Planning Grant Objective

Objective

This planning grant proposal seeks support for a meeting with potential industrial partners in the Dallas/Fort Worth (DFW) Metropolitan area. In addition to UNT, two other universities in the DFW area (Southern Methodist University and University of Texas at Dallas) are seeking to establish Research Sites and join with an NSF Embedded System Industry-University Cooperative Research Center (I/UCRC) that is currently being formed with Arizona State University as the lead institution. The academic institutions will collectively organize a meeting. The meeting using the planning grant funds will be an important step in further developing and transforming an existing consortium among the universities into I/UCRC research sites to undertake and promote research and development of service-oriented, network-centric systems. The establishment of new research sites at Dallas/Ft. Worth metroplex (DFW) universities will extend the capabilities of the I/UCRC led by Arizona State and enhance the participating universities' ability to open doors to many industrial sponsors, as well as obtain funding from other federal and state agencies. The NSF I/UCRC program will provide the needed infrastructure, organization, and evaluation mechanism to ensure the growth and success of the research cites and the consortium.

Strategy

The collective vision of the DFW area research sites is that they will provide a primary source for fundamental systems research for the modeling, analysis, design, implementation, verification and validation, deployment, and evolution of net-centric systems. The sites will enable coordinated hardware/software and systems-engineering research and development as well as education and training of US citizens to meet our nation's future software and systems workforce needs.

The DFW research sites will also establish a leading research alliance that is capable of conducting significant research projects for government and industrial customers. By joining the forces of the participating academic institutions and high tech companies, the NSF Embedded Systems I/UCRC will greatly enhance the research capabilities of the participants and revolutionize our national research competence. As a leading technology innovator, a technology incubator, and a center for technology commercialization, it will contribute significantly to the economic growth of the nation.

The DFW universities along with Arizona State University have long-standing collaborations with industrial partners in terms of research and technology workforce development. These collaborations range across software license agreements, industry-funded research, faculty consulting, student-internships, scholarships, monetary gifts, and equipment and software donations. Research collaborations among faculty from these institutions in terms of joint research projects, co-authored publications, and jointly supervised theses and dissertations already exist.

The planning grant provides an opportunity to explore the establishment of a UNT Research site of a NSF Industry/University Cooperative Research Center and to expand participation to additional industrial concerns. The research of the UNT site (and the companion sites at SMU and UTD) focuses on the industry that is involved in the design of highly dependable net-centric software/hardware systems and on the industry that needs such systems. This includes software developers and users in defense, energy, transportation, health care, homeland security, and emergency preparedness/emergency response.

During the planning period, the UNT research site (along with those at UTD and SMU) will focus on several short-term research needs identified in collaboration with prospective industrial partners, including Lockheed Martin, Raytheon, Tektronix and Metallect. In addition, we have based our projects on the needs several industries, among them the military, national security, transportation, and health care. The immediate goal is to investigate technologies for the rapid development of communication layer protocols, including sensor networks, data compression protocols, and network monitoring infrastructures. We also plan to identify and analyze a set of important net-centric applications and investigate the research needs of these systems that require extensions of current technologies or demand novel techniques that are beyond the current state of the practice. We will use these short-term projects as demonstration projects in our efforts to recruit industrial partners to the proposed research sites.

Potential Members

At present, there exists a consortium of the major universities in the Dallas/Fort Worth Metropolitan area exists. The Net Centric Software Consortium includes as its founding members, the University of North Texas, the University of Texas at Arlington, the University of Texas at Dallas, and Southern Methodist University, Lockheed Martin Aero, Raytheon, and Metallect. Other industrial concerns, including Boeing, Burlington Northern – Santa Fe (BNSF), Dell, EDS, Honeywell, Masergy, Nokia, Nortel, Rockwell-Collins, Samsung, Tektronix, Texas Instruments, and UGS-Siemans, as well as several smaller companies have expressed interest in joining the consortium. The establishment of research sites at 3 DFW area universities with a NSF I/UCRC will bring prestige to the consortium and enhance the participating universities' ability to open doors to many industrial sponsors, as well as in obtaining funding from other federal and state agencies. The active role played by the NSF staff will be invaluable for the success of our center.

Meeting Planning Arrangements

The three DFW area universities (UNT, UTD and SMU) will use the planning grant funds to organize a joint meeting to recruit industrial partners to the their research sites. A meeting in late March or early April 2008 in the Dallas-Fort Worth metroplex is being planned for this purpose. The meeting will be organized along the lines suggested by Gray and Walters[1] (a tentative agenda for the meeting is included in supplemental documents). The meeting will bring together potential industrial members, university administrators, center and site directors, faculty researchers, and students involved in demonstration projects. We will invite at least 50 potential industrial partners to the meeting. The industrial members will be presented with the governance structure of the center, research thrusts, and results from our short-term demonstration projects, and industrial agreements. The governance structure and industrial partner agreements will be in accordance with the NSF Embedded Systems I/UCRC that is being formed with Arizona State University as the lead institution. The details about the meeting organization, responsibilities of staff and presenters, and draft agenda can be found in the "Draft Agenda" among the supplementary documents for this proposal

[1] Dennis Gray and George Walters: Managing the Industry/University Cooperative Research Center, Battelle Press, 1998.

Project Description

1. Overview

1.1. Industrial Needs and Our Research Focus

The UNT research site (along with research sites being proposed by SMU and UTD) aims to develop effective tools and techniques for rapidly developing highly dependable, and adaptable net-centric systems for safety-critical and/or mission-critical applications, such as command and control systems, emergency preparedness infrastructure, large-scale medical operations, health monitoring, and healthcare systems. These applications typically harness a large number of autonomous subsystems, including sensors and actuators, remote-controlled mobile units, embedded systems, data storage and information processing entities, control centers; interconnected via Internet, wireless, and satellite networks, that collaborate to achieve critical tasks. A major characteristic of net-centric systems is that they are deployed in dynamically changing environments and they must constantly evolve to meet changing goals. Currently, the service-oriented architecture (SOA) technique, which allows autonomy in individual subsystems and supports the infrastructure for asynchronous collaboration among the entities, is considered a promising paradigm for modeling, developing, and rapidly deploying net-centric systems. In the SOA based net-centric paradigm, systems are no longer designed with a fixed set of capabilities but as a set of services that will be dynamically acquired or created, replaced, composed, verified, and validated in real-time in the field without human intervention. Due to the mission-critical nature of many emerging net-centric applications, these systems must always be available, reliable, dependable, and fault-tolerant, while simultaneously providing high security assurance. However, the researchers at the research sites of the I/UCRC will explore alternate paradigms, including agent-oriented architectures, for implementing net-centric systems.

The DFW research sites of the Embedded Systems I/UCRC will be the primary source of collaborative research in the design, implementation, and operational support of net-centric application systems. The center and the research sites will enable coordinated software engineering and systems research and development as well as education and training to meet future technology workforce needs of our nation. The research sites will undertake both basic and applied research needed for net-centric systems. More specifically, the research sites will develop research and associated technologies for highly dependable, rapidly composable, fully analyzable net-centric systems. To achieve these objectives, the research focuses on several layers of net-centric systems, including the communication layer, the service layer, and the application layer. At the communication layer, the research sites seek to develop rapid integration techniques that select and assemble existing wired and wireless technologies and automatically resolve interoperability issues to support real-time communication protocols among the large number of subsystems and services in net-centric applications. At the service layer, we will develop advanced research for 1) automated service assembly for net-centric systems that can be verified for functional, performance, and security compliance; 2) continuous monitoring and adaptation to tolerate unexpected events and adapt to dynamically evolving requirements; 3) verification and validation techniques that can assess the system dependability, survivability, and performance to a high degree of confidence. At the application layer, we will develop appropriate framework for rapid development of highly survivable and user-friendly net-centric systems in a variety of application domains.

During the planning period, the UNT Research Site (along with companion sites at SMU and UTD) will focus on several short-term research needs identified in collaboration with prospective industrial partners, including Lockheed Martin, Raytheon, Tektronix and Metallect. In addition, we have based our projects on the needs several industries, among them the military, national security, transportation and health care. The immediate goal is to investigate technologies for the rapid development of communication layer protocols, including sensor networks, data compression protocols, and network monitoring infrastructures. We also plan to identify and analyze a set of important net-centric applications and investigate the research needs of these systems that require extensions of current technologies or demand novel techniques that are beyond the current state of the practice.

1.2. General Policies, Guidelines, Organizational Structure, and Operational Procedures

The UNT Research Site will be a part of the Embedded Systems I/UCRC. The current director of the Center is Professor Sarma Vrudhula of Arizona State. Krishna Kavi will be the Site Coordinator of the UNT Research Site. The industrial partners of UNT Research Site will join the existing Industrial Advisory Board of the Embedded Systems I/UCRC. Together, the IAB will evaluate proposed and on-going projects. UNT will be responsible for accounting, invoicing and auditing of the costs associated with the research and operation of the UNT site. The UNT research site will adapt the organizational structure and procedures of the I/UCRC (which follows the procedures described in Gray and Walters[2]).

1.3. UNT Site Coordinator's Qualification

Krishna Kavi, Chair, Department of Computer Science and Engineering, University of North Texas, will be the Site Coordinator of the UNT Research Site. In addition to being a chair of a department with 20 faculty members, 5 staff members, 600 undergraduate students, 100 MS and 30 PhD students, Kavi was a program manager at NSF for two years (1993-1995), managing the Operating Systems and Compilers program within CISE, CCR division. He was the Eminent Scholar Chair professor of Computer Engineering at the University of Alabama, Huntsville for 4 years (1997-2001). Kavi has supervised 12 PhD dissertations and nearly 40 MS theses. Four of his doctoral students have been involved in startup IT companies in the US and Taiwan.

Kavi's research has been supported by both governmental agencies (NSF, NASA, DoD) and industrial concerns. Texas Instruments, Motorola, General Dynamics (now Lockheed Martin), Sigma Tech. of Huntsville, Alabama, Recognition Equipment (acquired by General Motors), and several other small companies are among the industrial concerns that have supported his research. He was a researcher on 3 SBIR efforts. One such effort with Intravurt of San Jose, California led to the development of very effective intrusion detection software. Intravurt was acquired by McAfee in 2002. His work with a former PhD student on visual requirements has led to user friendly software that is currently being used in Taiwan to develop educational games for elementary and secondary school children. Kavi has also served as a technical expert on several patent infringement cases. Involvement in these cases has provided him with invaluable knowledge regarding IP and patent issues.

1.4. Management and Staffing Plan

The UNT Site Coordinator will work with the Director of the I/UCRC, Sarma Vrudhula, who will be responsible for all aspects of the Center operations and has the primary responsibility for administering the NSF award. The Director is also responsible for implementing the recommendations of the Advisory Board, for promoting the Center activities and recruiting new members. UNT Site Coordinator (Krishna Kavi) will be responsible for the administrative activities and technical research development at the University of North Texas. He is also responsible for coordinating inter-university research, aiding the center Director in marketing the center activities, and recruiting new members to UNT Site.

The Advisory Board (IAB) is composed of representatives of industrial sponsors. The industrial members of the UNT Research Site will join the IAB of the Embedded Systems I/UCRC. The entire IAB is responsible for reviewing and evaluating proposed and ongoing research projects. The IAB selects a Chairperson who runs the executive sessions of the semi-annual meetings and represents the IAB interests

[2] Dennis Gray and George Walters: Managing the Industry/University Cooperative Research Center, Battelle Press, 1998.

between the meetings. In addition to an IAB representative, an industrial partner may designate and send one or more industrial monitors to provide oversight and mentoring on specific research projects at UNT Research Site.

During the planning phase, Kavi, along with other Site Coordinators, and the director of the I/UCRC, will promote the center and recruit industrial members and seek financial support for the research projects to be undertaken by the research sites of the center. In particular, Kavi will be responsible for recruiting at least five company members to UNT Research Site. Other UNT researchers will aid him in identifying key research projects and their industrial relevance. Several large industrial concerns, who have been working with UNT (including Lockheed-Martin, Raytheon, Honeywell, Masergy, Nortel ,UGS) will help recruit other industrial concerns by providing initial contacts.

1.5. Cost and Sources of Funding.

UNT Research Site activities will be supported by industrial memberships, by NSF fund and cost-sharing funds provided by UNT (as per NSF I/UCRC requirements). In accordance with the I/UCRC eligibility requirements, UNT site will recruit industrial members and raise a minimum of $150,000 per year in memberships. If funded, each university site with 10 or more industrial members will receive $70,000 per year from NSF, and any university that brings 5 to 9 industrial partners will receive $50,000 per year from NSF. UNT will provide funds equivalent to a minimum of 25% of industrial memberships. The majority of the funds from industrial memberships and NSF will be used for research. Some of the funds (primarily those form UNT and NSF) will be used to support the marketing activities of the site. All members of the I/UCRC will have non-exclusive, royalty free access to the research supported by the center.

The UNT Research Site, along with other research sites of the I/UCRC will also seek funding to support its research from other sources, including federal, state, and local governments. We anticipate that the industry/university cooperative research will greatly enhance the capability of UNT site in obtaining funding from NSF, DoE, DoD, NASA, NIST, as well as the Texas Advanced Research Program, Texas Emerging Technology Funds, and Texas Technology Workforce Development programs.

2. Research Projects

This section presents some projects proposed by prospective industrial partners for completion within the planning period. These projects will be undertaken by UNT and other DFW are universities. Each project description is designed to satisfy a number of "Collaborative Work Package Criteria":

a) The scale and scope of the project should be appropriate for being successfully completed within one year.

b) The research should fall within the expertise of one or more center members, so that researchers in the center can effectively perform the specified tasks.

c) Results of the research should be suitable for sharing among all center members.

d) Industrial sponsors of the research must be willing to provide oversight and review of the work as it progresses.

The research projects proposed for the planning period are motivated by the core research focus of the proposed Research Sites by DFW area universities and the Embedded Systems I/UCRC center in the area of dependable net-centric systems, especially for real-time applications. There are numerous important applications of real-time network-centric systems. Boeing has developed an architecture to support integration of systems for military avionics applications [Sch01][Sch02], including coordinated control of various weapon systems [Doe99, Sha98]. Another class of applications is related to interactive multiplayer network games involving real-time inputs from multiple users [She03] and real-time tele-collaboration systems that enable participants to share a variety of multimedia documents and applications [Shi98]. Several issues related to tele-collaboration environments are explored in [Kim97][Oli03] and applications range from support for virtual labs [Hoy04] [San02] and interactive collaborative learning in distance education [Cur02] to various business applications. The Player system [Ger03] provides development support for network-centric control and coordination of a group of heterogeneous robots with applications ranging from task allocation [Ger02] to human-robot interaction [Tew03], multi-robot teams [Cha05], coordinated transportation of objects [Vau02], tracking of targets [Jun02], telepresence [Fer99][Fer01], etc.

Major embedded net-centric software and systems engineering research issues include the development of practical techniques for achieving high reliability, availability, and safety for mission-critical net-centric systems [Ajm89] [Arl90] [Ben03] [Car98] [Cha00] [Daw96] [Ead02]. The distributed nature of these systems also requires the use of techniques that can achieve high security. Assurance of high maintainability is also needed to ensure that a deployed net-centric system can be adapted as needed to cope with the rapid advances in technology as well as evolution in the requirements of these systems. Given the complexity and wide range of technologies that must be harnessed to build dependable net-centric systems, the center will also investigate techniques that enable developers to rapidly build highly dependable net-centric systems [Jon95] [Lov00] [Man92] [Mor90] [Mor94] [Mor96] [Pav01] [Roa02] [Smi90] [Tro99] [Whi01] and to deploy them in the field. These will require the development of automated tools and techniques [Crn00][Don01][Mor94] for composing systems from existing services [Bas01] [Bas02] [Bat00] [Bat04] [Che97] [Fur02] [Gam95] [Hat03] [Joh97] [Kel98] [Kim03] [Kim04] [Mor04] [Nei80] [Yen01], for dynamically monitoring the health of net-centric systems, for adapting the system in real-time at run-time, and for enabling the certification of the system to a high degree of confidence using automated validation and verification methods [Han95] [Hex03] [Hsu97] [Jen94] [Luc03].

The projects identified for the planning phase are first steps towards the research sites' vision and will help lay a solid foundation for continuing research to enable the rapid development of highly dependable net-centric systems. More importantly, these projects are of immediate interest to industry and have been formulated by prospective industrial partners of the proposed center. These include projects that will analyze the characteristics of potential net-centric applications and their communication requirements, development of network management and data compression methods for net-centric applications, software safety, and dynamic system monitoring methods using MEMS technology. Each project description lists the budget estimate as well as the investigators who will be primarily responsible for the project, though all the three universities will actively participate in all aspects of each project.

2.1. Analysis of Net-Centric Applications

Development of automated tools and techniques for achieving rapid design and synthesis of highly dependable net-centric systems requires a detailed understanding of the characteristics and attributes of potential net-centric systems, especially with respect to the demands they place on the net-centric platform. This includes issues such as assurance of network connectivity, real-time communication, security, reliability, availability, etc. In the following project, several net-centric scenarios will be analyzed to gain a thorough understanding of the requirements and challenges of emerging net-centric applications. The project will also analyze the "value-added" benefits of net-centric approaches to these application domains. The results will be shared with potential industrial members so that they can strategically plan their business plans and product directions.

2.1.1. Net-Centric Scenarios in Domain-Specific Applications

(Sponsor: M. F. Siok, Lockheed Martin Aeronautics Company)

Case for Action: The center proposes to develop solutions for net centric systems application. These potential solutions address specific apparent problems that arise in domain-specific situations that can be identified and described completely. Specific cases for potential study should be elaborated and documented in presentation format at a minimum to begin a repository of potential project application. Customers and prospective center members require concrete examples of net-centric applications in order to assess how the benefits of center membership can be leveraged within their companies. Documented cases, either completed projects or specific elaborated scenarios illustrating example application provide these concrete examples that can be leveraged by industry members of the center during key budgeting cycles.

Problem Scope: The elaborated examples of net centric application collected and presented by the net-centric center are at present few. Elaboration of the problem and potential available solution would lend credence to the center as well as provide a wealth of potential projects from which to develop, seek customers, and potentially develop. Since net-centric applications are potentially wide and varied, the center should choose a half-dozen or so domains of interest and develop one or more comprehensive scenarios that elaborate examples of problems that can be solved with net-centric solutions. Candidate domains of interest include 1) Emergency Preparedness/Emergency Response, 2) Transportation and Transshipment, and 3) Energy Distribution and/or Management.

Major Tasks: 1) Identify relevant domains of interest for the center for elaboration of scenarios. 2) For each domain of interest, identify at least 1, preferably 2 scenarios that can be elaborated to identify and describe net centric application and their utility. 3) Elaborate chosen scenarios identifying the problems and net-centric solutions/features posed. Scope potential projects to develop potential solutions identifying resource requirements including engineering, capital, schedule, tests, etc. 4) Present each scenario for review to prospective center members; package scenarios for delivery to these members for follow-up action, if any.

Leverage/Reuse/Collaboration: It is expected that this work will leverage existing standards, guidelines, and research as well as all existing literature on the subject. The purpose is to learn how to apply what is known, and scale it to new application. The expertise of center members should be exploited and collaboration with center members is expected.

Key Deliverables and Delivery Dates: 1) Identify domains of interest (1 week after go-ahead). 2) Identify scenarios to be elaborated for each domain of interest (3 weeks after go-ahead). 3) Elaborate each scenario; document with detailed presentation. (2 weeks for each scenario 3 months total). 4) Review each elaborated scenario with center members in a net-centric review fashion (2 weeks for each review/correct cycle, 2 months to review all scenario packages). 5) Package reviewed scenarios and deliver to the center for distribution and reuse as applicable (one week following completion of last scenario).

Cost and Personnel. This project will be led by I-Ling Yen of UTD with collaborations from K. Kavi and R. Brazile of UNT. Two graduate students will assist the PIs on this project. The estimated cost of this project for a 1 year period is $100,000 to support the graduate students and one month support for the PIs.

2.1.2. REP – Radiological Emergency Preparedness

(Sponsor: Scott Flowerday, FEMA in Denton)

Case for Action: The Office of Homeland Security provides for the monitoring of nuclear power generating sites in the United States. The REP group of OHS provides this monitoring. On a regular schedule, REP conducts exercises to determine if all nuclear sites are ready to properly handle an emergency at their site. All personnel at the site plus emergency personnel in the police, fire and other state agencies participate, as well as medical personnel. The exercise simulates an emergency and REP observers grade whether the participants response was acceptable. If not, a note is made and the deficiency must be fixed within a certain time. REP is divided into nine regions over the US. There is little or no overlap among the regions. There is considerable time spent on setting up the exercises, recording the results and writing a report.

Problem Scope: REP needs a system that allows for creating exercises and recording the results of those exercises. The systems should be independent for each region, yet integrate all regions for review by headquarters in Washington. There is also a need to integrate the emergency response plans of REP with the plans for fire police and medical emergencies. By integrating these plans, redundant and inefficient use of emergency resources may be avoided.

Major Tasks: The first phase, that is a system for one region, has been done. Techniques for replicating the basic system and integrating the data and services of that system to all regions must be developed. A format and procedure for describing and integrating the various emergency response plans must be designed and implemented.

Key Deliverables and Delivery Dates: Replicating and integrating the data and services for a second region can be finished in six months. A plan for completing the other regions could be developed three months after that. Integration of emergency response plans can be planned in six months. The implementation of the integration of two plans would take another three months.

Cost and Personnel: Robert Brazile and Kathleen Swigger will lead the project in collaboration with Gopal Gupta of UTD. One graduate student for 1.5 years will be needed. Total cost $50,000.

2.1.3. Service-Oriented Computing

(Sponsor: Tom Hite, Metallect)

Case for Action: Software designers are moving towards a *service-oriented software architecture* [paul04a, tsai04a], while software is increasingly being made available as a service accessible over the Web (*Web service*). One can envision a world where such *Web services* will be as ubiquitously available and as easily searchable as documents are today on the World Wide Web. For Web services to become practical, an infrastructure needs to be developed that allows users and applications to deploy, discover, select [Gao05a] [Gao05b] [Lia04] [Ma03] [Ma05] [Ma06] [Rao03] [Ser05] [Zha06], compose, and synthesize services automatically [Bas01] [Gao06] [Liu03] [Liu05] [Wan04a] [Wan04b] [Wan05a] [Wan05b] [Yen02]. For this automation to be effective, formal semantic descriptions of Web services should be available [W3C04]. In this project, we aim to conduct research to develop an infrastructure that will permit this automation.

Problem Scope: We start by first defining a mark-up language for describing the semantics of Web services. This language, called Universal Service-semantics Description Languages (USDL), can be thought of as the semantic counter-part of the syntactic Web Services Description Language (WSDL). Next, we have to formally define the discovery and composition problems for Web services and develop solutions for these problems. These solutions are based on semantic specifications of Web services. For our project, the semantic description is, of course, assumed to be specified in USDL (Universal Service-Semantics Description Language). Finally, the algorithms that realize the solutions to the service discovery and composition problems need to be efficiently implemented. An additional critical issue that arises is developing tool support for generating USDL annotations for a Web service. In this project, our goal is conduct research in all aspects related to automated discovery and composition of Web services.

Major Tasks: 1) Refine the design of Universal Service-Semantics Description Language (USDL) to capture the various types of actions that a service can accomplish. Right now the set of actions is very limited (*create, find, delete, update*). Actions that cannot be captured have to be approximated by the reserve word *affects*. Our plan will be to allow any verb to describe the action achieved by a service. 2) Develop the complete Discovery and Composition Engines for USDL: Our preliminary implementations reported earlier restricted themselves to service repositories that had syntactic descriptions of services via WSDL. We would generalize our algorithms to take full USDL into account. The development of efficient discovery and composition engines will constitute the major part of the project. 3) Develop tools for semi-automatically generating USDL annotations: We will develop preliminary tools for semi-automatically generating USDL annotation for atomic and legacy services from its code and English documentation. We will also develop the necessary theory and implementation of a tool that will synthesize the USDL annotation of a composite service from the USDL description of its component services and the Composition query.

Leverage/Reuse/Collaboration: This research will have an impact on industry and society, as it will lead to new standards being developed. Additionally, it will ultimately result in development of search engines for Web services which can have a significant impact on industry/society similar to the impact that search engines for documents have had in our industry/society. Our research in discovery and composition can potentially result in productivity gains as in many instances new applications need not be developed from scratch, they can be realized by discovering and composing existing components.

Key Deliverables and Delivery Dates: 1) Refine the design of Universal Service-Semantics Description Language (USDL) (2 months after go-ahead). 2) Develop the complete Discovery and Composition Engines for USDL (6 months after go-ahead). 3) Develop preliminary tools for semi-automatically generating USDL annotations (12 months after go-ahead).

Cost and Personnel. Gopal Gupta of UTD will lead this project with collaboration from Phil Sweany of UNT and 1 graduate student for an estimated cost of $50,000 (1 year support for the graduate student and 1 month support for the PI).

2.2. High-Assurance Enabling Platform for Net-Centric Systems

The objective here is to investigate techniques for meeting the demands of mission-critical net-centric applications, including real-time factors [Far00] [Fra95] [Leu00] [Mar95] [Swe94], hardware/software co-design [Ziv96], minimization of bandwidth required by the applications [And92] [Che05] [Xia05] [Yen07], energy requirements [Ye00], handling the loss of some connections [Tim94] [Tsa96], etc.

2.2.1. Network-Management

(Sponsor: Geoff Bourne, Tektronix Texas)

Case for Action: Net-centric applications are likely to be distributed in various arrangements and communicate large volumes of control and data plane messaging [Pet03]. Various messaging frameworks are presently in use, which include TCP and a proprietary protocol, both of which are used in a point-to-point manner. The diversity of these frameworks makes current software development/testing activities inefficient. The point-to-point manner of communication is also prone to single-point communication failures. Therefore, a new, generic, robust communication framework is needed

Problem Scope: A distributed generic messaging framework, COMS, based heavily on hierarchical paradigms is presently being developed at Tektronix and presents many opportunities to optimize communication paths. Specifically, high volume messaging demands utmost consideration of network resources and may require bypassing the normal COMS communication paths in a deterministic manner. The scope of this project is to design and implement a path discovery algorithm and integrate the solution into existing COMS prototype framework.

Major Tasks: 1) Acclimate to existing COMS implementation. 2) Identify applicable algorithms to solve the stated problems. 3) Develop a report that describes the work including at least integration strategy and a survey of algorithms to be developed. 4) Develop and integrate the solution into the existing framework. 5) Provide demonstration.

Key Deliverables and Delivery Dates: 1) Statement of work including an integration strategy and a survey of related algorithms (two months from start). 2) Monthly project reviews (once a month). 3) Final project review (1 year after project starts). Constraint: It is highly desirable to use an Agile (or iterative) approach with cycles of demonstrable software prototypes every 2 weeks.

Cost and Personnel. Neeraj Mittal of UTD will lead this project with assistance from Suku Nair of SMU. The estimated cost, including support for 1 graduate student for one year and 1 month support for the PI, is $50,000.

2.2.2. Data Compression in Sensor Systems

(Sponsor: D. Struble, Raytheon)

Case for Action: The volume of data and the number of data sources to be negotiated on the battlefield is significant. High resolution video and audio images from multiple sensing elements requires high order compression to achieve meaningful levels of situational understanding within tactical mission time frames. Research areas include: (1) Bandwidth compression tools, including CODECs, frame rate reduction, and windowing; (2) scenario dependent control of the application of BW compression tools; (3) video transport including tactical RF transmission medium characteristics; and (4) decompression tools, all of which must be compatible.

Problem Scope: There is interest in determining the availability of software components and/or algorithms that can be used across multiple sensor systems. Common image compression and decompression schema would provide significant economies of scale when deploying multiple sensor platforms to achieve higher levels of sensor fusion using combinations of video and audio data. Contemporary methods offer at best 10:1 compression ratios for specific types of data, while the problem spaces for sensor fusion and cognitive workload reduction requires ratios better than 100:1, an order of magnitude improvement.

Major Tasks: 1) Investigate the state-of-the-art in compression algorithms and software components and develop a tutorial explaining each in quantitative terms, i.e., method (null suppression, pattern substitution, etc.), ratios by data type, etc. 2) Evaluate applicability of compression algorithms or tools to the problem space of sensor fusion and improvement of situational awareness. 3) Document proposed trade studies or prototyping activities that could be used to determine applicability of specific methods.

Key Deliverables and Delivery Dates: 1) Results of literature search and documentation of available compression and decompression methods with analysis by specific data types (DUE: 90DA CA). 2) Empirical, quantitative data showing results of specific algorithms or tools as applied to identified data types of interest (DUE: 180DA CA). 3) White paper documenting proposed trade studies, prototyping activities, and simulations to be used to improve compression/decompression speeds and ratios based on specific proposed improvements (DUE: 1 year after CA).

Cost and Personnel. Steve Tate of UNT will lead this project with collaborations from Hesham El Rewini of SMU and Neeraj Mittal of UTD. The cost to support one graduate student for one year and 0.5 month for the PI is estimated at $40,000.

2.3. High-Assurance Systems Engineering for Net-Centric Systems

The above projects are largely focused on the network and communications aspects of net-centric systems. The following projects focus on the software and systems issues involved in the development of highly dependable net-centric systems [Bin03] [Bro00].

2.3.1. Software Safety Tutorial

(Sponsor: Dennis J. Frailey, Raytheon)

Case for Action: An increasing number of RFPs, especially from the military services, call for software safety analysis and for addressing software hazards within the requirements, design, development and testing of software. In many cases, appropriate standards, guidelines or conventions are called out or assumed. Yet few software developers are familiar with software safety standards or procedures. Software safety can add significant cost to a software development program, thus it is essential that one apply a well-planned set of procedures that are appropriately tailored to fit the specifics of a given application without adding undue cost.

Problem Scope: There are several applicable standards and guidelines for software safety [Chi92] [Ham77] [How82] [Voa95] [Voa98], but no universally accepted ones. Contractors need help in determining which standards and guidelines to apply in any given situation. In many cases they also need basic training on software safety. We require a basic tutorial document that covers the fundamentals of SW safety, what the various standards are, how one goes about applying them on real software projects, how they compare (so one can select among them), etc. Some literature exists in this field, but its value and relevance are typically unknown to software practitioners. We need an annotated bibliography of this literature to serve as a guide to publications that are of most value and relevance.

Major Tasks: 1) Investigate the state-of-the-art in software safety including what standards exist, how they compare, who applies them under what conditions, etc. 2) Determine what research or literature exists that pertains to this subject. 3) Prepare a tutorial report on software safety aimed at experienced software developers who have not previously had to address software safety issues. 4) As an adjunct to the tutorial, prepare an annotated bibliography of major publications on software safety.

Leverage/Reuse/Collaboration: It is expected that this work will leverage existing standards and guidelines as well as existing literature on the subject [Chi89] [Cia89] [Gei90] [Hil96] [Joh88]. The purpose is to learn how to apply what is known, not to expand into new techniques. The expertise of consortium members should be exploited and collaboration with consortium members is expected [Che91] [Des89] [Kav87] [Kav91] [Kav94] [Kav96] [Tia00] [Tia05].

Key Deliverables and Delivery Dates: 1) Interim review of literature search and assessment activities. This will be a presentation and discussion with software development organizations within the consortium that need the information (DUE: 90DA CA). 2) First Draft Tutorial Report (DUE: 180DA CA). 3)

Semifinal Draft Tutorial Report, including annotated bibliography (DUE: 270DA CA). 4) Final Tutorial Report, including annotated bibliography (DUE: 1 year after CA).

Cost and Personnel. Jeff Tian of SMU along with D.T. Huynh and Eric Wong of UTD will lead the project. Estimated cost to include 2 graduate students for 6 months and ½ month for the PI is $40,000.

2.3.2. Self-Configuring Wireless MEMS Network

(Sponsor: M. F. Siok, Lockheed Martin Aeronautics Company)

Case for Action: There is considerable interest in miniature, lightweight, self-powered, wireless sensors which can be installed in practically any location in the aircraft and left to function without requiring access panels for maintenance. Envisioned use includes condition-based maintenance and flight-test and customer instrumentation applications. Given that small enough sensor packaging can be achieved, the ability of the sensors to configure themselves into a secure, viable, working network, and reconfigure the network at any time without loss of information (not necessarily data) is desired. Large cost savings can be achieved in predictive maintenance rather than timed or reactive maintenance and also in configuring aircraft for instrumented tests.

Problem Scope: There is currently on-going research into MEMS device scale sensor packages with self-contained power generation, e.g., energy harvesting and storage (TCU, UC Berkeley). Given that device volumetric packaging is targeted for less than 5 cubic mm with operations in hostile EMI and working environments, software needs to be small and efficient. Networking software should sense other compatible sensors in close proximity and establish a secure wireless communications environment. Network should be adaptable to real-time changes in number of networked devices, changes in operational environment, etc., without loss of service

Major Tasks: 1) Given that the MEMS mechanical and physical constraint problem can be solved separately, investigate, prototype, and test small networking software package that meets requirements of secure, viable, and adaptable communication for MEMS scale sensor devices delivering information packages to requesting systems also in secure, viable manner assuming all mobile devices. 2) Develop report that identifies and describes work, including requirements, designs, tests, and source. Deliver all items of software development lifecycle in electronic form only. 3) Provide software demonstration.

Leverage/Reuse/Collaboration: This work will leverage existing standards, guidelines, and existing literature on the subject. The purpose is to learn how to apply what is known and scale it to new applications. Use model-based engineering techniques for requirements and designs; use IDEs, automated test and test generation tools. The expertise of consortium members should be exploited and collaboration with consortium members is expected.

Key Deliverables and Delivery Dates: 1) Interim review of literature search and assessment activities. Deliver presentation and discussion with software development organizations within the consortium (3 months after go-ahead). 2) Perform incremental review of designs, tests, and prototype code using IDEs. Deliver final review in presentation form to consortium in net-centric manner. (6 months after go-ahead, incremental reviews to be coordinated separately by development team prior to presentation delivery). 3) Provide demonstration of network capability to deliver information to requesting system illustrating network reconfigurability with nodes dropping out and adding in, with loss of data due to ambient environmental noise, and independent mobility of all network nodes in same. Demonstration may involve actual equipment, prototypes, and/or simulations (9 months after go-ahead).

Cost and Personnel. Robert Akl and Krishna Kavi of UNT will lead this project with collaborations from Farokh Bastani of UTD and Hesham El Rewini of SMU. Support for 2 graduate students for one year and 2 months for PIs will be needed, for an estimated cost of $100,000.

2.4. Project Summary and Schedule

The projects proposed by prospective industrial partners will be completed within the planning period. The results from these projects will not only be shared among all center members, but will also be used to recruit other industrial members. Table 1 summarizes these projects and their tentative schedule, together with the academic member institutions that will conduct the projects.

Table 1: Example Project

Project Name	Sponsor	Duration	Milestones	Staff	Estimated Cost
Net-Centric Scenarios in Domain Specific Applications (Section 2.1.1)	Mike Siok, Lockheed-Martin	7-12 months	1 week: Domain identification; 3 weeks: Scenario identification; 4: Scenario elaboration 6: Scenario review 7: Packaging	I.-L. Yen (UTD), K. Kavi and R. Brazile (UNT)	2 Students and 2 month for PIs, $100,000
REP – Radiological Emergency Preparedness (Section 2.1.2)	Scott Flowerday, FEMA in Denton	18 months	6 months: Extension to a second region 9 months: Plan for other regions 15 months: Plan for integration with emergency response systems 18 months: Prototype implementation	R. Brazille, K. Swigger (UNT), G. Gupta (UTD)	1 student for 1.5 years, $50K
Service-Oriented Computing (Section 2.1.3)	Tom Hite, Metallect	12 months	2 months: Refine/design USD 4 months: Develop engines for USDL 1 year: Tools for semi-automatic generation of USDL annotations	G. Gupta (UTD), P. Sweany (UNT)	1 Student and 1 month support for PI, $50,000
Network Management (Section 2.2.1)	Geoff Bourne, Tektronix	12 months	2 weeks: Statement of work; Monthly reviews thereafter until completion.	N. Mittal (UTD), S. Nair (SMU)	1 Student and 1 month for PI, $50,000
Data Compression in Sensor Systems (Section 2.2.2)	Dave Struble, Raytheon	12 months	3 monhts: Survey; 6 months: Empirical, quantitative data 1 year: Completion	S. Tate (UNT), H.E. Rewini (SMU), N. Mittal (UTD)	1 Student and ½ month for PI, $40,000
Software Safety Tutorial (Section 2.3.1)	Dennis Frailey, Raytheon	6 months	3 months: Interim review 6 months: Draft report	J. Tian (SMU), D.T. Huynh and E. Wong (UTD)	2 Student and ½ month for PI $40,000
Self-Configuring Wireless MEMS Network (Section 2.3.2)	Mike Siok, Lockheed-Martin	9 months	3: Interim review 6: Final review 9: Demonstration	R. Akl , K. Kavi(UNT) , F.B. Bastani (UTD), H.E. Rewini (SMU)	2 Grad Students and 2 months for PIs, $100,000

3. UNT Research Site Participants

The University of North Texas is a comprehensive, Doctoral/Research Extensive public institution, located in the Dallas-Ft. Worth Metropolitan area. With over 32,000 students, UNT is the fourth largest University in Texas. The Computer Science and Engineering Department offers BS, MS, and PhD degrees in Computer Science and BS and MS degrees in Computer Engineering. In addition to the faculty listed below, other faculty from the College of Engineering will participate in the UNT Research Site activities based on their expertise and project needs. The UNT Site will draw upon a large number of faculty with expertise and ongoing projects that benefit the Center's research focus and mission.

a. Krishna Kavi, chair of UNT's Computer Science and Engineering Department, conducts research primarily in computer systems architecture, agent-oriented system and formal methods of software engineering. Kavi will be the Site Coordinator of the UNT Research Site

b. Philip Sweany, associate professor in UNT's Computer Science and Engineering Department, maintains a research focus in both compiler optimization for architectures exhibiting fine-grained parallelism and application of compiler optimization algorithms to automated synthesis of net-centric systems.

c. Robert Brazile, associate professor in UNT's Computer Science and Engineering Department, conducts research in Information Integration and Computer Supported Collaborative Work. His current interest is in shared access to both distributed data and distributed applications using XML and Web Services.

d. Robert Akl, assistant professor in UNT's Computer Science and Engineering Department, conducts research on wireless sensors and WIFI network design and optimization. His current work has been on energy-aware routing and synchronization in sensor networks.

e. Steve Tate, associate professor in UNT's Computer Science and Engineering Department, conducts research on Computer Security, cryptography, algorithms and data compression.

4. Results from Prior Support

Projects from the University of North Texas (UNT). The purpose of the NSF award NSF-0222628, "Computational Science and Engineering – Information acquisition and management of Infrastructure (PI: Kavi) was to acquire computational and storage resources to support several scientific and engineering projects. Among the acquired resources are i) a 4-Node SMP, ii) a tera-byte RAID storage server, iii) a 8-node Beowulf cluster, and iv) 4 workstations, including a HP/DEC Alpha for running ATOM tools. The resources were used to perform extensive instruction set level simulations of both conventional superscalar architecture and innovative decoupled multithreaded architecture. In addition, we also performed experimental evaluation of innovative cache memory organizations, decoupling memory management from main processing elements and dynamically suppressing redundant function calls.

The following publications of PI Kavi resulted from this award: [Li05][Li06][Li07][Lin07][Naz06][Rez06][She04][Son03].

Supplementary Documents

The supplementary documents include letters from potential industrial members and a tentative agenda for the meeting with potential industrial partners in the Dallas/Fort Worth (DFW) Metropolitan area.