Special Notice 11-SN-0004

Special Program Announcement for 2011 Office of Naval Research Applied Research

"Data to Decision"

I. INTRODUCTION

This announcement describes an applied research thrust, entitled "Data to Decisions," to be launched on behalf of the Office of the Secretary of Defense (OSD), Office of the Director of Defense Research and Engineering (DDR&E) under the Office of Naval Research (ONR) Broad Agency Announcement (BAA), ONRBAA11-001, Long Range Broad Agency Announcement for Navy and Marine Corps Science and Technology which can be found at http://www.onr.navy.mil/Contracts-Grants/Funding-Opportunities/Broad-Agency-Announcements.aspx. The research opportunity described in this announcement specifically falls under numbered paragraph 1 of the "Command, Control, Communications, Computers, Intelligence, Surveillance and Reconnaissance (Code 31)" sub-section. The submission of full proposals, their evaluation and the placement of funding vehicles will be carried out as described in that Broad Agency Announcement.

The purpose of this announcement is to focus attention of the scientific community on (1) the area to be studied, (2) a workshop for dialogue amongst those interested in this arena, and (3) the planned timetable for the submission of full proposals.

II. TOPIC DESCRIPTION: Data to Decision

Background

Mathematical decision science is an enduring research field pervasive throughout many aspects of the US economy, and critical to National Security. In National Security applications, rapidly changing and adaptive threats have driven recent advances in sensors leading to petabyte and larger datasets rich in spectral, spatial and temporal features. With this abundance in data, the need to discover and identify threat signatures in complex, incomplete, imprecise and potentially contradictory large data sets has become a critical issue in military decision-making.

Systems that support decision-making are an integral and critical piece of nearly all defense and intelligence related missions. These systems span a broad spectrum of defense needs ranging from tactical operations to strategic intelligence ones. For tactical operations, the decision support system is generally fully automatic and externally cued. These systems are typically closed loop control systems and rely on narrow field-of-view sensors to support low-latency real-time operations. These systems include end-game seekers, counter air and weapons like AIM-9X, JDAM and cruise missiles. Operational Intelligence missions are usually associated with locating threat systems like Surface-to-Air missiles and submarines, but can also include missions like Space Situational Awareness. These missions are characterized by medium temporal latency, have larger spatial fields of regard and are primarily analytic-intensive with simple "bell-ringer" user interfaces. Strategic Intelligence missions tend to dominate much of asymmetric warfare and are characterized by the requirement to operate in an open and partially unknown environment. This mission includes data mining and interpretation, multiple hypothesis generation and adaptable interfaces to support discovery, integration and collaboration.

Although each of these missions is different in its application, there are many similarities throughout; for instance, the common goal in all these systems is to provide recommendations that lead to high value actions consistent with mission goals. These recommendations occur through a user interaction layer that allows the user to task the system and evaluate its recommendations. When tasked, the system first postulates options, or hypotheses, which are used to drive a model. Data is collected and compared to each model in the analytic layer, and decisions are made as part of that layer on whether to collect more data, generate new options, or provide a recommendation to the decision maker. Data collection and management occur through a data management layer.

With decision support systems being so critical to defense needs, it is not uncommon for these systems to go through the conventional 10-year development cycle of most Major Defense Acquisition Programs. Not only does this 10-year cycle not match the continuous and compelling need for new decision support systems, it also results in systems that are hard-wired and not adaptable to changing countermeasures. Additionally, it becomes difficult to incorporate scientific and technical innovations developed in academia, small businesses, government laboratories, and large industries; thus, these systems generally do not integrate new advances across the defense enterprise. This approach is reminiscent of the early computer industry in late 1950s when DEC built the PDP-1, which had a hard-wired instruction set and took many years of development. In the 1960s, IBM recognized a better way to build computers was to modularize their components and the first successful attempt using this approach was the IBM System 360. System 360 was built using modular components under closed standards and resulted in a system much simpler to build, maintain and upgrade. One could easily argue that the introduction of Service Oriented Architectures has allowed Decision Support Systems to reach this same level of maturation, since many of these new systems can be composed and built using modules cast as services.

But the modern computer industry is much further along in its maturation now. New Information Technology (IT) systems are rapidly built through the horizontal integration of open standard modular components that are easily reconfigured across a wide array of products. For instance, a

hard drive manufacturer can now supply the same drive for a laptop, desktop, MP3 player and even a smart mobile device. By having matured the "lego" approach to building new systems, the IT industry has enabled true rapid development; but just as important, they've provided a way to capture innovations across small innovative companies that specialize in critical components.

Objective

The primary purpose of the Data-to-Decision (D2D) Program is to develop a capability to rapidly develop, evaluate, and field prototype Decision Support Systems. To accomplish this goal, the D2D program has identified the creation of a library of flexible modules that can be easily repurposed, or modified, to meet the needs of many National Security applications as a critical need. This library will form an open architecture basis for the rapid development of new prototype decision support systems. Additionally, the architecture should allow these modules to be rapidly reconfigured, should be sufficiently flexible to enable easy upgrades of existing modules and should allow insertion of new or alternative modules created by innovative organizations.

To help focus on the development of these broadly applicable modules, the program will be conducted using a series of cross-service challenge problems centered on specific sensing modalities. Data will be provided at the unclassified level, to ensure all innovative ideas can be captured, and at the operational level to make sure the solutions work with real data development that will generally occur through open consortiums delivering specific capabilities in contrast to a program in which the offeror has unique measures of success and operates independently.

The system will be built on an open-architecture testbed running a Service Oriented Architecture with distributed nodes. The testbed will be built, owned and operated by a government team and will be used to perform research and module development in advanced user interfaces, analytics and data management. The testbed itself will consist of 100+ high performance nodes on a 10 GB Ethernet backbone. Between 0.5-1.0 petabyte of storage will be provided, and there will be a mix of different memory types to support research in storage architectures and efficient data access and retrieval. The system will be fully partitionable and reconfigurable to support research in grid vs. cloud architectures, algorithm scalability, data indexing, common data representation and other modules as identified or needed. Details regarding the Application Program Interfaces (API's) and hardware will be supplied to offerors that have been selected for an award.

The D2D Program will develop a library of modules through a spiral development process in support of each of these three layers. This process is modeled after other successful programs and will consist of both government and contractor teams working together to develop, identify and deliver modules to support future rapid development capabilities. This model is outlined in Figure 1 and consists of four functions. The CHALLENGE team chooses the set of cross-Service challenge problems associated with a particular data source and is responsible for collecting data to support that mission. It is also responsible for making connections with

operational partners to insure proposed solutions are consistent with needs. Both unclassified and operational data sets are collected, and this team also develops Concept of Operations

(CONOPs). Data is then delivered to the DEVELOP team whose responsibility is to develop solutions against the challenge problem. The ANALYZE team serves as an honest broker and tests each solution to determine how well it works and under what conditions. Each solution is tested against several sensors and target types to determine its extensibility. The BUILD team then integrates the modules that perform well into the library.

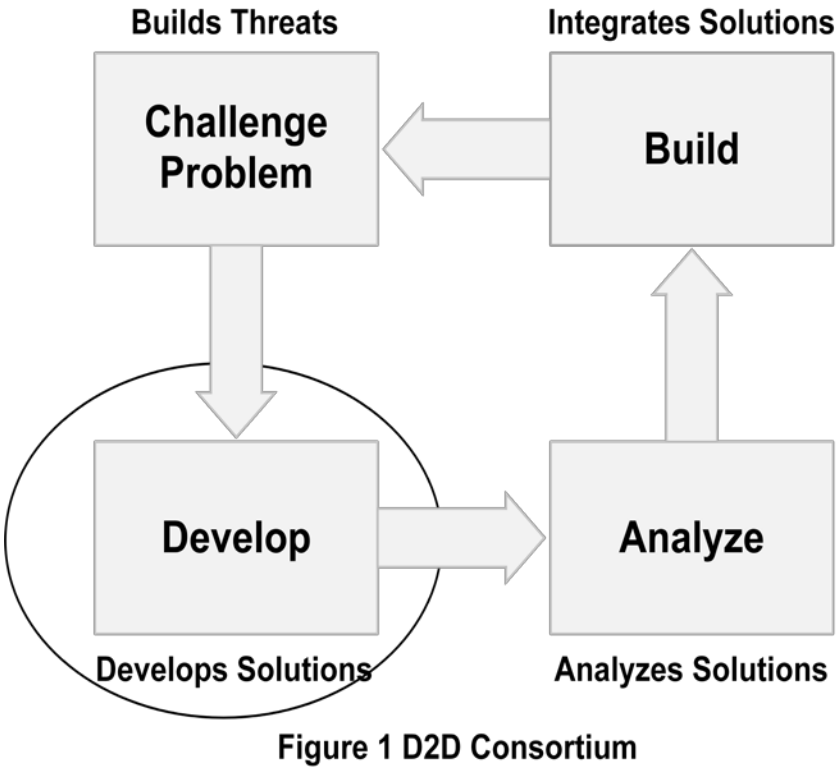

Figure 1: D2D Development Consortium and Government Team.

Analytics will be a significant focus throughout this program and will be a primary means for converting the varied and many unstructured data sources to structured databases. Unstructured data sources to be considered include Wide Area Motion Imagery, Ground Moving Target Indication (GMTI), text and imagery, but others sources will also be considered as needed. A common theme throughout this program will be to develop analytics that can be used broadly across many sensing modalities and target types. Module extensibility might be achieved through decomposing analytics into granular elements that span the set of data sets and targets. Development of simple algorithms that attain performance through fusion of uncorrelated or

contextual data, or through statistics, or model-based approaches might also be viable candidates. Offerors are free to propose any algorithm or method that supports the conversion on unstructured data to structured data in support of the Challenge Problem. The government does not have a preferred approach or paradigm for Analytics and wishes to collect, develop, and evaluate the performance of algorithms and approaches in the context of the Challenge Problem. Algorithms, approaches, and components may have been developed in the context of another challenge problem, and offerors are free to submit these algorithms, approaches, and components for evaluation if they believe that the performance of these algorithms, approaches, and components is warranted. The government may also explore the performance of the proposed algorithms, components, and approaches in the context of other cross-service mission areas beyond the challenge problem described below.

Structured databases will be mined and aggregated in the user interface layer, which will develop more than just traditional user interfaces. Data discovery, mining, integration, and collaboration will all be important components of this aggregator function in which individual data sources are fused. Research in this area will focus on broad issues like search and mining of tagged structured data, provenance management, metadata structures and collaboration toolsets. Advanced work in end-user programming methods, workflow automation, templates and ontology-generic interfaces will also be conducted.

The D2D program will be primarily conducted through awards to a consortium consisting of industry, academia, and government laboratories. Teaming arrangements are welcome and can be formed as part of the proposal process, but each team member will be separately funded. To ensure all innovative ideas are captured, teaming arrangements are not a requirement for participation in this award. Consortium members will be incentivized to work together as a team and assist each other as necessary. Non performers will have a diminished role or will be replaced. The business set up will be sufficiently robust such that success, or lack thereof, will not be based on the performance of one organization. The consortium either wins or fails as a team.

The awardees will jointly address the challenge problems using open source development, and openly share and provide documentation of all modules among team members. The aim will be to develop and deliver “best of breed” capabilities in a short period of time. To offset biases inherent to all organizations, "peer review” groups will be established. These groups will be composed of government, industry and academia and will oversee a build-test-build process using “real world” data sets to evaluate new capabilities. These evaluations will be done in a transparent fashion in which the offerors, the government, and appropriately selected government support contractors and FFRDC's will be invited to participate.

Even though open source development is the goal of this program, DDR&E and ONR recognize some Offerors are not comfortable with open source development. Some Offerors may have intellectual property, specific methods for graph construction, or analytic modules that help solve the challenge problems, but may wish to keep the software proprietary. These capabilities can be tested and evaluated by the government team. Proprietary products will be treated as a “black

box" during the peer review group evaluation process. Funding may be setaside for these Offerors to demonstrate the performance of their algorithms on operational data. In the event these capabilities exceed the performance of the consortium's capabilities with similar functionality, the government will consider licensing the intellectual property for inclusion in the D2D architecture while preserving the intellectual property of the offeror.

Goal

Because of the importance of the Counter Insurgency (COIN) mission and the focused concentration it has received over the past several years, the program will begin with an emphasis on the application of MOVINT data to COIN. Moving Intelligence (MOVINT) data for this mission primarily consists of Wide Area Motion Imagery (WAMI) persistent sensing, but can also include GMTI derived from radars and other sources. The MOVINT data is typically used to connect events, or places, with other places in a graph theoretic framework in support of building social networks. There are many potential approaches to building this network, and the offerors should not feel constrained to any specific solution. The overall goal of this program phase is to develop, evaluate, and provide early prototypes of capabilities that ingest MOVINT data and build a social relationship graph consistent with supplied data. Offerors are free to propose methods for the automatic construction of graphs, detection schemes, and anomaly detection techniques that yield high probability of detection with low false alarm rates. Additional areas of interest to the program are methods and supporting technologies that abet the automatic generation of graphs.

III. WORKSHOPS

The Office of Naval Research intends to host two (2) workshops. Details of the workshops including dates, locations, and registration information is available at http://data-to-decisions.ida.org.

IV. FULL PROPOSAL SUBMISSION AND AWARD INFORMATION

Full proposals (including one technical volume and one cost volume) should be submitted under ONRBAA11-001 by **29 March 2011**. Full Proposals received after that date will be considered as time and availability of funding permit.

ONR anticipates that **grants and contracts** will be issued for this effort. Full proposals for grants must be submitted through http://www.grants.gov. The following information must be completed as follows in the SF 424 to ensure that the application is directed to the correct individual for review: Block 4a, Federal Identifier: enter N00014; Block 4b, Agency Routing Number: Enter the Program Office Code (i.e., 312) and the Program Officer's name, last name first, in brackets (i.e., [Toth, Gary]). Applicants who fail to provide a Department code identifier may receive a notice that their proposal will be rejected.

ONR, on behalf of DDR&E, plans to fund eight (8) to ten (10) individual awards with a value of $300,000 to $750,000 per year. However, lower and higher cost proposals will be considered.

The Research and Development (R&D) efforts to be funded consist of Applied Research. The funds available to support awards are budget Activity 6.2. The period of performance for projects may be from one (1) to four (4) years.

Although ONR expects the above described program plan to be executed, ONR reserve the right to make changes.

Funding decisions should be made on or about **29 April 2011**. Projects will have an estimated contract award or grant date of **01 August 2011**.

VI. POINTS OF CONTACT

In addition to the points of contact listed in ONRBAA11-001, the specific points of contact for this announcement are listed below:

Technical Point of Contact:

Dr. Gary Toth, Program Officer, Code 312, Email Address: gary.toth@navy.mil

Business Point of Contact:

Ms. Kenesha Y. Hargrave, Contract Specialist BD 0251, Email Address: kenesha.y.hargrave@navy.mil