ADMINISTRATIVE NOTE:
NEW REQUIREMENTS/PROCEDURES

==

BAA 04-29 PROPOSER INFORMATION PAMPHLET

==

Deleted: 03-xx

Deleted: 3

Deleted: 44

The Defense Advanced Research Projects Agency (DARPA) often selects its research efforts through the Broad Agency Announcement (BAA) process. The BAA will be posted directly to FedBizOpps.gov, the single government point-of-entry (GPE) for Federal government procurement opportunities over $25,000. The following information is for those wishing to respond to the Broad Agency Announcement.

Coordination Decision Support Assistants (COORDINATORs), SOL BAA 04-29, Proposals Due: Initial Closing: August 23, 2004, Final Closing: June 1, 2005, POC: Dr. Tom Wagner, DARPA/IPTO; FAX: (703) 741-7804

Deleted: Architectures for Cognitive Information Processing (ACIP), SOL BAA 04-XX, Proposals Due: Initial Closing: February 27, 2004, Final Closing: January 7, 2005, POC: Mr. Robert Graybill, DARPA/IPTO; FAX: (703) 807-1720

Deleted: Self-Regenerative Systems (SRS), SOL BAA

Deleted: 03-xx

Deleted: 03-44, Proposals Due: Initial Closing: November 26, 2003, Final Closing: September 24, 2004, POC: Mr. Lee Badger, DARPA/IPTO; FAX: (703) 741-7804

Deleted: ¶

PROGRAM OBJECTIVES AND DESCRIPTION.

Introduction

The Defense Advanced Research Projects Agency (DARPA) Information Processing Technology Office (IPTO) is soliciting proposals for Coordination Decision Support Assistants (COORDINATORs), a new program to develop software coordination managers that provide coordination support to humans. The goal of COORDINATORs is to help fielded human units (e.g., soldiers, tactical teams) adapt their mission plans online in response to change. Coordination support will enable fielded units to respond more rapidly and more accurately to the dynamics of the situation while incurring less cognitive load and performing with a greater degree of coordinated action.

There are five primary technical areas in this program: distributed activity coordination, context-dependent coordination autonomy, machine learning, organizational reasoning, and meta-cognition. Hard research problems include distributed coordination over large interconnected mission structures that change dynamically, supporting coordination of large-scale operations where units may have roles in multiple missions, learning to support the units better by automating decision making when data is potentially sparse, responding in (fast enough) "real time" to change, and reasoning about military decision-making policies and procedures during coordination.

The program is expected to have four 12-month phases. Only Phase I will be funded initially. DARPA will host an Industry Day for the COORDINATORs program on July 14, 2004. For more details and registration information please go to http://www.darpa.mil/ipto/Solicitations/solicitations.htm. Additional BAA details follow.

Placement and Motivation

The focus of the program is to create distributed intelligent computational systems that adapt existing mission plans online, in real time, by making changes to task timings and allocations and by selecting from pre-planned contingencies. Replanning from first principles is not part of this program. Plans are formed off-line, *a priori*, by human planners using existing military planning techniques. COORDINATORs will solve the online adaptation problem using the aforementioned techniques (adjusting task timings, changing task allocations or assignments, and selecting from pre-planned contingencies). For solutions that fall outside of this space, human input will be required. This will enable COORDINATORs to provide effective support without large amounts of domain knowledge. This will also enable COORDINATORs to operate within the existing military structure by supporting existing processes/procedures rather than replacing them.

Figure 1 illustrates the concept. Currently the military has effective human processes for mission planning that incorporate a wide range of factors from target selection to support logistics. Where coordination technology can pay the highest dividends is in what happens to those static (often paper) mission plans and pre-planned contingencies when the units deploy. Once deployed, the game changes. The units are physically distributed, authority is distributed, and information is distributed. When change occurs the units must gather and exchange the change information, evaluate the implications of the change, generate candidate response options by considering their initial plans / contingencies, evaluate downstream implications of each option, and evaluate the trade-offs of the different options. They must do this in a distributed setting and generally do this using radios to communicate with each other and up/down the chain of command. When humans perform this coordination, the results are error-prone, suboptimal, time consuming to produce, and the process of coordination itself incurs great cognitive load. The key issue is that it distracts the human units from focusing on the big picture and from focusing on the enemy, and divides their attention between the high-level tasks (at which humans are proficient) and the low-level information exchange and analysis.

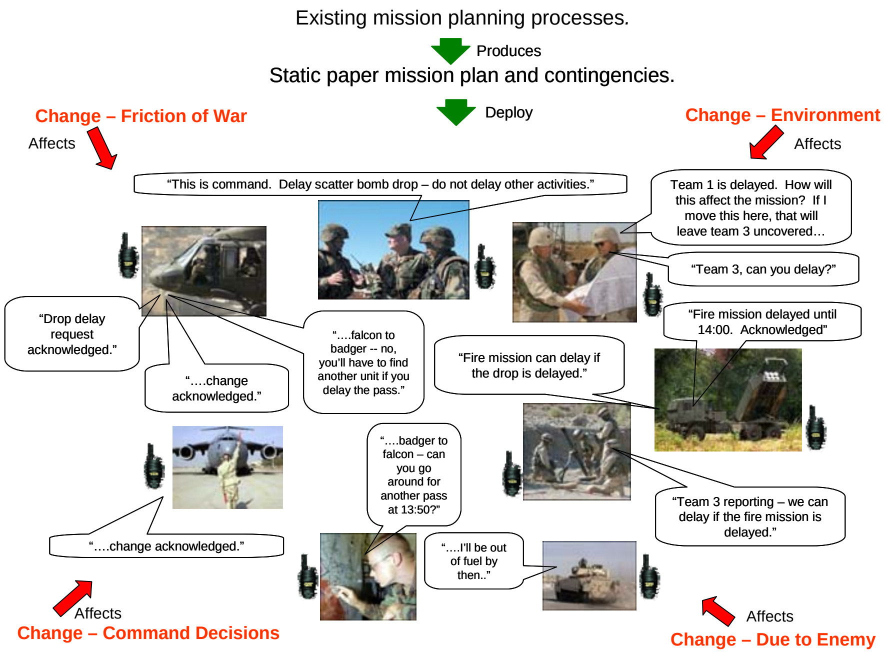

Figure 1 - Today Coordination Is Manual And Distracts Human Units

Our goal is to replace human labor with COORDINATORs / cognitive coordination managers, as shown in Figure 2. With intelligent support, responding to change will be fast, precise, and not labor intensive. The humans can focus on managing the high-level picture while the COORDINATORs handle information exchange, reasoning about the implications of change, option generation, option evaluation, and over time, even learning to make decisions for the human user when he/she is occupied with other tasks.

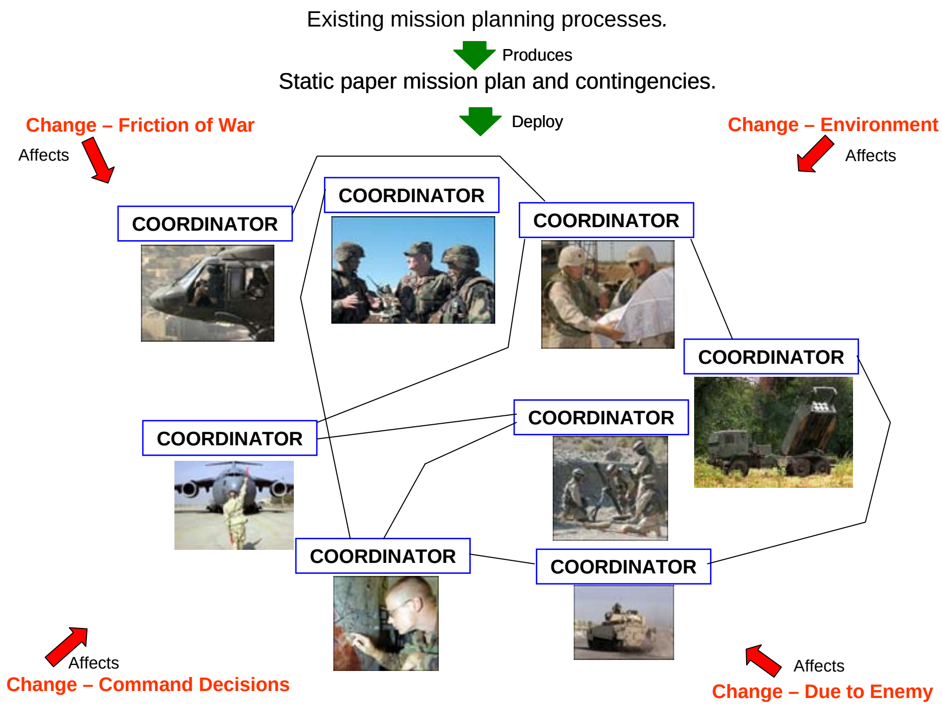

Figure 2 - With COORDINATORs Humans Can Focus On The Big Picture

COORDINATORs are about making sure the right tasks are performed by the right people, at the right times, for the current and *changing* circumstances.

To better understand the anticipated role of COORDINATORs, imagine a fielded human unit, e.g., a solider, coupled with a COORDINATOR running on a wearable computing platform. The soldier's COORDINATOR will use wireless networking technology to interact with the COORDINATORs of other soldiers to coordinate their actions. The COORDINATORs will do this by reasoning about both individual and joint action – the tasks assigned to their respective units and the temporal constraints placed on the tasks (e.g., deadlines), and how the tasks of their units interact with the tasks of other units. Note that this program is not concerned with developing new device technology or with developing new networking technologies. The necessary infrastructure either exists or will be developed elsewhere.

The term *coordination* is sometimes subject to broad interpretation. This program is not about collaboration in general rather about managing the interdependencies between the activities of different distributed parties. To illustrate the class of problems and DARPA's interest in this problem, let us consider a **hypothetical** example.

In this scenario, a political hostage is captured by a terrorist force. If the demands of the terrorists are not met by time *T*, the hostage will be killed. Intelligence indicates that the hostage is being held in one of six different locations – three land facilities and three oceangoing vessels. One possible military response is to set up a joint forces headquarters and to engage in a multi-service synchronized strike against all six targets. For the purpose of this example, we will assume this is the case. The synchronized strike is necessary because the hostage's exact location is unknown and if the different sites are forewarned they could move the hostage, kill the hostage, or be better prepared. For this mission the military deploys a company of Army Special Forces (SF), a platoon of SEALs (Navy), four Navy MK-V boats, a detachment of Air Force MH-53J troop transport helicopters, and two Air Force AC-130U gunships.

The high-level strike plan is shown in Figure 3. The Air Force MH-J helicopters will take the SF units to their respective drop points and return to base. The SF units will move into position while the SEAL teams also move into position. When all the teams are ready, they will engage. Concurrently with this the Air Force gunships will fly patterns over the region ready to lend fire support as needed.

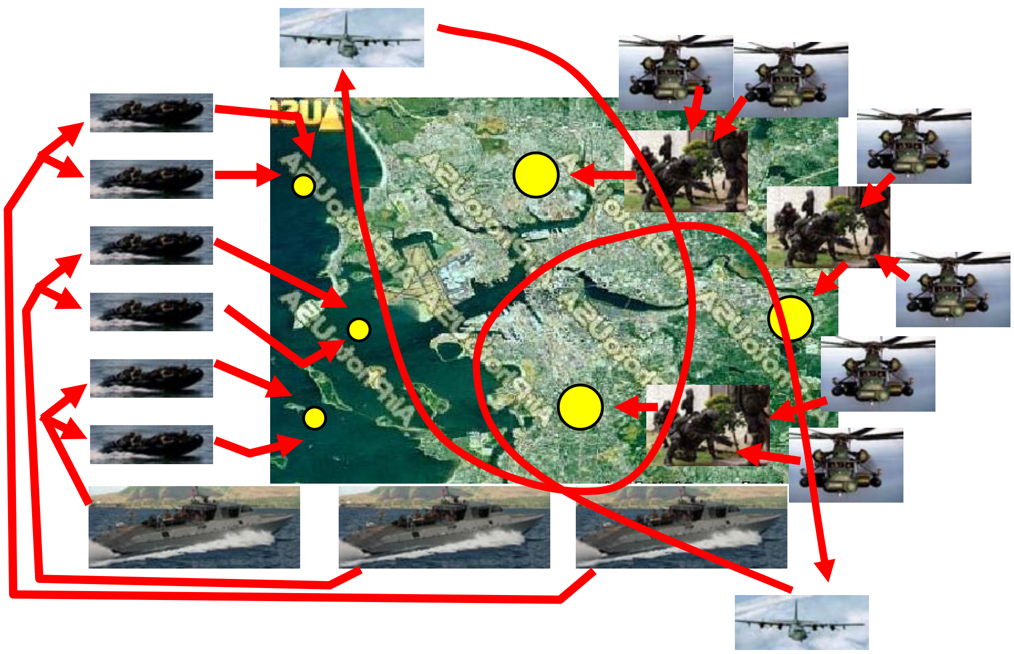

Figure 3 - High-level Strike Plan For Hostage Rescue Mission

Figure 4 shows what happens to the mission from a planning perspective. The initial mission plan is formed off-line, *a priori,* using intelligence, experience, and assumptions about how the enemy will be configured and will respond. The military has many heavyweight and effective planning processes and procedures for forming the initial plans. The output of the planning process is a set of static mission plans and a set of pre-planned contingencies. These plans are then deployed in a dynamic environment and the game changes.

After deployment, change occurs and impacts the mission – change from friction of war, the environment, the enemy, and even from command decisions. The distributed units must respond to the change online, in real-time, by adapting their mission plans and often the response is to make changes to task timings, allocations, or to select from pre-planned contingencies. This is the focus of COORDINATORs – online plan adaptation. COORDINATORs make sure the right tasks are performed by the right people at the right times for the current and changing circumstances.

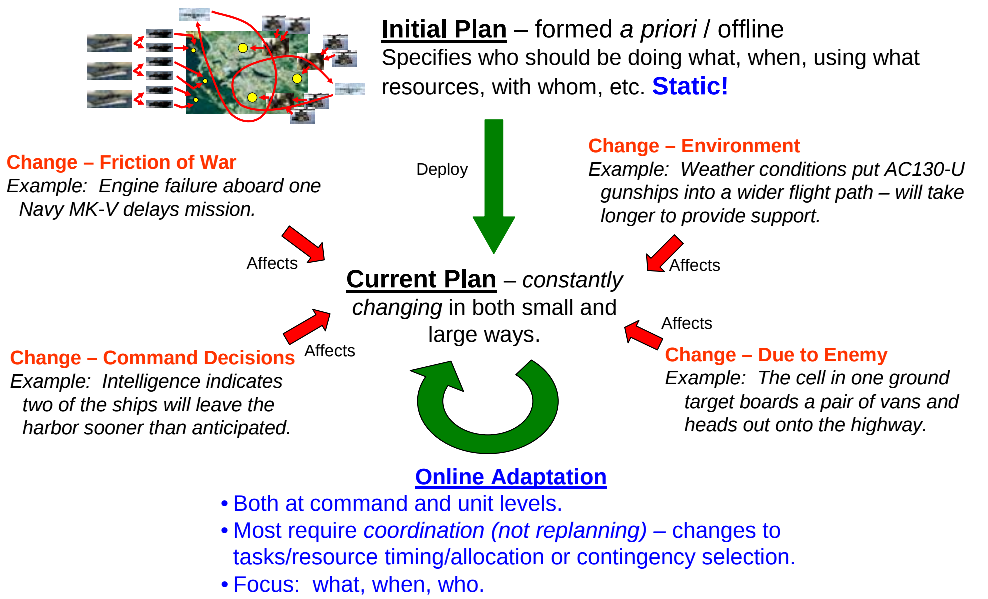

Figure 4 - The Mission Planning Process Through The Mission Life-cycle

Returning to the mission, Figure 5 shows a high-level view of the task interconnections and timings of the mission. Note that not all the team activities are spelled out or fully broken down. The general flow is that the MH-Js drop the SF teams and return to base. The SF teams move into position and when they are in position they deploy their sniper-observer units. At that same point (the land/sea operations synchronization point), the Navy SEALs, who have concurrently been moving into position on their MK-V boats, are deployed on their CRRCs (rubber rafts) to make their way stealthily toward their targets (deployment times are staggered based on distance to target). When all the teams are in place, they engage. The scheduled time for engagement is called *H-hour* and that synchronization point is also identified in the figure. Note that the teams engage before *T*, the time at which the hostage is to be killed.

Figure 5 - The Initial Plan

Figure 5 shows the plans as they are formed off-line / *a priori*. Once the teams are deployed of course, change occurs. What happens if SF team 3 is delayed, as shown in Figure 6? Today the teams must manually exchange information (the delay) and reason about the implications of any change being considered. With COORDINATORs this process would be automated. In this example, the COORDINATORs could handle the information exchange, reason about the interactions between tasks, analyze the implications of the delay, and suggest a change (eventually learning to make the decisions autonomously when appropriate). In this case, they would flag the affected tasks, and recommend a right-shift, as shown in Figure 7. Note that the scope of the change in this case is limited – all the teams have to do is to move the two land/sea operations synchronization points and they can still engage the targets before time *T*. The revised timings are shown in Figure 8. For large efforts, however, the information exchange, analysis, and option generation processes can be non-trivial.

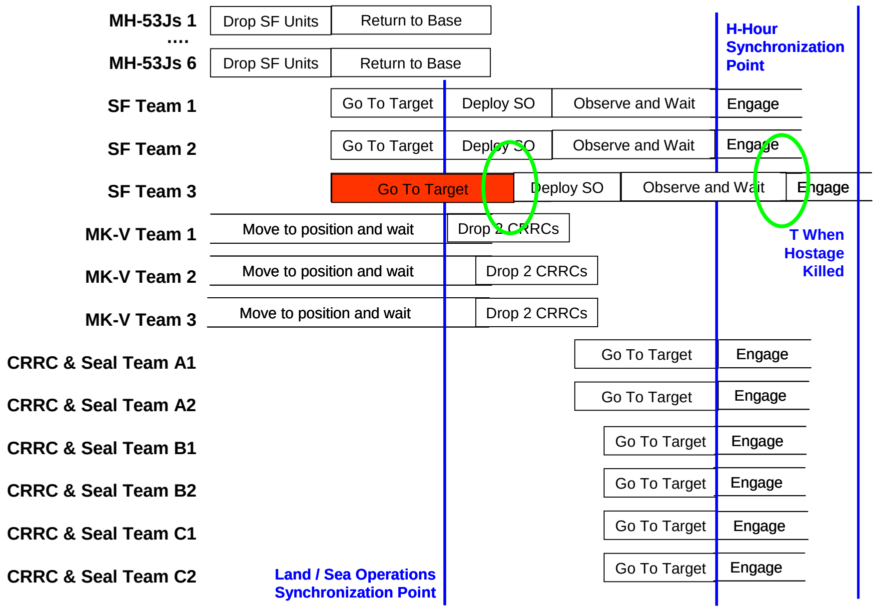

Figure 6 - SF Team 3 Is Delayed

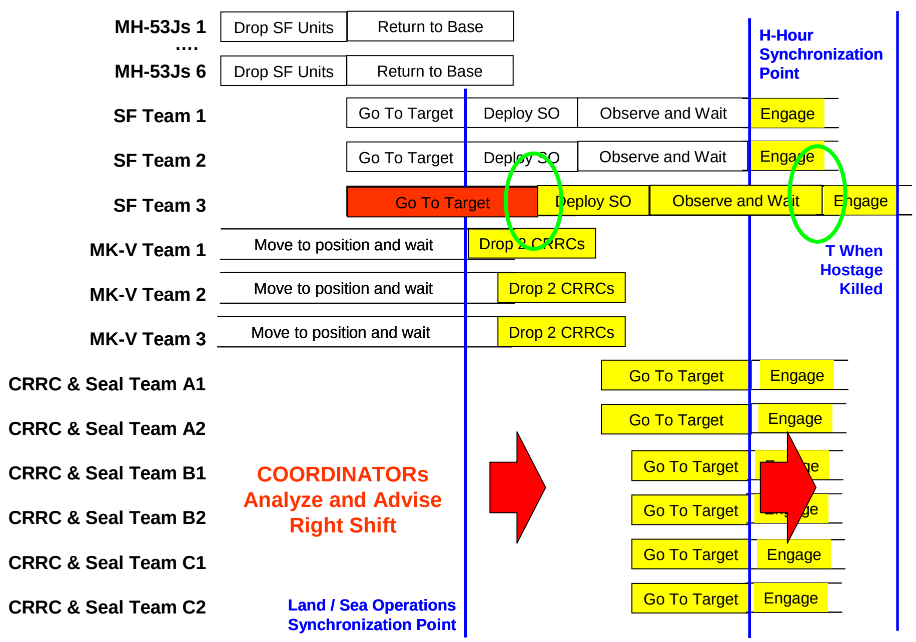

Figure 7 - COORDINATORs Analyze The Change And Suggest Revised Timings

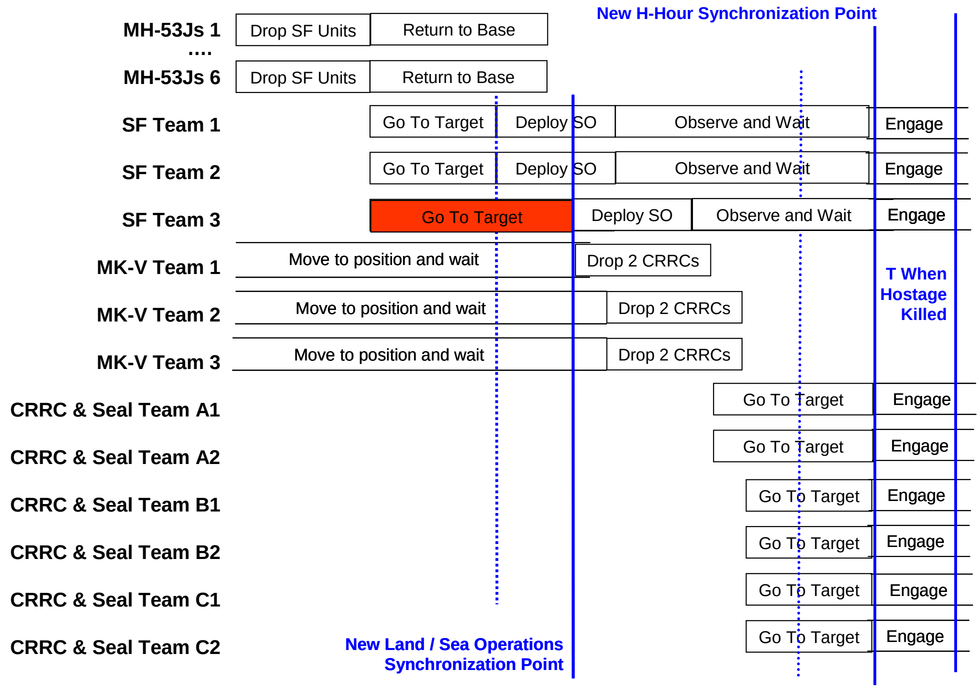

Figure 8 - Mission With Revised Timings

Now assume another change occurs – new intelligence indicates that one of the ships is preparing to leave harbor sooner than anticipated. In response, command asks the question "What happens if we move H-hour sooner?" Again, COORDINATORs could replace the manual analysis and situation verification from all the distributed units. In this case the affected tasks would be evaluated and, as shown in Figure 9 and Figure 10, the COORDINATORs could select from pre-planned contingencies to step-up the tempo of operations. The key change being that instead of deploying the SEALs on their CRRCs the SEALs would speed directly to their respective targets on the MK-V boats and deploy.

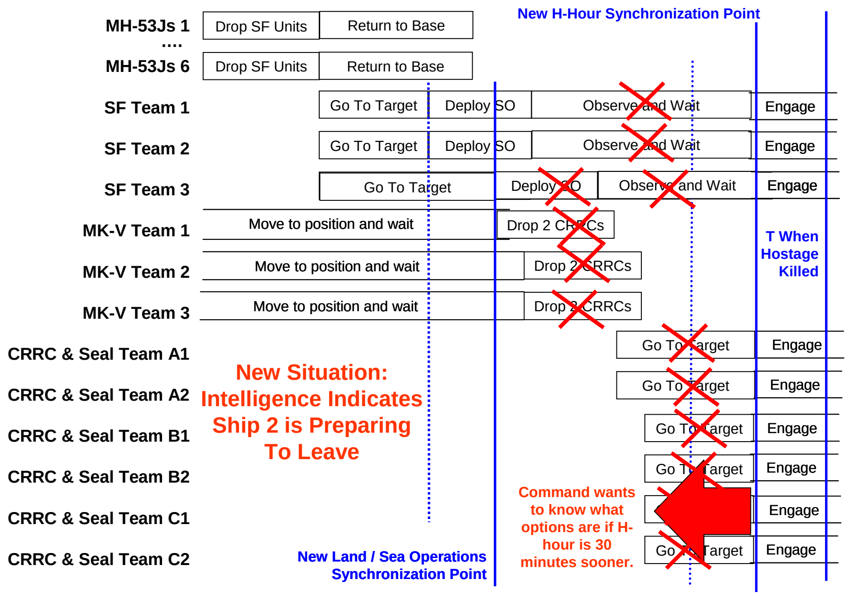

Figure 9 - COORDINATORs Perform The Analysis

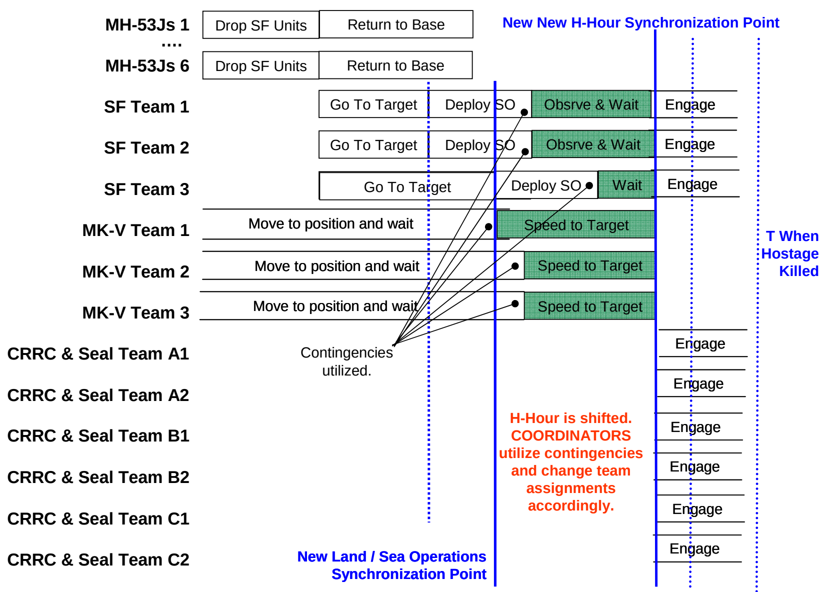

Figure 10 - COORDINATORs Adapt Mission - Engage Sooner

In this scenario we have illustrated two types of change and two adaptations. One adaptation is a *when* change or an adjustment to task timings. The other change is a *how* change where the change in how is achieved via selection from pre-planned contingencies, not from generative planning. Imagine a large scale scenario with 40-50 teams or operations scaling to 10,000 fielded units and the difficulties with detailed coordination are readily apparent.

Functional Architecture and Technical Vision

Figure 11 shows a *functional architecture* of a single COORDINATOR. The expression *functional architecture* is used instead of *prototypical architecture* because there may be many alternative ways to achieve the desired functionality – this is just one alternative. <u>Discussion will be in terms of software modules but keep in mind that as a performer, you may have an entirely different architecture or different way to achieve the desired functionality.</u>

Regardless of the actual architecture used, a distributed (or partially distributed) solution is desired[1]. The vision is having a distributed, large-scale, partially connected network of COORDINATORs where each coordinator is paired with a human unit or team and managing their activities. Other solutions may be acceptable though a strong argument for the non-distributed choice must be made (see Footnote 1 for points to address).

[1] Motivation for a distributed approach includes the fact that the authority, information, and units themselves are already distributed. Additional motivation includes avoiding having a single point of failure (read "target"), addressing likely limited node connectivity, addressing wireless/radio communication bandwidth issues, avoiding a centralized processing bottleneck, and adhering to current military preferences toward less linear command models.

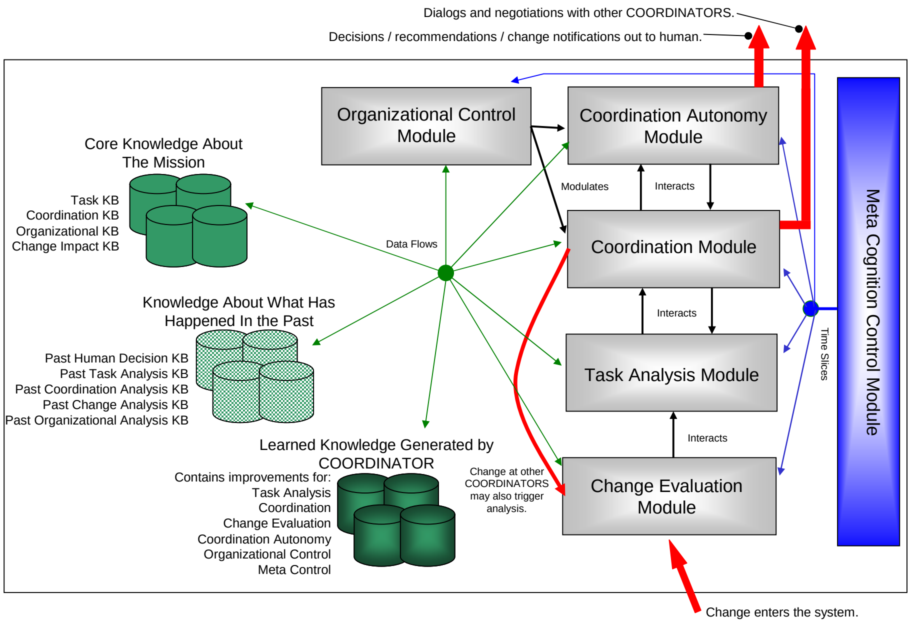

Figure 11 - Functional Architecture Of A Single COORDINATOR

In the functional architecture the drums represent data sources and the rectangles represent software modules. Note that all of the envisioned modules may be *cognitive* – having elements of representation, reasoning, learning, and self-reflection. The data sources that a COORDINATOR is likely to need include core knowledge about the mission, history of what has happened in the past (COORDINATORs are persistent), and actionable knowledge that the COORDINATOR has learned over time.

In general, change enters the system from the bottom via the change evaluation module – change event notifications are sent to this module. The change evaluation module is responsible for determining if the change requires deeper analysis or if it can be ignored. For instance, on a mission involving underwater submarine operations, a notification that it is raining above the surface may be something that can be ignored. In contrast, if requisite tank support for a tactical operation is delayed, deeper analysis is needed.

If deeper analysis is needed, the change evaluation module notifies the task analysis module. The task analysis module is envisioned as the COORDINATOR's local analysis expert – a component that reasons about the tasks assigned to the unit and how its local tasks interact with those of other units. The task analysis module should probably first do a local evaluation of the change and the way it impacts the human unit's planned course of action. For instance, it may be that the environmental change does not actually require a response or that the response is limited to the local unit (and does not impact this unit's interconnections with other units). If the change does impact other units, with which this unit is interconnected in some fashion, then the coordination module is invoked.

The coordination module is the COORDINATOR's communication and interaction expert. This is the module that is envisioned communicating with other COORDINATORs to exchange information, reason about the implications of change, and to generate response options. This module may interact with the task analysis module to evaluate different candidate responses.

The output of the coordination/task analysis module pair is a set of options for the unit to use to respond to the change that has occurred. This set of options is routed to the coordination autonomy module. This module is conceptually the "human on-board." It is this module's job to learn to rank order the options and eventually to make decisions for the human unit when he/she is unable to do so and the circumstances require a response. This is the module that interacts directly with the human unit.

The organizational control module is what makes sure the COORDINATOR follows military decision-making policies and procedures. It modulates the decisions being made by the coordination autonomy module and the communications conducted by the coordination module. Certain classes of mission change require approval one level up, others require approval two levels up, and some require local approval but information dissemination up the chain of command and so forth. The organizational control module must represent these policies and procedures and make sure that its local COORDINATOR communicates with the necessary other COORDINATORs and that the decisions being made are within the scope of the COORDINATOR's authority.

The meta-cognition module is responsible for deciding when to allocate processing time to the other modules, how much to allocate, and in what order. Recall that COORDINATORs are for on-line use, thus response time is important and regulating the (potentially exponential) problem solving of the other components is needed. The coordination module provides a good demonstration vehicle for this issue. In response to a given change the local COORDINATOR might well like to have conversations with hundreds of other COORDINATORs to collect the most complete picture possible and to generate the best possible response options. In most cases there will not be time for this kind of process (exhaustive search). It is the job of the meta-cognition module to decide when to coordinate and for how long – possibly by asking the coordination module to report its status or make its own performance estimates. The coordination module is likely to be the one to decide the details of how to spend its allocated processing time, e.g., deciding how much time to spend communicating and how much time to spend doing coordination analysis activities. This holds true for the other modules as well. Meta-cognition will provide the high level direction to keep the COORDINATOR responding to change and progressively improving via learning.

To realize the COORDINATOR vision, a number of difficult research problems must be addressed in each of the modules or functional areas. A representative set of research issues follow:

Module(s): Coordination and task analysis.

- What: functional core of the coordinator.
- Module goals/processing:
    - Perform in-depth analysis of mission tasks, interactions with other teams, intended course of action, and change(s) that have occurred.
    - Communicate with other COORDINATORs to understand full implications of change and to develop response options.
    - Generate a set of options/responses for the unit that adapts the mission to current circumstances.
- Envisioned input:
    - Mission task structures (from knowledge bases), including:
        - Specification of interactions with other units.
        - Temporal/resource constraints and interactions.
        - Pre-planned options/contingencies
    - Current intended course of action (previously generated, stored in knowledge base).
    - Change that occurred (from change evaluation module).
    - Time allocated for analysis/coordination (from meta-cognition).
- Envisioned output:
    - Set of options for the unit (to coordination autonomy module). Annotated with trade-offs of different options, downstream repercussions of particular choices.
- Hard problems: combinatorics, scale, dynamism and partial information.
- Possible solution paths:
    - Computationally adjustable coordination and analysis mechanisms.
    - Learning right adjustments to use for a given situation.
    - Learning with whom to communicate for a given problem instance or class.
    - Learning which options to generate for the humans (feedback being provided by coordination autonomy module as it learns this information) for a given problem instance or class.

Module(s): Coordination autonomy

- What: conceptually the “human on board.”
- Module goals/processing:
    - Decide how to handle the options produced by coordination/task analysis.
    - Reason about the current decision-making context and determine if module should: 1) Simply pass through all options (not enough knowledge to rank order options) and interrupt human for decision making. 2) Provide a ranked list to the human user and interrupt human for decision making. 3) Make the decision for the human unit (is it certain enough?). 4) Play for time (if unable to make a decision but the options are expiring). 5) Instruct lower levels (coordination/task analysis) to generate more options if the human wants more/different options.

    - Learning - process historical data, make coordination autonomy control rules from past experience, possibly share experience with other COORDINATORs to amass more data for learning.
- Envisioned input:
    - Set of options or responses for the unit that adapt its mission plans to the current circumstances (from coordination/task analysis).
    - Knowledge about what the human unit is currently doing, e.g., is activity interruptible for decision making? (Provided / encoded in the knowledge base – heavyweight inference combined with sensors and such will not be required).
- Envisioned output:
    - One of: list of unranked options, list of ranked options, notice that decision made by module, notice of playing for time, notice that more options being generated, etc. (all to human).
- Hard problems: complex decision space and learning to take the right action when data may be sparse and complex.
- Possible solution paths:
    - Representing a rich coordination context and decision space, e.g.: What is the human doing? Is he/she interruptible? If not now, when? How pressing is the decision? Can we play for time? How important is the decision compared to what the person is doing? What if I decide but make a wrong choice? What if I interrupt the human inappropriately? How certain am I?
    - Learning to abstract and group/cluster instances – potentially sharing with other COORDINATORs.

Module(s): Meta-cognition
- What: thinking about thinking.
- Module goals/processing:
    - Determine how much time to allocate to each module, when, and how much time to allocate to learning within the modules (learning may be an activity that should be performed when the COORDINATORs have idle time).
    - Do this by reasoning about the current state (e.g., change that occurred, mission, what the human is doing, state of coordination), and potentially abstracting and generalizing in order to compare the current state to past states / experience.
    - Learn performance profiles of the different modules by past experience.
    - Processing and allocation decisions are made on a continuous basis as the state of the world and problem solving change.
- Envisioned input:
    - COORDINATOR state information, including change being processed (if any), intended mission plans, interconnections with other units, relevant decision policies, state of all modules under its control (from KB or from modules directly).
    - History of time allocations given and performance profiles of different components (from KB).
- Envisioned output:

  - Time slices to the different modules. Each slice tagged as "primary processing" or "learning."
- Hard problems:
  - Learning to control cognitive artifacts (change + sparse data).
  - Two interconnected problems:
    - Learning to predict performance of each artifact on a given class of problem (performance may be highly problem-class dependent).
    - Adapting allocations over time as components learn and improve.
- Possible solution paths:
  - Represent rich context (mission, times allocated to components, performance).
  - Learn to abstract and group instances to cope with sparse data.
  - Learn to make allocations sans improvement-by-learning considerations.
  - Monitor and learn to adjust performance expectations as modules improve.

Module(s): Organizational control
- What: ensures the COORDINATOR follows military decision-making policies and procedures.
- Module goals/processing:
  - Reason about current state and policies. Decide which policies apply.
  - Manage the application of the policies by modulating activities of coordination module and coordination autonomy module.
- Envisioned input:
  - Specification of known decision-making policies and procedures. Four elements of this: 1) protocol / communication structure information. 2) decision-making authority information. 3) implications of not following procedure, e.g., court marshal. 4) context in which procedure applies (from KB).
  - Coordination decision or change currently being contemplated (from coordination/task analysis modules and the change evaluation module).
  - State of coordination, state of decision-making process (different steps may have different procedures/policies) (from coordination module).
- Envisioned output:
  - Communication structure / process information (to the coordination module).
  - Decision-making authority information (to the coordination autonomy module).
- Hard problems: control modulation based on organizational knowledge, learning when to circumvent when data will be sparse.
- Possible solution paths:
  - Sharing information between COORDINATORs to amass enough data for learning.
- Advanced concept to consider: Learning when to break the rules. Sometimes situations require a unit to act without following all necessary policies and procedures. It may be reasonable for the organizational control module to learn when to do this and to suggest to the other modules when a similar situation arises. Reasoning about

the importance of following the policy relative to the importance of making the decision may be appropriate.

Module(s): Change evaluation

- What: knowing what change is important and anticipating it.
- Module goals/processing:
    - Reasoning about change from the environment, enemy, friction of war, etc., and deciding when it is necessary to (re) analyze tasks and (re) coordinate.
- Envisioned input:
    - Change notifications from "information feeds" (from external sources).
    - Current mission plan and contingencies (generated during previous processing and stored in local KBs).
- Envisioned output:
    - Change notification message to task analysis module (notification more sophisticated than raw one received by this module – explicitly identify tasks or coordination decisions that are impacted).
- Hard problems:
    - Affects analysis / contextual evaluation.
    - Learning to anticipate larger change (see below).
- Possible solution paths:
    - For learning, consider sharing data with other COORDINATORs.
- Advanced concept to consider: Learning to anticipate significant events / change. It may be that major events always have smaller precursors. If the COORDINATORs could learn to identify the precursors, they could "pre" coordinate and be ready with an instant response if/when the larger event occurs.

One of the important issues to address is being able to handle potentially large and potentially complex task structures. The term *complex* is used here to denote a richer construct than AND/OR trees. Important features include the ability to group activities under tasks and specify ordering and selection functions over the activities. For instance, there might be three different ways to perform task X where each alternative way consists of some number of steps and each alternative has different performance characteristics (e.g., utility, risk, duration). Modeling task/action utility and having some way of expressing the value of a particular task/action to mission objectives is important for making choices. The models should also support temporal constraints (e.g., earliest start times, deadlines) assigned to individual actions, tasks, or entire missions. Interactions between tasks and actions should also be supported so that there is some way to explicitly reason about how the task of one performer affects the task of another performer (and the effects). For example, if team A must provide fire support to team B in order for team B to perform its task, there should be some way of representing this interaction and reasoning about it.

In this problem space, tasks and interactions are *not static* and are not always known *a priori*. The reason for this is that responding to change may require selecting from pre-planned contingencies and (re)tasking in addition to making changes to task timing. In a large-scale effort it is unreasonable for every unit to somehow be seeded *a priori* with all possible contingencies and tasks that they may be asked to perform. The technology may also be used on a continuous basis – again calling into question methods that are predicated on having complete information *a priori* about possible tasks and possible interactions. Proposers should provide a convincing case that their technical approach will support dynamic on-line addition of candidate actions and dynamic on-line change to the units with which they coordinate.

Proposers should also make a convincing case that their approach will scale beyond small tightly interconnected teams (this is the focus of year 2 / Phase II). COORDINATORs are for large-scale wide-spread deployment. Imagine 10,000 fielded units operating with coordination support. Not all of the units will need to coordinate but a given individual unit may be interconnected with many different subsets of that space of 10,000 units. Do not assume that missions, activities, and units are discretely partitioned. For instance, if unit X is part of missions M1, M2, and M3, and a change is made to a task of M1, unit X may have to coordinate with all units assigned to all missions M1, M2, and M3 (not just those of M1) because the planned change impacts X's ability (e.g., time availability) to perform mission tasks relating to M2 and M3. Note that the issue of choice and utility appears in this small example also – if all tasks cannot be performed then their value to M1, M2, and M3 must be assessed and potentially the value of M1, M2, and M3 compared with one another.

As alluded to elsewhere, real time in this program is not "hard" real time. In this program *real time* is being used to convey the concept that COORDINATORs must be fast enough for the grain-size of the application on which they are used and for the tasks over which they are coordinating. For example, in an operational level application where tasks take days to perform a reasonable response grain-size for a network of COORDINATORs is on the order of 5-10 minutes. For tactical team coordination of a small number of units over a set of missions whose total duration is a few hours, coordination response should be on the order of 1 minute. The two important facets of this discussion for proposers are that 1) they should provide a convincing case that their planned solution path can be made to function on the order of 30 seconds to some marginally acceptable upper bound, e.g., 5 minutes, and 2) that they should discuss algorithmic approaches that are time-adjustable in some fashion, i.e., given more time to respond to change (where time is set externally / defined by the application), they can do more, process more fully, etc.

Resources are not central to this program – the focus of this program is on human-to-human activity coordination. Note that certain classes of resource versus task issues are a question of modeling. For instance, if a COORDINATOR is managing the activities of a tank it could be viewed as handling the allocation/assignment of a given resource (the tank) to candidate missions. Methods from resource coordination may apply to this space though a convincing case must be made for their applicability. The expression *human activity coordination* is used deliberately in this program to communicate the richness of the problem space, the high degree of autonomy present in human systems, and the *active* role that humans will have in providing direction to the system. Other resource issues are less transformable, e.g., an aircraft not having enough of the right type of armament to perform all possible missions and the process of managing resources (missiles) to supply them to the aircraft. This latter group of resource issues will not be part of the program.

Technical Tasks and Program Structure

There are four tasks in this program. The same proposer may write for more than one task though separate proposals are required. If the proposer chooses to write for multiple tasks and Task I is part of that mix, then the proposer must articulate how the organization will establish firewall procedures to preclude conflict of interest issues and must describe the firewall procedures in its proposals. The tasks are:

Task I – Technical

The technical task contains all of the technical areas (modules or functionalities) discussed previously, i.e.:

- Coordination and task analysis.
- Meta-cognition.
- Coordination autonomy.
- Organizational control.
- Change evaluation.

As stated, performers are encouraged to consider their own approaches for addressing the technical task. Figure 11 is a reference architecture that describes the envisioned functionality. Performers may use the architecture or propose their own -- opportunities for innovation exist not only within the individual modules (desired functionalities) but within the architecture as well.

Teaming is strongly encouraged to address all of the technical areas. Each team will be expected to build a complete COORDINATOR solution, i.e., construct a deployable network of COORDINATORs. The COORDINATORs from different teams will not be integrated so that teams can focus on their own unique COORDINATOR solutions. Each team's COORDINATOR solution will, however, be integrated with a common testbed / experimentation framework to support common evaluation. The testbed will be developed under Task II. Task I performers will handle the integration of their solution with the common testbed environment. Other reasonable software specifics, e.g., the ability to read/write standard file formats for interacting with the testbed and reporting logging information, porting to demonstration platform machines, etc., will also be required of Task I performers.

The evaluation plans (later sections) contain more details that may be of interest to Task I performers. Unique proposals that address only a portion of the space above may be considered but there is a strong preference for teams that create a complete solution. Awards to multiple teams are anticipated.

Task II – Testbed Development / Integration

The technical work will be supported by the construction of an experimental framework or testbed that supports evaluation of COORDINATOR solutions by simulating the activities of human units in a dynamic environment. The testbed will be used in Phases I – IV. To supplement evaluation using the testbed, in Phases III and IV the same performer will modify the framework for use in live tactical coordination exercises using portable/wearable computing hardware.

The envisioned environment should support instantiation of a network of COORDINATORs (running in different processes on either the same machine or distributed machines) and be able to route information/change events to the COORDINATORs, monitor the communication between the COORDINATORs, simulate human interaction with the COORDINATORs, and note the solutions generated by the COORDINATORs. The testbed should support rudimentary display/graphing of the data using internal tools and should also support exportation of the data for use in third-party analysis or visualization tools. There should also be support for either visualizing Task I performer logfiles or instrumenting a given COORDINATOR to determine where its processing time is spent and to monitor the interactions between its components.

The testbed should also be able to run/evaluate skeletal COORDINATORs, i.e., those with a subset of the possible components. For instance, to test coordination/task analysis, the testbed would need to initiate and seed multiple distributed skeletal COORDINATORs, each containing a coordination/task analysis module rather than a complete COORDINATOR solution. Similarly for testing coordination autonomy, the COORDINATORs would contain only that module and the interactions with humans and other modules would be simulated. For additional details, see the evaluation plans for each of the modules in the sections that follow.

The testbed must also execute the centralized coordination oracle (Task III), collect its output, and support automated comparison between its output and the output generated by the distributed COORDINATORs. Comparison in this case will probably consist of assigning (and recording) a percentile ranking to the distributed COORDINATORs solution based on the range of solutions as enumerated by the centralized coordination oracle.

When the framework is modified in Phase III and IV for COORDINATOR use in a live tactical exercise, its role is likely to be primarily COORDINATOR network instantiation and logging / data collection. The performers of this task will handle porting their technologies to the selected portable computing platform (likely to be a mainstream computing device with a mainstream OS).

One award is anticipated for this task.

Task III – Optimal Centralized Coordination Oracle

To evaluate the distributed technologies a centralized coordination oracle is desired. Because the centralized coordination process is intractable, the oracle is envisioned being used on *small* problem instances. The role of the oracle is to produce an optimal set of schedules for each of the simulated units during experimental evaluation. These schedules will support distributed technology assessment, i.e., give us an optimal value for a mission metric (e.g., # of mission goals completed) that can be compared to the values for said metric returned by the distributed technologies. The preference is for the optimal centralized approach to engage in exhaustive search (on small problem instances) so that the complete space of solutions is identified. Having the complete space enables assignment of percentile rankings to distributed solutions. The centralized oracle technology does not have to meet the same temporal requirements as the distributed technology. However, the anticipated time to produce solutions will be considered when evaluating proposals for this particular task. We would prefer for the oracle not to be a bottleneck during testing – particularly when the problem instances are small. (To address bottleneck issues the oracle may be run off-line, *a priori*, to produce a solution set and be given very long time periods to execute.) Because the centralized oracle will be facing an intractable problem space, proposals should discuss the upper-bound on the anticipated problem space size that will be solvable by their approach. One award is anticipated for this task. Effort on this task will occur primarily in Phase I with gradually decreasing effort during Phases II through IV.

Task IV - Scenario Generation

Performers of this task will create the mission structures used for evaluation. Responsibilities include working with military experts to create realistic tactical team scenarios and working with other performers to abstract/translate the scenarios into representations/structures appropriate for the common testbed. Performers of this task will also randomize said structures in various ways as required by the evaluation plans and will be responsible for encoding the structures in various machine readable formats. An automated approach to generating many problem instances from a set of mission structures is suggested. Performers will be expected to produce both (a) information-use scenarios for unofficial testing/evaluation by the individual performers of Tasks I-III, and (b) evaluation-use scenario information for program-wide testing/evaluation. Performers may also contribute to the scenarios used in the deployed tactical team exercises. The timing of scenario generation will be driven by the evaluation schedule and the needs of the program. The evaluation plans (later sections) contain more relevant information. One award is anticipated for this task.

The program is expected to have four 12 month phases. Proposers should address all phases but the level of detail for Phases I and II should be somewhat greater than that for other phases. **<u>Initial funding will be for Phase I only.</u>** Funding will occur in accordance with the phases, i.e., assuming a successful Phase I and available resources, Phase II will then be funded, etc. The program plan is shown in Figure 12. Evaluation plans correspond with each phase – details on evaluation and metrics appear later.

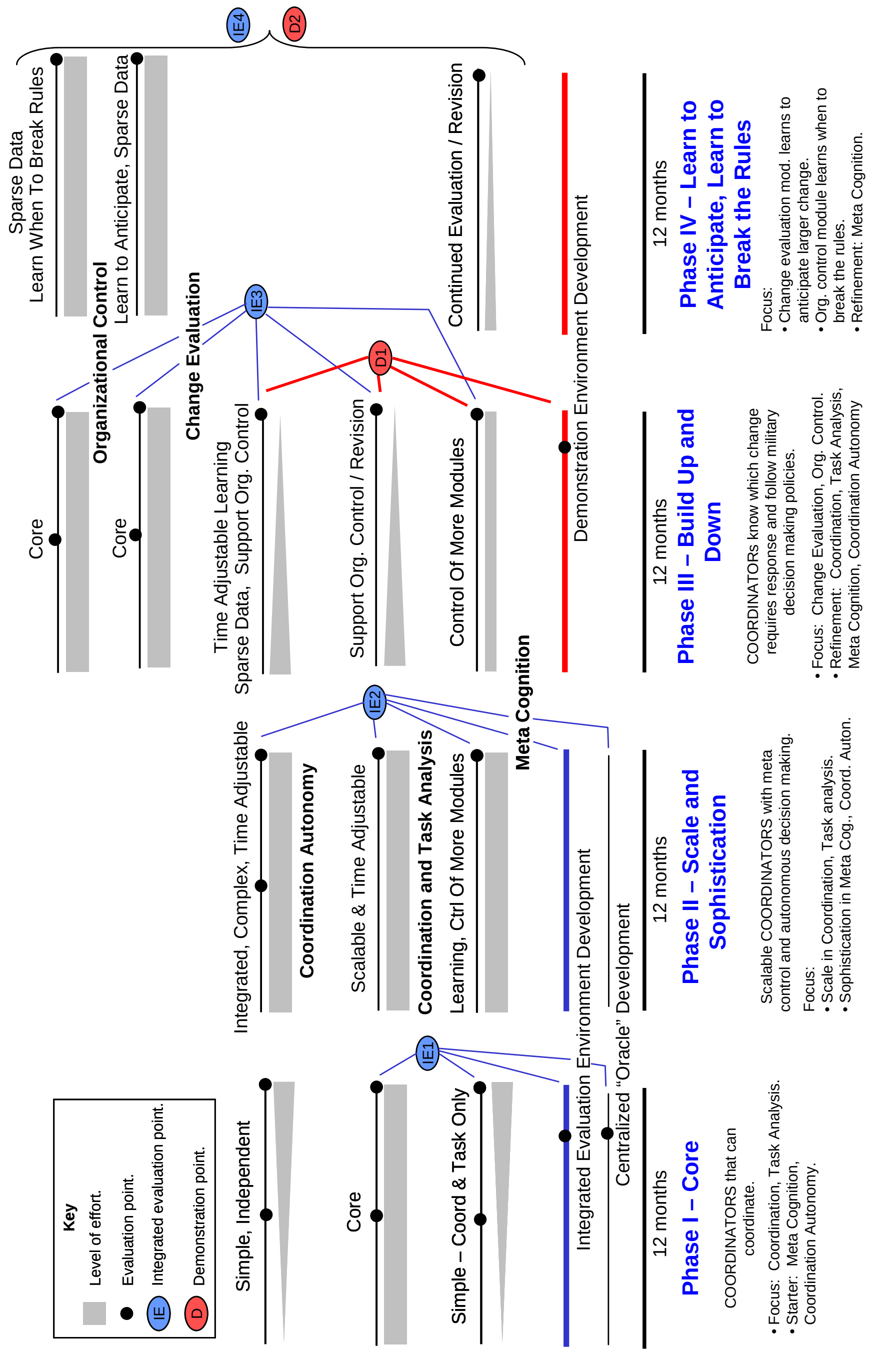

Figure 12 - Program Plan

The different phases and their goals/objectives are:

- Phase I – Concentration on the functional core of a COORDINATOR. Create a skeletal COORDINATOR that integrates coordination/task analysis and meta-cognition only. Coordination autonomy will begin during this phase also though the work may be decoupled from the coordination/task analysis problem space to simplify research if desired. The emphasis of this phase is on coordination/task analysis with preliminary approaches to meta-cognition and coordination autonomy. The goal of this phase is to show proficiency at coping with change by making changes to task timings, task assignments (including dynamic addition of new tasks), and selecting from pre-planned contingencies. A successful Phase I will increase the likelihood of a Phase II.

- Phase II – In this phase the focus is on rounding out the technologies developed in Phase I. For coordination/task analysis, this means addressing issues of scale and making the algorithms so they are time/computationally-adjustable. For coordination autonomy, this means integrating the module into the COORDINATOR (interfacing with the other modules), enriching its problem space, and also making it time-adjustable. For meta-cognition, this phase entails managing more modules and addressing the problem space more completely, e.g., learning the performance profiles of the underlying components as they change. The goal of this phase is to create COORDINATORs that can handle large problem instances and demonstrate the ability to learn to automate decision making when appropriate. A successful Phase II will increase the likelihood of a Phase III.

- Phase III – In this phase performers will continue to enhance the coordination/task analysis, coordination autonomy, and meta-cognition modules and begin work on organizational control and change evaluation. Enhancements to the pre-existing modules include: addressing sparse data (coordination autonomy), supporting organizational control (coordination autonomy, coordination/task analysis), control of more modules (meta-cognition), and general coordination improvements (coordination/task analysis). For organizational control and change evaluation, the focus is on creating the functional core of each. The goal of this phase is to create COORDINATORs that can follow military decision-making policies and procedures and are more efficient – responding to change only when necessary. A successful Phase III will increase the likelihood of a Phase IV.

- Phase IV – In this phase the focus is on advanced concepts for organizational control and change evaluation, namely learning when to break the rules / not follow military decision-making policies and procedures (organizational control) and learning to anticipate larger change (change evaluation).

Anticipated Application Space and Field Demonstrations

While the problems being solved by COORDINATORs are ubiquitous, for this program we will focus on tactical or near-tactical applications as the pay-off for coordination support in time-constrained situations is more readily assessed. Specifically, this program will concentrate on human activity coordination of fielded units in tactical settings, e.g., tactical team commander coordination in a complex mission space. In this space, one might imagine fielded units with wearable computing platforms interacting with their COORDINATORs via heads-up displays and speech input. Note that the focus of this program is on coordination technology not the devices or hands-free interfaces, etc. The latter are beyond the scope of this program. For the demonstrations, off-the-shelf portable computing devices (e.g., tablet PCs) with mainstream operating systems are likely to be used.

Evaluation and Metrics

Figure 12 identifies the placement of evaluation points. Evaluations will be both *in-the-small* (modules separated from entire COORDINATOR framework) and *in-the-large* (across an integrated COORDINATOR solution with multiple COORDINATORs executing). Evaluation plans may change/evolve with the program, however, anticipated plans appear in the following tables. Proposers are welcome to propose different evaluation plans and metrics.

In the event of budgetary or other considerations, performance rankings generated by evaluations may be used as a basis for down-selection or may serve as a component of the down-selection criteria.

Coordination and Task Analysis Evaluation	
Program Month	Evaluation Description
6	What: *Documentation of planned approach.* Performers will document their planned approach and provide a convincing argument that their algorithms will address requirements and will scale. The metric is plausibility as determined by PM and/or PM's consultants.
12	What: *Verify that core coordination algorithms are functioning.* Each team's algorithms/modules will be empirically tested in the common test environment using a series of N experiments with increasing levels of difficulty (more nodes, connections, changes), each having M trials. The metric is optimality, the goal is to achieve 90%-tile or better given the range of possible solutions as determined by the centralized coordination oracle.
24	What: *Verify that algorithms are time-adjustable (primary) and scale (secondary).* • Time-adjustable – Each team's algorithms/modules will be tested in the common test environment. As above, there will be M experiments, with N trials, where each experiment is of increasing difficulty. Each experiment will be performed K times with different time allocations.

	Degradation in solution quality (as time decreases) will be computed by comparing solutions against the solution space defined by the centralized coordination oracle. The rate of performance degradation will be compared across all performers. The metric is optimality (as in the 12 month experiment); the interested property is degradation in solution quality. Teams may be performance ranked for comparison. • Scalable – Each team's algorithms will be tested in the common test environment. As above, there will be M experiments, each having N trials, where the number of tasks in each experiment grows by a factor of X, where X is determined during the program. A likely X is 2 where the number of tasks increases from 32, to 64, to 128, to 256, to 512. The performance of each solution is compared to optimal as in the 12 month experiment. The time required to produce a solution in each case is recorded but is informational at this stage. The metric is optimality; the goal is 90%-tile. In the likely event that the centralized scheduling oracle will not scale to the desired range of experiments, even given large amounts of time, this evaluation will migrate to a relative performance ranking of the teams.
36	What: *Harder test on scalability (combining scale and time adjustability) and test support for organizational control.*

Table 1 - Coordination / Task Analysis Evaluation

Coordination Autonomy Evaluation	
Program Month	Evaluation Description
6	What: *Documentation of planned approach.* Performers will document their planned approach and provide a convincing argument that their algorithms will learn / address requirements and will be useful when data is sparse. The metric is plausibility as determined by PM and/or PM's consultants.
12	What: *Verify learning and assess sparse data handling.* • Verify learning – The goal is to verify that each team's algorithms/modules can learn to automate simple decision making. In this set of tests, the possible actions for the module will be limited to: *decide* and *do not decide*. (If making the decision, the module must also select from the candidate option set.) A partitioned experimental approach will be used. The metric is error rate. Teams may be performance ranked for comparison. • Assess sparse data handling – This is an informational test on error rate increase when training data is increasingly sparse. The metric is error rate change. Teams may be performance ranked for comparison.
18	What: *Verify module is integrated and verify that module can make a richer set of decisions.* • Integration check – This is a Boolean test to verify that the module is able to

	interact with the meta-controller (receive time slices for learning or decision making) and to provide high level direction to the coordination/task analysis module. • Richer decision space evaluation – This is similar to the 12 month test but in this case the possible actions for module are: *provide human with unranked list, provide human a ranked list,* and *make decision* (if the certainty is over a defined threshold). For these experiments we will assume that the human is interruptible. As in the 12 month experiments these will be partitioned experiments and the metric will be error rate. Teams may be performance ranked for comparison.
24	What: *Verify richer decision making, verify time-adjustability, informational check on sparse data handling.* • Verify that modules can make decisions over complete space – These experiments follow in the lines of previous ones but in this case the possible actions for the module are: *provide the human with an unranked list, provide the human with a ranked list, make a decision for the human, play for time to give human more time to respond,* and *request more options of the lower-level components*. In this set of experiments, the interruptability of the human's current task will be factored in to the decision process where interruptability will be marked as *low*, *medium*, or *high*. As above, partitioned experiments will be used. In this case the experiments may be factored based on the state of the human, e.g., 5 decisions, 3 human states = 15 experiments in total. The metric is error rate. Teams may be performance ranked for comparison. • Time adjustability in decision making – In these experiments, the factored experiments of above will be repeated N times each with different time allocations given to decision making. Any degradation in decision quality (as time decreases) is measured and the change in error rate compared across all performers. The metric is error rate. The interested property is degradation. The goal is to show no degradation during decision making as we do not anticipate the act of deciding to be computationally expensive. • Informational test on sparsity handling (see 12 month test).
36	What: *Test on sparsity handling, test of support for organizational control, test on time adjustability of learning algorithms.*

Table 2 - Coordination Autonomy Evaluation

Meta-cognition Evaluation	
Program Month	Evaluation Description
6	What: *Documentation of planned approach.* Performers will produce a two page written discussion of their planned approach, predict strengths / weaknesses, and specify which COORDINATOR state elements are being used by the approach. The metric is plausibility as determined by PM

	and/or PM’s consultants.
12	What: *Basic test of core algorithms.* Experiments will assess how well the meta-cognition module can identify and learn the performance profiles for the different modules where the performance profiles of the other modules are simulated – generated off-line – and a partitioned experiment used to assess meta-cognition learning. The module set will be limited to coordination and task analysis. The metric is % error rate. Teams may be performance ranked for comparison.
24	What: *Integrated test with real data.* In this set of experiments the performance of the meta-controller will be assessed by comparing the solutions generated by a given team’s network of COORDINATORs in a time-constrained situation against that same network in a non-time-constrained situation. To accomplish this, we will perform a set of N integrated experiments with live coordination and task analysis components. Each experiment will test the performance of the meta-control module for a class of problems (kind of mission, coordination problem, change, etc.). In each experiment, M trials, each belonging to the same class, will be tested. Before each experiment, the coordination autonomy module will be trained to make decisions over the class of problems being tested (so that it can run without humans in the loop). During the experiments, the system will be tested twice with each trial. In one case, the system will be given the trial and an accompanying time deadline, e.g., 30 seconds, by which a response is required. In the other case the system will be given the trial and given a relatively long time, e.g., five minutes, to produce a result (for all practical purposes this will be a non-time-constrained case). The results produced in each case will be compared to one another. Teams will be scored on the quality of the solution produced in the time-constrained case relative to the solution produced in the non-time-constrained case. Teams may also be scored on the quality of their solutions relative to those produced by the coordination oracle, however, the focus is on comparison between the same artifacts to isolate meta-control issues from coordination/task analysis performance. Teams may be performance ranked for comparison.
36	What: *Integrated test with more modules.*
48	What: *Integrated test modified to induce more change in the performance profiles of the other modules.*

Table 3 - Meta-cognition Evaluation

Organizational Control Evaluation	
Program Month	Evaluation Description
30	What: *Documentation of planned approach.* Performers will produce a two page written discussion of the algorithmic approach being used to modulate coordination and decision making. Document will include

	a specification of the representations being used for policy/procedure information and a discussion of algorithmic plans for learning to break the rules. The metric is plausibility as determined by PM and/or PM's consultants.
36	What: *Basic test of core functionality.* These experiments will verify that the module modulates coordination and decision making as appropriate. In the experiments, the COORDINATORs network will be instrumented so that the communication flow can be monitored. For each policy/procedure class, we will verify that the network follows the proper directives by testing it with N (e.g., 32) trials that belong to said policy/procedure. In each trial, the communication flows and decisions made will be recorded. Performers will be scored on the number of decisions made inappropriately (beyond the scope of authorized decision making) and the number of times communication flows were not up to specification. For each experiment, the metric is error rate in two dimensions (decisions/authority and coordination). Teams may be performance ranked for comparison.
48	What: *Learning to break the rules and learning to cope with sparse data*. • Learning to break the rules – These experiments will verify that the module is learning when to break the rules. To simplify the experiments, we will assume that the coordination autonomy module has already learned to make a given decision if it is given the authority to do so. The module will first be trained with data from each class of policy that it knows. To do this, for each of N policies, a set of randomly generated problem instances will be created and marked, by hand, to indicate where rules should be broken and instances where rules should not be broken. The problem instances will include information about timeliness of the given decision and other mission characteristics necessary for decision making. The module will then be trained with a portion of the data and tested with the other. The metric is decision error rate. Teams may be performance ranked for comparison. • Learning to cope with sparse data – See "Learning to make decisions with sparse data" in the coordination autonomy evaluation section. The evaluation is similar here.

Table 4 - Organizational Control Evaluation Plan

Change Evaluation Module Evaluation Plan	
Program Month	Evaluation Description
30	What: *Documentation of planned approach.* Performers will produce a two page written discussion of their planned approach to identify when tasks, coordination decisions, etc., are impacted by change. Document will include specification of information that will be passed from the change evaluation module to the task analysis module and a discussion of algorithmic plans for learning to anticipate change. The metric is plausibility as determined by PM and/or PM's consultants.
36	What: *Basic test of core functionality.*

	The goal of these experiments is verify that the module correctly passes relevant change notifications "up" and correctly filters irrelevant notifications. To test this capability, we will first create a set of M mission structures complete with commitments (agreements) with other teams. For each of the M mission structures, a list of impact points will be created – points where the plans can be affected by outside change. We will then iterate over the impact points for each mission structure and create N randomized change notifications that affect the impact points. Similarly, we will create R randomized change notifications that are orthogonal to the impact points. The resulting data will form M experiments with N+R trials (N of which module should "pass up," R of which module should ignore). The module will be tested with the data. The metric is error rate tabulation on an impact point basis. Performers should be without error.
48	What: *Learning to anticipate change, learning to cope with sparse data.* • Learning to anticipate change – The goal is to verify that module learns to correlate small changes with subsequent larger changes. To do this, we will create data sets in which smaller changes are followed by larger changes with varying temporal distance between the two changes and with varying whether or not the smaller changes are impactful. The module will be tested using a partitioned experiment approach. The metric is error rate. Teams may be performance ranked for comparison. • Learning to cope with sparse data – See "Learning to make decisions with sparse data" in the coordination autonomy evaluation plan. Evaluation is similar here.

Table 5 - Change Evaluation Module Evaluation Plan

Evaluation plans for in-the-large focus on the integrated system as a whole. These are identified as IE (Integrated Experiment) points on Figure 12. Demonstration points (D) also indicate in-the-large / cohesive system evaluations though the focus is not on detailed evaluation but providing feedback from actual application deployment. IE and D evaluation sketches follow.

Integrated Evaluation Plans		
Program Month	Evaluation Point	Description
12	IE1	What: *Simple integration test.* Tests will verify integration of early meta-cognition controller with coordination and task analysis. Test will consist of simple Boolean checks on data passing and control modulation.
24	IE2	What: *Test existence of a capable, functioning core.* In these experiments, each COORDINATOR will include the meta-cognition, coordination autonomy, and coordination/task analysis modules. For the evaluation, each team's COORDINATOR network

		will be tested against generated coordination instances (hand designed missions combined with randomization of events/changes). The metric is proximity to optimal response as determined by the centralized coordination oracle (see coordination/task analysis module evaluation plans, month 12). Goal is to be 90%-tile or better.
36	IE3	What: *Test improved core plus organizational reasoning and change evaluation.* As with IE2 but this round will also test the organizational control module and the change evaluation module. Tests will evaluate adherence to decision-making policies and procedures and change filtering capabilities. For instance, in some test cases COORDINATORs that properly ignore noise events will conserve computation for important changes and will have more time to respond to the important changes – possibly producing a higher quality result. The metric is proximity to optimal (see IE2). The goal is to be 90%-tile or better.
48	IE4	What: *Test improved core with enhanced abilities.* As with IE3 but this round will include evaluation of learning to anticipate change and learning when to break the rules. For instance, in some test cases COORDINATORs that can anticipate change will have more time to prepare and may be more likely to produce a higher quality final result. Metric is proximity to optimal (see IE2). Goal is to be 90%-tile or better.

Table 6 - Integrated Evaluation Plans

Demonstration Plans		
Program Month	Evaluation Point	Description
36	D1	What: *Use of COORDINATORs in basic tactical training exercise.* The goal is to produce anecdotal/hands-on evaluation by military units engaged in a tactical training exercise. One team (more if possible) will be selected to have their COORDINATORs deployed in a live exercise. Selection criteria may include team performance on the other tests. If possible, the exercise will be performed twice – once with and once without coordination support and the performance in each exercise will be compared.
48	D2	What: *Use of COORDINATORs in a more complex training exercise.* As with D1 though missions will be designed to be more challenging, larger scale, and require a faster operations tempo.

Table 7 - Demonstration Plans

PROGRAM SCOPE.

Proposed research should investigate innovative approaches and techniques that lead to or enable revolutionary advances in the state-of-the-art. Proposals are not limited to the specific strategies listed above, and alternative visions will be considered. However, proposals should be for research that substantially contributes towards the goals stated. Specifically excluded is research that primarily results in minor evolutionary improvement to the existing state of practice or focuses on special-purpose systems or narrow applications.

GENERAL INFORMATION

Proposals not meeting the format described in this pamphlet may not be reviewed. Proposals **MUST NOT** be submitted by fax or e-mail; any so sent will be disregarded. This notice, in conjunction with the BAA 04-29 FBO Announcement and all references, constitutes the total BAA. A Frequently Asked Questions (FAQ) list may be provided. The URL for the FAQ will be specified on the DARPA/IPTO BAA Solicitation page. No additional information is available, nor will a formal Request for Proposal (RFP) or other solicitation regarding this announcement be issued. Requests for same will be disregarded. All responsible sources capable of satisfying the Government's needs may submit a proposal that shall be considered by DARPA. Historically Black Colleges and Universities (HBCUs) and Minority Institutions (MIs) are encouraged to submit proposals and join others in submitting proposals. However, no portion of this BAA will be set aside for HBCU and MI participation due to the impracticality of reserving discrete or severable areas of this research for exclusive competition among these entities.

Security classification guidance on a DD Form 254 (DoD Contract Security Classification Specification) will not be provided at this time since DARPA is soliciting ideas only. After reviewing incoming proposals, if a determination is made that contract award may result in access to classified information, a DD Form 254 will be issued upon contract award. **If you choose to submit a classified proposal you must first receive the permission of the Original Classification Authority to use their information in replying to this BAA.**

SUBMISSION PROCESS

This Broad Agency Announcement (BAA) requires completion of a **BAA Cover Sheet** for each Proposal prior to submission. This cover sheet can be accessed at the following URL:

http://www.dyncorp-is.com/BAA/index.asp?BAAid=04-29

After finalizing the **BAA Cover Sheet**, the proposer must print the **BAA Confirmation Sheet** that will automatically appear on the web page. Each proposer is responsible for printing the BAA Confirmation Sheet and attaching it to every copy. The Confirmation Sheet should be the first page of the Proposal. If a proposer intends to submit more than one Proposal, a

Deleted: The full ACIP program is anticipated to comprised of three phases: 1) Early Architecture Concepts and In-Context Evaluation, 24 months; 2) Full-Scale Implementation and Demonstration, 48 months; 3) Cognitive Technology System Transitions to DoD, 24 months.¶
Phase I, the focus of this announcement has the following expected deliverables:¶
<#>Architecture and runtime requirements derived from a comprehensive spanning set of current and emerging reasoning, knowledge, and learning components that are representative of the types and scale of future candidate DoD cognitive information processing applications;.¶
<#>Participation in the Living Framework Fforum and a draft Living Framework document;¶
<#>Architecture concepts, models and supporting analysis;¶
<#>Composable runtime concepts;¶
<#>Device, software, and user development environment concepts and technology roadmap;¶
<#>Suggested Phase II challenge problem, evaluation metrics and success criteria;¶
<#>Interim and final technical concept description document;.¶
<#>DARPA/IPTO Quarterly Status Reports and Annual Project Summary Report.¶

Deleted: PROGRAM OBJECTIVES AND DESCRIPTION. The Defense Advanced Research Projects Agency (DARPA) is soliciting proposals for DARPA's Information Processing Technology Office to perform research, requirements and constraint analysis, architecture concept development and design, architectural modeling, in-context evaluations, and concept evaluations to support the initial phase of the Architectures for Cognitive Information Processing (ACIP) program. It is the intent of the DARPA IPTO office to develop cognitive information processes that will bring enabling embedded intelligence capabilities to aid the warfighter, as well as DoD supporting functions and activities – enabling machines that think to aid human performance. Current intelligent processing implementations depen ... [1]

Deleted: 03-xx

Deleted: 3

Deleted: 44

Deleted: ing

Deleted: 03-xx

Deleted: 3

Deleted: 44

unique UserId and password must be used in creating each BAA Cover Sheet. Failure to comply with these submission procedures may result in the submission not being evaluated.

Proposers must submit the original and 2 copies of the full proposal (3 total) *and* 2 electronic copies (i.e., 2 separate disks) of the full proposal (in PDF or Microsoft Word 2000 for IBM-compatible format on a 3.5-inch floppy disk, 100 MB Iomega Zip disk or cd). **Mac-formatted disks will not be accepted**. Each disk must be clearly labeled with BAA 04-29, proposer organization, proposal title (short title recommended) and "Copy <n>___ of 2". The full proposal (original and designated number of hard and electronic copies) must be submitted in time to reach DARPA by 12:00 PM (ET) **August 23, 2004**, in order to be considered during the initial evaluation phase. However, **BAA 04-29, COORDINATORs** will remain open until 12:00 NOON (ET) **June 1, 2005**. Thus, proposals may be submitted at any time from issuance of this BAA through **June 1, 2005**. While the proposals submitted after the **August 23, 2004**, deadline will be evaluated by the Government, proposers should keep in mind that the likelihood of funding such proposals is less than for those proposals submitted in connection with the initial evaluation and award schedule. DARPA will acknowledge receipt of submissions and assign control numbers that should be used in all further correspondence regarding proposals.

Deleted: 03-xx
Deleted: 4
Deleted: Wednesday
Deleted: Friday, November 26Feb
Deleted: 3
Deleted: 03-xx
Deleted: 3-44
Deleted: SR
Deleted: September 24
Deleted: January
Deleted: 4
Deleted: September 24
Deleted: Janua
Deleted: 4
Deleted: 3

Restrictive notices notwithstanding, proposals may be handled for administrative purposes by support contractors. These support contractors are prohibited from competition in DARPA technical research and are bound by appropriate non-disclosure requirements. Input on technical aspects of the proposals may be solicited by DARPA from non-Government consultants /experts who are also bound by appropriate non-disclosure requirements. However, non-Government technical consultants/experts will not have access to proposals that are labeled by their offerors as "Government Only". Use of non-government personnel is covered in FAR 37.203(d).

REPORTING REQUIREMENTS/PROCEDURES: The Award Document for each proposal selected and funded will contain a mandatory requirement for submission of DARPA/IPTO Quarterly Status Reports and an Annual Project Summary Report. These reports, described below, will be electronically submitted by each awardee under this BAA via the DARPA/IPTO Technical – Financial Information Management System (T-FIMS).

The T-FIMS URL will be furnished by the government upon award. Detailed data requirements can be found in the Data Item Description (DID) DI-MISC-81612A available on the Government's ASSIST database (http://assist.daps.dla.mil/quicksearch/). Sample instructions that specify how information in the DID may be collected (content and frequency requirements) can be found in Appendix A. An outline of T-FIMS report requirements is as follows:

(a) Status Report: Due at least three (3) times per year – Jan, Apr, & Oct
  1) Technical Report
    a) Project General Information
    b) Technical Approach
      - Accomplishments

- Goals
- Significant changes / improvements

c) Deliverables
d) Transition Plan
e) Publications
f) Meetings and Presentations
g) Project Plans
h) Near term Objectives

2) Financial Report
3) Project Status / Schedule

(b) Project Summary (PSum): Due once each fiscal year in July

1) All Sections of the Status Report
2) QUAD Chart
   a) Visual Graphic
   b) Impact
   c) New Technical Ideas
   d) Schedule

PROPOSAL FORMAT

Proposals shall include the following sections, each starting on a new page (where a "page" is 8-1/2 by 11 inches with type not smaller than 12 point) and with text on one side only. The submission of other supporting materials along with the proposal is strongly discouraged. Sections I and II (excluding the submission cover/confirmation sheet and section M) of the proposal shall not exceed 39 pages. Maximum page lengths for each section are shown in braces { } below.

Deleted: 6

Deleted: 4

Section I. Administrative

The BAA Confirmation Sheet {1 page} described under “Submission Process” will include the following:

A. BAA number;
B. Technical topic area;
C. Proposal title;
D. Technical point of contact including: name, telephone number, electronic mail address, fax (if available) and mailing address;
E. Administrative point of contact including: name, telephone number, electronic mail address, fax (if available) and mailing address;
F. Summary of the costs of the proposed research, including total base cost, estimates of base cost in each year of the effort, estimates of itemized options in each year of the effort, and cost sharing if relevant;
G. Contractor's type of business, selected from among the following categories: "WOMEN-OWNED LARGE BUSINESS," "OTHER LARGE BUSINESS," "SMALL DISADVANTAGED BUSINESS [*Identify ethnic group from among the following:*

Asian-Indian American, Asian-Pacific American, Black American, Hispanic American, Native American, or Other]," "WOMEN-OWNED SMALL BUSINESS," "OTHER SMALL BUSINESS," "HBCU," "MI," "OTHER EDUCATIONAL," "OTHER NONPROFIT", or "FOREIGN CONCERN/ENTITY."

Section II. Detailed Proposal Information

This section provides the detailed discussion of the proposed work necessary to enable an in-depth review of the specific technical and managerial issues. Specific attention must be given to addressing both risk and payoff of the proposed work that make it desirable to DARPA.

[IMPORTANT NOTE: WITH THE EXCEPTION OF E, C THROUGH H HAVE BEEN REVISED.] Page-counts are maximums.

A. {1 Page} Innovative claims for the proposed research.
This page is the centerpiece of the proposal and should succinctly describe the unique proposed contribution.

B. {1 Page} Proposal Roadmap
The roadmap provides a top-level view of the content and structure of the proposal. It contains a synopsis (or "sound bite") for each of the nine areas defined below. It is important to make the synopses as explicit and informative as possible. The roadmap must also cross-reference the proposal page number(s) where each area is elaborated. The nine roadmap areas are:

1. Main goals of the proposed research (stated in terms of new, operational capabilities for assuring that critical information is available to key users).

2. Tangible benefits to end users (i.e., benefits of the capabilities afforded if the proposed technology is successful).

3. Critical technical barriers (i.e., technical limitations that have, in the past, prevented achieving the proposed results).

4. Main elements of the proposed approach.

5. Rationale that builds confidence that the proposed approach will overcome the technical barriers. ("We have a good team and good technology" is not a useful statement.)

6. Nature of expected results (unique/innovative/critical capabilities to result from this effort, and form in which they will be defined).

7. The risk if the work is not done.

8. Criteria for scientifically evaluating progress and capabilities on an annual basis.

9. Cost of the proposed effort for each performance year.

C. {2 Pages} Research Objectives:

1. Problem Description. Provide concise description of problem area addressed by this research project.

2. Research Goals. Identify specific research goals of this project. Identify and quantify expected performance improvements from this research. Identify new capabilities enabled by this research. Identify and discuss salient features and capabilities of developmental hardware and software prototypes**.**

Deleted: **Provide a set of metrics and sucesssuccess criteria for the concepts proposed under Phase I.**

3. Expected Impact. Describe expected impact of the research project, if successful, to problem area.

D. Technical Approach:

1. {12 Pages} Detailed Description of Technical Approach. Provide detailed description of technical approach that will be used in this project to achieve research goals

Deleted: **. Specifically identify and discuss the innovative aspects of the ACIP technical approach for . two or more diverse (full scale) reasoning techniques integrated as part of a complete cognitive system comprised of reasoning, knowledge representation, and learning subsystems. This section should be well motivated and should clearly articulate the need for innovative architecture advances in-context of full scale applications, the proposed innovative solution, and the payoff relative to todays'today's COTS computing solutions. new idea and reasoning/evidence giving confidence that the idea can work. Full spectrum (major classes) of techniques for reasoning, knowledge representation, and learning and their architectural impact should be discussed. Note: An optional technical viewgraph summary in MS Power Point format (maximum of 8 vgs) may also be included as part of the Technical Volume and will not be considered as part of the volume page count.¶**
<#>¶

2. {2 Pages} Comparison with Current Technology. Describe state-of-the-art approaches and the limitations within the context of the problem area addressed by this research.

E. {3 Pages} Statement of Work (SOW) written in plain English, outlining the scope of the effort and citing specific tasks to be performed, references to specific subcontractors if applicable, and specific contractor requirements.

Deleted:

F. Schedule and Milestones:

1. {1 Page} Schedule Graphic. Provide a graphic representation of project schedule including detail down to the individual effort level. This should include but not be limited to, a multi-phase development plan, which demonstrates a clear understanding of the proposed research; and a plan for periodic and increasingly robust experiments over the project life that will show applicability to the overall program concept. Show all project milestones. Use absolute designations for all dates.

2. {3 Pages} Detailed Individual Effort Descriptions. Provide detailed task descriptions for each individual effort and/or subcontractor in schedule graphic.

Deleted: 2

G. {2 Pages} Deliverables Description. List and provide detailed description for each proposed deliverable. Include in this section all proprietary claims to results, prototypes, or systems supporting and/or necessary for the use of the research, results, and/or prototype. If there are no proprietary claims, this should be stated. The offeror must submit a separate list of all technical data or computer software that will be furnished to the Government with other than unlimited rights (see DFARS 227.) Specify receiving organization and expected delivery date for each deliverable.

H. {1 Page} Technology Transition and Technology Transfer Targets and Plans. Discuss plans for technology transition and transfer. Identify specific military and commercial organizations for technology transition or transfer. Specify anticipated dates for transition or transfer.

Deleted: If software developed by the project will not be released under an Open Source license, provide clear reasoning showing that the technology transition plan is likely to be more successful than Open Source would be at making the software available to interested researchers and commercial enterprises.

I. {3 Pages} Personnel and Qualifications. List of key personnel, concise summary of their qualifications, and discussion of proposer's previous accomplishments and work in this or closely related research areas. Indicate the level of effort to be expended by each person during each contract year and other (current and proposed) major sources of support for them and/or commitments of their efforts. DARPA expects all key personnel associated with a proposal to make substantial time commitment to the proposed activity.

Deleted: 2

J. {1 Page} Facilities. Description of the facilities that would be used for the proposed effort. If any portion of the research is predicated upon the use of Government Owned Resources of any type, the offeror shall specifically identify the property or other resource required, the date the property or resource is required, the duration of the requirement, the source from which the resource is required, if known, and the impact on the research if the resource cannot be provided. If no Government Furnished Property is required for conduct of the proposed research, the proposal shall so state.

K. {1 Page} Experimentation Plans. Offerors should identify experiments to test the hypotheses of their approaches – these may be in addition to those expressed in this PIP and/or alternative suggested approaches. Offers should be willing to work with other contractors in order to develop joint experiments in a common testbed environment. Offerors should expect to participate in teams and workshops to provide specific technical background information to DARPA, attend semi-annual Principal Investigator (PI) meetings, and participate in numerous other coordination meetings via teleconference or Video Teleconference (VTC). Funding to support these various group experimentation efforts should be included in technology project bids.

L. {5 Pages} Cost. Cost proposals shall provide a detailed cost breakdown of all direct costs, including cost by task, with breakdown into accounting categories (labor, material, travel, computer, subcontracting costs, labor and overhead rates, and equipment), for the entire contract and for each **calendar year, divided into quarters**. Where the effort consists of multiple portions that could reasonably be partitioned for purposes of funding, these should be identified as contract options with separate cost estimates for each.

Deleted: 2

Deleted: Government fiscal year (October 1 – September 30)

M. Contractors requiring the purchase of information technology (IT) resources as Government Furnished Property (GFP) **MUST** attach to the submitted proposals the following information:

1. A letter on Corporate letterhead signed by a senior corporate official and addressed to Dr. Tom Wagner, DARPA/IPTO, stating that you either can not or will not provide the information technology (IT) resources necessary to conduct the said research.

2. An explanation of the method of competitive acquisition or a sole source justification, as appropriate, for each IT resource item.

3. If the resource is leased, a lease purchase analysis clearly showing the reason for the lease decision.

4. The cost for each IT resource item.

<u>IMPORTANT NOTE</u>: IF THE OFFEROR DOES NOT COMPLY WITH THE ABOVE STATED REQUIREMENTS, THE PROPOSAL WILL BE REJECTED.

Awards made under this BAA may be subject to the provisions of the Federal Acquisition Regulation (FAR) Subpart 9.5, Organizational Conflict of Interest. All offerors and proposed subcontractors must affirmatively state whether they are supporting any DARPA technical office(s) through an active contract or subcontract. All affirmations must state which office(s) the offeror supports, and identify the prime contract number. Affirmations should be furnished at the time of proposal submission. All facts relevant to the existence or potential existence of organizational conflicts of interest, as that term is defined in FAR 2.101, must be disclosed in Section II, I. of the proposal, organized by task and year. This disclosure shall include a description of the action the Contractor has taken, or proposes to take, to avoid, neutralize, or mitigate such conflict.

Section III. Additional Information

A bibliography of relevant technical papers and research notes (published and unpublished) that document the technical ideas, upon which the proposal is based, may be included in the proposal submission. Provide one set for the original full proposal and one set for each of the **2** full proposal hard copies. Please note: The materials provided in this section, and submitted with the proposal, will be considered for the reviewer's convenience only and not considered as part of the proposal for evaluation purposes.

Deleted: 3

EVALUATION AND FUNDING PROCESSES

Proposals will not be evaluated against each other, since they are not submitted in accordance with a common work statement. DARPA's intent is to review proposals as soon as possible after they arrive; however, proposals may be reviewed periodically for administrative reasons. For evaluation purposes, a proposal is the document described in PROPOSAL FORMAT Section I and Section II (see below). Other supporting or background materials submitted with the proposal will be considered for the reviewer's convenience only and not considered as part of the proposal.

Evaluation of proposals will be accomplished through a scientific review of each proposal using the following criteria, which are listed in descending order of relative importance:

(1) Overall Scientific and Technical Merit: The overall scientific and technical merit must be clearly identifiable and compelling. The technical concepts should be clearly defined and developed. The technical approach must be sufficiently detailed to support the proposed concepts and technical claims. Evaluation will also consider the effectiveness of the system integration and management plan.

(2) Innovative Technical Solution to the Problem: Offerors should apply new and/or existing technology in an innovative way that supports the objectives of the proposed effort. The proposed concepts and systems should show breadth of innovation across all the dimensions of the proposed solution. Offerors must also specify quantitative experimental methods and metrics for measuring progress of the effort.

(3) Potential Contribution and Relevance to DARPA/IPTO Mission: The offeror must clearly address how the proposed effort will meet the goals of the undertaking and how the proposed effort contributes to significant advances to DARPA/IPTO.

(4) Offeror's Capabilities and Related Experience: The qualifications, capabilities, and demonstrated achievements of the proposed principals and other key personnel for the primary and subcontractor organizations must be clearly shown.

(5) Plans and Capability to Accomplish Technology Transition: The offeror should provide a clear strategy and plan for transition to military forces (and commercial sector, where applicable). Offerors should consider involving potential military transition partners, as appropriate, in any proposed experiments, tests and demonstrations. Offerors should also provide a plan for transition of appropriate technology components and information to the user community.

(6) Cost Realism: The overall estimated costs should be clearly justified and appropriate for the technical complexity of the effort. Evaluation will consider the value of the research to the government and the extent to which the proposed management plan will effectively allocate resources to achieve the capabilities proposed.

The Government reserves the right to select all, some, or none of the proposals received in response to this solicitation and to make awards without discussions with offerors; however, the Government reserves the right to conduct discussions if the Source Selection Authority later determines them to be necessary. Proposals identified for funding may result in a contract, grant, cooperative agreement, or other transaction depending upon the nature of the work proposed, the required degree of interaction between parties, and other factors. If warranted, portions of resulting awards may be segregated into pre-priced options.

The administrative addresses for this BAA are:

Fax: 703-741-7804 Addressed to: DARPA/IPTO, BAA 04-29
Electronic Mail: baa04-29@darpa.mil

Formatted: Bullets and Numbering

Deleted: <#>Overall Scientific and Technical Merit: The overall scientific and technical merit must be clearly identifiable and compelling. The technical concept should be clearly defined, developed and defensibly innovative. Emphasis should be placed on the technical excellence of the development and experimentation approach. ¶
¶
(2) Innovative Technical Solution to the Problem: Proposed efforts should apply new or existing technology in an innovative way such as is advantageous to the objectives. The plan on how offeror intends to get developed technology artifacts and information to the user community should be considered. The offeror shall specify quantitative experimental methods and metrics by which the proposed technical effort's progress shall be measured.¶
¶
<#>Potential Contribution and Relevance to DARPA/IPTO Mission: The offeror must clearly address how the proposed effort will meet the goals of the undertaking and how the proposed effort contributes to significant advances to the DARPA/IPTO mission of preventing strategic surprise. ¶
¶
<#>Offeror's Capabilities and Related Experience: The qualifications, capabilities, and demonstrated achievements of the proposed principals and other key personnel for the primary and subcontractor organizations must be clearly shown.¶
¶
(5) Plans and Capability to Accomplish Technology Transition: The offeror should provide a clear explanation of how the technologies to be developed will be transitioned to capabilities for military forces. Technology transition should be a major consideration in the design of experiments, particularly considering the potential for involving potential transition organizations in the experimentation process.¶
¶
(6) Cost Realism: The overall estimated cost to accomplish the effort should be clearly shown as well as the substantiation of the costs for the ... [2]

Deleted: 03-xx

Deleted: 3

Deleted: 44

Deleted: 03-xx

Deleted: 3

Deleted: 44

Electronic File Retrieval: http://www.darpa.mil/ipto/Solicitations/solicitations.htm
Mail to: DARPA/IPTO
ATTN: BAA 04-29
3701 N. Fairfax Drive
Arlington, VA 22203-1714

Deleted: 03-xx

Deleted: 3

Deleted: 44

Formatted: Bullets and Numbering

PROGRAM OBJECTIVES AND DESCRIPTION. The Defense Advanced Research Projects Agency (DARPA) is soliciting proposals for DARPA's Information Processing Technology Office to perform research, requirements and constraint analysis, architecture concept development and design, architectural modeling, in-context evaluations, and concept evaluations to support the initial phase of the Architectures for Cognitive Information Processing (ACIP) program. It is the intent of the DARPA IPTO office to develop cognitive information processes that will bring enabling embedded intelligence capabilities to aid the warfighter, as well as DoD supporting functions and activities – enabling machines that think to aid human performance. Current intelligent processing implementations depend on the use of existing COTS computing architectures that were developed and are best suited for numeric processing applications. To enable the performance of cognitive capabilities in real-time, dynamic, data-intensive, embedded environments and scenarios an underlying processing infrastructure optimized to perform the required cognitive processing is essential. The Architectures for Cognitive Information Processing (ACIP) program seeks to address these deficiencies by developing processing architectures and structures that are uniquely optimized for cognitive computing. The overarching goals of the ACIP program are to develop architectures, processing approaches, and supporting development tools and environments to enable the efficient implementation of embedded realtime cognitive processing and the application of cognitive processing to dynamic, real-world, embedded utilization. This announcement addresses the first phase of the ACIP program. The first phase will address the definition of cognitive computing components requirements specification and runtime requirements; definition of architecture concepts, models and evaluations; development of a concept device specification and technology roadmap; development of composable run-time concepts, and the definition of a living framework approach.

The intent of the ACIP program is to drive the development of a new class of cognitive information computing architectures, data structures, development frameworks, and implementations that efficiently address and instantiate cognitive computing for information processing systems and real-time DoD missions. ACIP will incorporate biological, cognitive algorithm, and DOD mission challenge clues as inputs to establish the concepts of the effort. ACIP will address specific topic areas such as cognitive architectures, alternate representations, composable runtime software, active processing and memory retrieval hardware, and living frameworks to create cognitive information processing solutions. These solutions will be influenced and incorporate concurrent IPTO initiatives in the areas of functional demonstrations and algorithm developments and MTO initiatives addressing physical interconnect and packaging advances. An overall goal of the ACIP Focal Challenge is to provide the computing infrastructure and realtime implementations to enable the IPTO overall goal of "systems that know what they are doing".

The ACIP BAA Focal Challenge will place special emphasis on real and efficient cognitive physical implementations, not just functionality, by developing and demonstrating cognitive information computing system architectures, cognitive computing frameworks, and implementation development environments within DoD application contexts. ACIP will close the cognitive system engineering design loop between algorithms and physical computing structures and lay the foundations for cognitive innovation. Current intelligent processing implementations depend on the use of existing COTS computing architectures that are best suited for numeric processing applications. Today's knowledge representations, abstraction (processing objects), architectures, and implementations are adhoc, awkward and inefficient. Transformation from today's cognitive techniques running on conventional computers is required to develop innovative DoD cognitive computing solutions. To realize the impact and promise of cognitive information processing approaches, computing architectures and development frameworks attuned to cognitive processing fundamentals need to be established that will implement uniquely cognitive structures efficiently. Cognitive computing systems will require: decoupling of languages from underlying structures, composable runtime systems, higher level goal/motive oriented descriptive languages, agile micro-architectures, adaptive morphware, and multi-dimensional memory structures. Cognitive solutions in areas such as cognitive architectures, composable runtime software, alternative representations, active processing and memory retrieval hardware, and a living framework must be addressed. Without a special emphasis on the total "cognitive information processing" context and structures, cognitive techniques and implementations will always be limited by the use of COTS computing architectures that are inefficient for cognitive processing. ACIP will develop revolutionary and efficient cognitive computing architectures and fundamental computing infrastructures including the abstraction representation/storage/retrieval necessary to efficiently implement real-time DoD cognitive approaches and systems. ACIP will create the computing capabilities to meet the goal of computing systems that adapt to emerging threats.

The intent of the ACIP program is to establish cognitive computing capabilities that significantly advance the state of the art and enable efficient computing at all levels of cognitive processing – cognitive threads, cognitive modules, and cognitive systems, and provide the underlying cognitive computing infrastructure and architectures to support efficient cognitive implementations. These developments will be evaluated in terms of complexity, cost, and platform constraints. An important element of the ACIP program will be a Cognitive Information Framework Forum (CIFF) that will be established to promote and pursue common cognitive computing development environments, tools, and evaluation methods across multiple ACIP efforts and provide an enduring basis for wide community use and application.

In order to focus and establish context for the ACIP program, ACIP will pursue processing requirements, realtime constraints, and innovative architectural concepts incorporating concurrent IPTO cognitive processing activities and within in-context DoD mission areas. Such representative areas could include, but are not limited to: resource management, unmanned combat platforms, intelligent analyst assistant, and unattended distributed sensors systems. Such in-context mission areas could provide the context or challenge space relevant to the development of ACIP. Equivalent alternate in-context mission areas will also be considered.

Resource management could address aided and self-management of system computing resources in terms of system introspection. This would include aided and self-aware system computing resource management and optimization and include robustness and validation and verification of system configurations in a dynamic mission environment. Unmanned combat platform missions could address the dynamic use of system resources for mission performance and optimization. This could include dynamic system resource decisions, allocation, and optimization across mission requirements and performance options as well as the aided and self-aware performance of mission requirements within a dynamically changing mission environment. Intelligence analyst assistant development could address the cognitive architectures and computing requirements necessary to perform aided and self-aware analysis such as signal and image analysis. The cognitive organization, coordination, and utilization of diverse and disparate information sources could be addressed. Automated analysis activities would be enabled. Unattended ground sensor activities could include the aided or self-aware dynamic utilization and optimization of varied sensor and computational resources for adaptive sensor fusion, intelligent and optimized interpretation computation based on system resources and conditions, and reactive and proactive exploitation of conditions and system resources. Overall this area could address the computing architectures, data structures and organization, and implementation frameworks to support adaptive and self-aware cognitive interactive processing utilizing an assemblage of sensors, computing, and communications resources across dynamic mission conditions for optimized mission performance. These four application areas are examples that could define constrained challenge spaces, the identification of key derived requirements and the basis for the developmental research testbeds.   The ability to support crisis constrained runtime "cognitive " responses is vital for DoD systems. The goal is to demonstrate for these specific examples "systems that know what they are doing." Successful pursuit, implementation, and integration of ACIP technologies, components, and architectures into a working overall system is paramount.

It is essential that the technologies, components, architectures, and frameworks developed in the course of this research be general enough to be viable across a broad range of applications (portability across cognitive applications) - the goal (as is the goal of the entire BAA) is to create powerful and reusable cognitive computing architectures, technologies, and techniques rather than simply to create a limited implementation that serves only as a single point demonstration

TEST AND EVALUATION. Performers will test and evaluate their technologies using their own facilities and report results at PI meetings. In addition, performers will provide software distributions and will document all test and evaluation choices and procedures (hardware, software environment, scenario, etc.) with enough clarity for a third party to repeat the evaluations. Regarding test and evaluation, an Independent Evaluation Team (IET) will collaborate with performers to foster out-of-the-box thinking and sharing of results among performers and the larger research community.

Within each effort, the performer must quantify the capability demonstrated and the capability to be realized through the cognitive processing approaches and capabilities being developed. Specific metrics and goals relevant to DoD missions and the cognitive requirements, constraints, and development goals being pursued must be established. Advances in cognitive computing capabilities must be quantified against the established metrics and goals.

The ACIP program will provide all contractors with selected kernels that will compose an evaluation and development set for cognitive computing activities. This will also enable a common evaluation process and analysis/evaluation for the ACIP program and support a common library of kernels and metrics for use by ACIP participants. All ACIP contractors will be expected to work collaboratively with these separately funded and neutral ACIP efforts.

PROGRAM SCOPE DARPA. Proposed research should investigate innovative approaches and techniques that lead to or enable revolutionary advances in the state-of-the-art. Proposals are not limited to the specific strategies listed above, and alternative visions will be considered. However, proposals should be for research that substantially contributes towards the goals stated. Specifically excluded is research that primarily results in minor evolutionary improvement to the existing state of practice or focuses on special-purpose systems or narrow applications.

The proposed ACIP program is intended to be broken into three phases for an anticipated total of a 108 month total performance period. Phase I, addressed in this BAA, will be a 33 month effort consisting of the development of cognitive computing components requirements, specifications and runtime requirements; architectural concepts, models, and evaluations; concept device specification and technology roadmap development, establishing composable runtime concepts; and developing a living framework draft. Each proposed cognitive architectural development effort will include the investigation and association of efforts with cognitive learning reasoning and knowledge modules, development of cognitive processing approaches within cognitive DoD applications, development of cognitive architectures and processing structures optimized to address identified cognitive module processing requirements, DoD application requirements, and realtime constraints via innovative architectural concepts. These efforts will develop early architectural concepts and perform in-context evaluations. Deliverables will include cognitive computing requirements specification and runtime requirements; architecture concepts, models, and evaluation; concept device specification with an associated technology roadmap; and the development of composable runtime concepts. Validated multi-level metrics and kernels will be developed for lower implementations at 21 months into the program and for a system level implementation at 33 months. Draft device specification and an implementation technology roadmap shall be delivered at 15 months with the final device specifications and technology roadmap delivered at 33 months. These shall be established within the context of DoD mission applications. Phase I is planned to be followed by a 48 month Phase II implementation, evaluation, and demonstration of the cognitive architectures developed in Phase I of the ACIP program. A Phase III 30 month effort is then planned for the implementation of full scale DoD ACIP system proof of concept efforts. Phase II is contingent on the results and performance of ACIP Phase I and ACIP Phase III is contingent of the results and performance of ACIP Phase II. Throughout Phase I the analysis and development of cognitive computing approaches, architectures, and implementations shall be pursued. In addition during the Phase I effort the evaluation of baseline kernels and metrics that represent cognitive computing within the context of DoD mission areas will be a critical set of activities. These evaluations will be critical in determining the value of an ACIP contractor's effort proceeding into ACIP Phase II. At 33 months the cognitive computing requirements, cognitive modules derived, approach, architectures, and initial implementation approaches developed will be presented and reviewed. These results will be evaluated as potential efforts to proceed with ACIP Phase II. The activities performed in Phase I will establish the viability of the cognitive computing approaches being proposed and developed and the viability and extensibility of the approaches developed. In conjunction with the cognitive architectures definition and architecture development, Living Frameworks will be pursued. At 9 months the architectural concepts for a Living Framework will be developed. At 21 months a Living Framework draft will be presented. At 33 months the baseline concepts necessary for a Living Framework to support cognitive architectures shall be completed and presented. Ongoing Living Framework definitions shall be developed and distributed among the Phase I efforts. Each contracted effort selected for ACIP Phase I shall support and provide inputs to Living Framework development activities and provide inputs to support the Living Framework baseline.

Selection of potential ACIP Phase II performers will be based on Phase I performance and proposed Phase II cognitive computing architecture development, implementation, and demonstration. Potential down selection of Phase I activities may occur at the transition into Phase II. Phase II will provide the actual development, implementation, and initial demonstration of the long term innovative cognitive computing approaches, structures, architectures, and supporting development frameworks developed in ACIP Phase I within in-context DoD mission areas. Phase II is anticipated to be a 48 month effort. Phase III will depend on the success of ACIP Phase II and will be composed of a 30 month full scale implementation of an ACIP system proof of concept. As in Phase I, during ACIP Phase II and II there will be a separate Living Framework activity to support the utilization of the cognitive processing being developed. Milestones for ACIP Phase II and II will be specifically developed based on the results of ACIP Phase I. Performers shall work closely and continuously throughout the ACIP program with the Living Framework development performers working cooperatively to provide the most flexible, supportive, and viable framework across the DoD mission area examples.

Concurrent with the ACIP Phase I cognitive architecture development efforts, and as mentioned above, this BAA also solicits proposals for the support and development of a Living Framework or Cognitive Information Framework Forum (CIFF). The CIFF will pursue and develop common cognitive information living frameworks, interfaces, functionality, adaptation, and modularity across cognitive development activities and provide a common cognitive computing architecture framework for real-time DoD systems. All ACIP cognitive development activities will be participants in the CIFF and work with the CIFF contractor(s) to support the development of the common cognitive computing environment. The proposer for CIFF activities would provide overall forum leadership and work to compose and develop cognitive frameworks, interfaces, tools, and elements across the ACIP program. The CIFF contractor would provide cognitive information computing Early Living Framework architectural concepts at 9 months, a Living Framework draft at 21 months, and baseline Living Framework concepts at 33 months. The CIFF would be carried into a Phase II and III ACIP efforts to provide common support to the cognitive architectures being pursued.

Awards for Phase I efforts are expected to be made during the first half of calendar year 2003. Deliverables, milestones, and demonstrations must be included and clearly defined in proposals with links to the Statement of Work. The establishment of detailed lower level milestones, while at the discretion of the proposer, should clearly provide demonstrable results of the research and integration cumulatively achieved by the team at the milestone described. Milestones of specific interest were briefly discussed earlier in this document. It is anticipated that there will be multiple awards for ACIP Phase I.

Proposers should propose a multi-organizational but integrated team comprising a Lead System Integration (LSI) function and a set of Technology Contributors (TC's). The LSI function will have overall project management responsibility, to include chief architect and interface control functions, system integration of concepts from the TC's, and concept validation and evaluation processes. A proposing LSI should be composed of a well balance team of performers that fully cover the topics of interest of this ACIP Phase I BAA. Multi disciplinary teams are highly encouraged. The teams assembled should incorporate the research disciplines, specifically address the cognitive computing approaches, structures, and architectures proposed, and provide the experience and knowledge of processing approaches, structures, and architectures deemed necessary to address ACIP. The LSI, the integration lead and system integrator, is anticipated to provide the DoD application context lead, specifically providing the expertise and in-context knowledge to support cognitive computing development for relevant in-context DoD mission areas, and or provide unique cognitive processing experience and capabilities (such as concurrent work in cognitive processing algorithm or technique definition and development). Technology Contributors themselves may be multi-organizational, and should reflect a broad and deep representation from the technical community with unique and enabling capabilities for major technical sub-areas key to ACIP Phase I success. They should participate in the design and development activities of the Lead System Integrator, recommend technology elements to the Lead System Integrator, and develop technology elements for all iterations of the architecture and technology concepts for all cognitive computing systems envisioned by the proposal. Proposers are encouraged to bid using this team approach. If multiple mission areas are to be bid, separate proposals should be submitted for each mission area. The CIFF area shall be bid separately from the architecture development activity and can be bid as a single entity or as a team depending on the proposer’s determination to provide the best approach.

Collaborative efforts/teaming are strongly encouraged. The program is designed for teams organized around members with ongoing cognitive experience and current cognitive development activities, relevant DoD mission and application area experience and knowledge, and supporting technology efforts. Additional information is provided in the BAA xx-xx Proposer Information Pamphlet referenced below. Cost sharing is not required and is not an evaluation criterion, but is encouraged where there is a reasonable probability of a potential commercial application related to the proposed research and development effort. Although proposals identified for funding under this effort may result in a contract, grant, cooperative agreement, or other transaction depending upon the nature of the work proposed, the required degree of interaction between parties, and other factors, the Government anticipates awarding only contracts in order to maintain the desired level of control over this research.

This solicitation is for Phase I only. A separate full and open solicitation is currently planned at a later date for a Phase II program. Offerors should not propose a base effort exceeding 36 months. Any such proposal doing so may be disregarded. Options for up to an additional twelve months over the base period will be acceptable. Any offeror may submit a proposal in accordance with the requirements and procedures identified in this BAA. These requirements and procedures include the form and format for proposals. Any classification requirements deemed necessary due to DoD content in any portion of the proposed effort need to be clearly stated and the handling of classified elements of the proposed effort specifically addressed.
PROGRAM OBJECTIVES AND DESCRIPTION. The Defense Advance Research Projects Agency (DARPA) is soliciting proposals for DARPA's Information Processing Technology Office to perform research, development, modeling, design, and testing to support the Self-Regenerative Systems (SRS) program. Network-centric warfare demands robust systems that can respond automatically and dynamically to both accidental and deliberate faults. Adaptation of fault-tolerant computing techniques has made computing and information systems intrusion-tolerant and much more survivable during cyber attacks, but even with these advancements, a system will inevitably exhaust all resources in the face of a sustained attack by a determined cyber adversary. Computing systems and information systems also have a tendency to become more fragile and susceptible to accidental faults and errors over time if manually applied maintenance or refresh routines are not administered regularly. The Self-Regenerative Systems (SRS) program seeks to address these deficiencies by creating a new generation of security and survivability technologies. These "fourth-generation" technologies will bring attributes of human cognition to bear on the problem of reconstituting systems that suffer the accumulated effects of imperfect software, human error, and accidental hardware faults, or the effects of a successful cyber attack. The overarching goals of the SRS program are to implement systems that always provide critical functionality and show a positive trend in reliability, actually exceeding initial operating capability and approaching a theoretical optimal performance level over long time intervals. Desired capabilities include self-optimization, self-diagnosis, and self-healing; it will be important for systems to support self-awareness and reflection in order to achieve these capabilities.

The approach of this program to constructing self-regenerative systems that meet the above needs is to create fourth generation survivability and security mechanisms to complement received first-generation security mechanisms (trusted computing bases, encryption, authentication and access control), second-generation security mechanisms (boundary controllers, intrusion detection systems, public key infrastructure, biometrics) and third-generation security and survivability mechanisms (real-time execution monitors, error detection and damage prevention, error compensation and repair). Among other things, new fourth generation technologies will draw on biological metaphors such as natural diversity and immune systems to achieve robustness and adaptability, the structure of organisms and ecosystems to achieve scalability, and human cognitive attributes (reasoning, learning and introspection) to achieve the capacity to predict, diagnose, heal and improve the ability to provide service.

The vulnerabilities of computing and information systems addressed by this program include mobile/malicious code, denial-of-service attacks, and misuse and malicious insider threats, as well as accidental faults introduced by human error and the problems associated with software aging. The program will build on the advances made in earlier programs addressing the DoD's operational needs for information systems, such as the ability to operate through attacks, maintenance of critical functionality, graceful degradation of non-critical functions in the face of intrusions and attacks when full functionality cannot be maintained, and the ability to dynamically trade off security, performance and functionality as a function of threat.

Fault-tolerant systems deal with accidental faults and errors while intrusion-tolerant systems cope with malicious, intentional faults caused by an intelligent adversary. Combining fault- and intrusion-tolerance technologies produces very robust and survivable systems, but these techniques depend upon resources that may eventually be depleted beyond the point required to maintain critical system functionality. The fourth generation technologies we seek will reconstitute and reconfigure these resources in such a manner that the systems are better protected in the process, reliability is continually improved as vulnerabilities and software bugs are discovered and fixed autonomously, and the ability to provide critical services is never lost.

Assessment and validation of self-regenerative approaches will be carried out to determine their efficacy. The challenge here is that security and survivability requirements have heretofore defied quantification and analytical approaches. Progress made in creating a practical framework for validating intrusion-tolerance techniques will be built upon and extended to validate SRS technologies.

The first phase of this effort is planned to be 18 months long. This is a solicitation for Phase I only. If results are promising, a Phase II follow-on program is a possibility.

Phase I program goals are to create the core technologies needed

- to design and develop systems that provide 100% critical functionality at all times in spite of attacks;
- for a system to learn its own vulnerabilities over time,
- to ameliorate those vulnerabilities,
- to regenerate service after attack, and

ultimately, to improve its survivability over time.

The ultimate goal at the end of a Phase II program would be to achieve sufficient system robustness and regenerative capacity to provide 100 per cent availability of critical functionality and system integrity in the face of sustained malicious attacks and accidental faults.

There will be four major research thrusts in the Phase I technology development of the program. These areas, along with their success criteria, are as follows:

Biologically-inspired diversity. This research thrust area will create a genetically diverse computing fabric in which diversity limits the impact of any given vulnerability. Coarse-grained diversity (e.g., using several different operating systems or server software packages in an architecture) has been used to achieve intrusion tolerance, but that approach was limited by the relatively small number of manually-created interchangeable operating systems, server packages, and similar software components. The technical approach of the SRS program is to achieve fine-grained diversity at the module level to remove common vulnerabilities and to automatically generate numerous diverse software versions. The success criterion for this thrust is the automatic production of 100 functionally-equivalent versions of a software component with no more than 33 having the same deficiency.

"Cognitive immunity" and self-healing. This research thrust area will show automated cyber immune response and system regeneration. The technical approach will include biologically-inspired response strategies, machine learning, and cognitively-inspired proactive automatic contingency planning. The success criterion for this thrust is the accurate diagnosis of at least 10% of the root causes of system problems and automatic effective corrective action for at least half of those diagnoses.

Granular, scalable redundancy. This research thrust area will increase the practicality of redundancy techniques by dramatically reducing the time required to achieve consistency among replicas after an update. This thrust area will attack the consistency problem in two distinct sub-areas—a centralized server setting, and a distributed publish/subscribe setting. Performers who propose to the scalable redundancy thrust area may address either or both sub-areas. Success criteria here include the following: in the centralized server setting, attain a three-fold reduction in latency for achieving consistency of replicated data while tolerating up to five Byzantine failures; in the distributed publish/subscribe setting, attain a fifteen-fold reduction in latency for achieving consistent values of data shared among one hundred to ten thousand participants while using robust epidemic algorithms, where all participants can send and receive events.

Reasoning about the insider threat to preempt insider attacks and detect system overrun. The technical approach will include inferring user goals, enabling anomaly detection, and combining and correlating information from system layers, direct user challenges, etc. The success criterion for this thrust is the thwarting or delaying of at least 10% of insider attacks.

These research areas will explore techniques that span the spectrum from autonomic/reflexive response through and including introspection and learning. These research areas will explore techniques that span the spectrum of human mental function, from autonomic/reflexive response through and including introspection and learning. Proposals should address only one research thrust area. A proposer may submit multiple proposals. The success criteria for the four thrust areas constitute the program's gating evaluation criteria for the possibility of a Phase II follow-on program. They are minimum requirements to gain confidence that self-regenerative systems are feasible. A Phase II program would seek much higher levels of performance. Phase I offerors are strongly encouraged to aim for performance that exceeds these criteria where possible.

It is envisioned that a Phase II program would integrate the more promising techniques into an exemplar system prototype to demonstrate the advantages of implementing these technologies in high value critical applications. The system demonstrated would exhibit the fourth generation capabilities of self-optimization, self-awareness, self-diagnosis, self-healing and reflection.

Offerors must state in their proposals a plan for providing deliverables for installation, training, manuals, etc. required for evaluation by the testing facility, as well as travel costs. Offerors should support the technical feasibility of their concept or idea and discuss the future development of their ideas, validation and transition.

TEST AND EVALUATION. Performers will test and evaluate their technologies using their own facilities and report results at PI meetings. In addition, performers will provide software distributions and will document all test and evaluation choices and procedures (hardware, software environment, scenario, etc.) with enough clarity for a third party to repeat the evaluations. Regarding test and evaluation, an Independent Evaluation Team (IET) will collaborate with performers to foster out-of-the-box thinking and sharing of results among performers and the larger research community. Because progress in the scalable, granular redundancy research thrust area is relative to a baseline that is very sensitive to the testing environment, performers in that area will construct a testbed environment, establish a test procedure, test the best available techniques to determine baseline performance in that testbed, and report their baseline results at the first PI meeting. Testing and evaluation for granular, scalable redundancy techniques developed in Phase I will be conducted on an identical testbed.

PROGRAM SCOPE. Proposed research should investigate innovative approaches and techniques that lead to or enable revolutionary advances in the state-of-the-art. Proposals are not limited to the specific strategies listed above, and alternative visions will be considered. However, proposals should be for research that substantially contributes towards the goals stated. Specifically excluded is research that primarily results in minor evolutionary improvement to the existing state of practice or focuses on special-purpose systems or narrow applications.

This solicitation is for Phase I only. A separate full and open solicitation is possible at a later date for a Phase II program. Offerors should not propose a base effort exceeding 18 months. Any such proposal doing so may be disregarded. Options for up to an additional twelve months over the base period will be acceptable. Any offeror may

submit a proposal in accordance with the requirements and procedures identified in this BAA. These requirements and procedures include the form and format for proposals. Phase I is planned to be unclassified, but Phase II is likely to be a classified program. Offerors who desire to be able to participate in a possible Phase II program are encouraged to be willing and able to obtain appropriate security clearances.
Offerors for the technology development of self-regenerative systems may be foreign firms or may team with foreign firms as long as the firm meets the criteria in this solicitation and the Government is permitted to conduct business with the firm. Offerors for the technology development of self-regenerative systems may also include foreign personnel as part of their proposed resources as long as these personnel qualify technically. It is strongly recommended that researchers in Phase I be willing and able to obtain security clearances in order to be able to continue their work in Phase II.

Page 38: [2] Deleted rgraybill 12/23/2003 3:05:00 PM

Overall Scientific and Technical Merit: The overall scientific and technical merit must be clearly identifiable and compelling. The technical concept should be clearly defined, developed and defensibly innovative. Emphasis should be placed on the technical excellence of the development and experimentation approach.

(2) Innovative Technical Solution to the Problem: Proposed efforts should apply new or existing technology in an innovative way such as is advantageous to the objectives. The plan on how offeror intends to get developed technology artifacts and information to the user community should be considered. The offeror shall specify quantitative experimental methods and metrics by which the proposed technical effort's progress shall be measured.

Potential Contribution and Relevance to DARPA/IPTO Mission: The offeror must clearly address how the proposed effort will meet the goals of the undertaking and how the proposed effort contributes to significant advances to the DARPA/IPTO mission of preventing strategic surprise.

Offeror's Capabilities and Related Experience: The qualifications, capabilities, and demonstrated achievements of the proposed principals and other key personnel for the primary and subcontractor organizations must be clearly shown.

(5) Plans and Capability to Accomplish Technology Transition: The offeror should provide a clear explanation of how the technologies to be developed will be transitioned to capabilities for military forces. Technology transition should be a major consideration in the design of experiments, particularly considering the potential for involving potential transition organizations in the experimentation process.

(6) Cost Realism: The overall estimated cost to accomplish the effort should be clearly shown as well as the substantiation of the costs for the technical complexity described. Evaluation will consider the value to Government of the research and the extent to which the proposed management plan will effectively allocate resources to

achieve the capabilities proposed. Cost is considered a substantial evaluation criterion but is secondary to technical excellence.