameworks;
- Tools and techniques to import data dictionaries, data schemas, lexicons, and legacy expert systems, other source material to create initial ontologies and skeleton theories for elaboration and completion;
- Human-computer interaction and visualization techniques to enable easy browsing, visualization, and understanding of large libraries of ontologies, domain theories, and problem-solving strategies;
- Knowledge Representation (KR) concepts, languages, and techniques needed to represent knowledge for reuse by multiple applications with diverse

problems solving strategies and for reasoning at different levels of abstraction – to the extent these representation techniques are needed to implement the knowledge-base editing and manipulation tools described above.

- Development of libraries of common ontologies and common domain theories for space, time, actions and processes, belief and uncertainty, simple physical substances and behavior, and other commonly used knowledge domains."

Knowledge-Acquisition Technology DARPA is interested in new techniques for "using foundation knowledge to generate domain-specific knowledge acquisition, data mining, and information extraction tools to enable collaborating teams of domain (non-computer) experts to build and maintain comprehensive knowledge bases by extending the foundation knowledge, defining additional domain theories and problem-solving strategies, and acquiring domain facts. This would include developing:

- Tools and techniques to generate, from foundation knowledge, domain-specific knowledge-base editors and visualization tools to enable groups of domain experts, with little or know knowledge-base or computer experience, to easily build large, comprehensive knowledge bases by extending the foundation theories, defining additional domain theories and problem-solving strategies, and acquiring domain facts;

- Machine learning and data mining tools and techniques to automatically discover and acquire new knowledge from a variety of information sources, including databases, textbooks, dictionaries, and observations of activities and behavior, while reducing the size of the training data sets required for effective learning;

- Human language understanding and information extraction tools and techniques to automatically discover and extract new knowledge from a variety of text documents (manuals, textbooks, dictionaries, reports, web pages, messages, and other informal documentation) or from natural language dialogs with domain experts;

- Tools and techniques to help collaborating teams of domain experts define new knowledge, compare independently developed ontologies and domain theories, and merge independently developed knowledge into coherent and consistent theories;

- Tools and techniques to support interactive, mixed-initiative knowledge definition and machine learning, from both human and computer sources;

- Tools and techniques to provide continuous knowledge acquisition, error checking, anomaly detection, and knowledge maintenance over the life-cycle of a knowledge base."

Efficient Problem-Solving Technology DARPA is interested in new techniques to "provide efficient problem solving machinery for large knowledge bases – either by providing efficient reasoning and inference procedures to operate on large knowledge bases, or by providing techniques to select and transform knowledge from a complete knowledge base into optimized problem-solving modules tailored to the unique requirements of the application. This would include developing:

- Tools and techniques to provide efficient reasoning and inference procedures for large knowledge bases;

- Tools and techniques to select, extract, transform, and compile knowledge from a large knowledge base into application specific problem solving modules tailored to the unique requirements of a particular problem-solving strategy or level of abstraction for a particular application;

- Tools and techniques to enable the translation, packaging, and integration of problem solving modules into diverse application environments;

- Integration of the above technologies into an integrated knowledge-base development environment."

Funding for the program is estimated at approximately $8M per year for four years. About $3-4M per year will be available for individual technology developments, $2-4M for integrated development environments, and $1-2M per year on challenge problem developments. Individual technology efforts will range from $200K to $600K per year, and integration efforts will range from $1M to $2M per year. Full proposals are due 12/02/96. Further information can be obtained at URL: http://www.darpa.mil/baa/#iso.

4 Survivability of Large Scale Information Systems

DARPA announced a second 3-year program on the survivability initiative (see the March 1996 issue of this column for excerpts of the first program). This program emphasizes the distinct issue of "survivability", meaning

that a large scale information system as a whole continues to perform its critical functions even when individual elements are compromised. Proposals are sought for research into three areas.

Detecting, Limiting, and Preventing Attacks for large scale computational "systems of systems". Topics of interest include "intrusion detection, diagnosis of the pattern of a coordinated attack, mechanisms for obtaining additional information about an attack, and mechanisms for secure, robust and authenticated coordination and communication between elements of the overall system."

Adaptive Architectures for Survivable Systems Survivable systems must continue performing their critical functions after sustaining loss of computational resources and data. Research is needed into "organizing large scale systems of systems so that remaining resources are preferentially allocated to more important tasks. [...] Possible solutions include (but are not limited to):

1. Economic mechanisms that regulate the allocation of services to applications by establishing a price for individual computational services and budgets for potential consumers, and
2. Locally adaptive mechanisms which achieve overall policy objectives as emergent behavior by guiding the choice to tasks to be performed by each individual computational resource."

Techniques for Creating Variability and Diversity In current computational systems, there is a high degree of dependence on a small number of operating systems, networking environments, etc. This means that a single flaw is replicated in virtually all computational elements. Research is needed into affordable techniques to engender variability within the computational mix, such as dynamic runtime optimization, alternative implementations of common services and APIs, and different copies of the same software by varying the memory layout or compiled code formats.

Proposal abstracts were due 10/18/96, and full proposals are due 01/10/97. For further details of the program, send email to baa9640@darpa.mil or browse URL: http://www.ito.darpa.mil/Solicitations.html.

5 DARPA Project Genoa

In July 1996, DARPA announced a program on information gathering, processing, and analyses in support of crisis management (Project Genoa). Although the deadline for proposal submission has already passed, the database requirements of the program might be of interest to SIGMOD readers.

The program seeks research and new technology development related to systems to assist crisis managers at the highest decision levels. The objectives of the program are to

- reduce to near zero the formation time for tailored, interdisciplinary, crisis action teams;
- provide crisis action officers access in seconds to relevant data using a consistent user interface across multiple data bases of mixed classification;
- develop methodologies for preparing templates to describe triggering events which can then be used by intelligent agents to automatically mine data bases rapidly and provide pointers to relevant data;
- provide tools for rapid quantitative assessment of emerging crises and long-term analysis;
- reduce the time to a few hours for producing a coherent summary of emerging crisis situations, initial options, reasons for differences, and description of uncertainties for presentation to decision makers in an interactive (and later a multimedia) environment;
- provide capability for a decision maker to drill down through the data to answer specific questions or check assumptions;
- development of crisis prototypes in an explicit and quantifiable manner which would provide some predictive capability based on a complex mix of economic, social, climatological, and/or political factors.

Four areas of interest have been identified: CrisisNet, CrisisBrowse, CrisisScope, and CrisisBrief. CrisisBrowse is of particular interest to database researchers. "CrisisBrowse is an application to run in the environment of the CrisisNet to provide seamless access to hundreds of heterogeneous data bases with a single standard request format for all data bases. The minimum requirements for CrisisBrowse are:

1. Seamless interface to multiple data bases and ability to add access to new data bases as they become available or relevant without conventional programming;
2. Option to return data or pointers to data;
3. Response time of minutes to search hundreds of data bases with thousands of new records per day;

4. Ability to retrieve and display records from multimedia sources such as messages, statistical records, published text, images, maps, motion pictures, video and audio recordings (and possibly real time transmissions);

5. Facility for knowledge rovers, mediators, and intelligent agents containing the user's goals and queries to roam the data bases defined for the CrisisNet, searching for, negotiating, accessing, processing and collecting information and data for return to the user; an option for a near real time monitoring function once a scan of historical data has been made should be considered;

6. Exploration of new, innovative, unique ways of displaying masses of data in ways that present visual clues of what might be important about the data for further investigation by the user;

7. Capability of making filtered sets of data available to other users in keeping with the collaborative nature of crisis management; this feature should not require retransmitting large volumes of data unless this is desired; consider creation of web pages on the fly for publication on the CrisisNet; consider voice annotation of data sets;

8. CrisisBrowse should have minimal impact on current operations of existing data bases."

The FY97 funding level is estimated at $6M. Proposals were due 09/16/96. The full text of the announcement can be found at URL: http://www.darpa.mil/baa/.