DARPA-SN-09-46

Request for Information (RFI)

Ubiquitous High Performance Computing (UHPC)

for

Information Processing Techniques Office (IPTO)

Defense Advanced Research Projects Agency (DARPA)

Table of Contents

DARPA-SN-09-46 Ubiquitous High Performance Computing (UHPC)

Request for Information (RFI)

PURPOSE

All current DoD sensors, platforms, and missions heavily depend on computer systems – from in-field distributed sensors to complex weapons system simulations. Revolutionary deployed computing systems will be essential to support new generations of advanced DoD capabilities. Current evolutionary approaches are inadequate. In order to provide the required revolutionary computer capabilities, DARPA is considering releasing a Broad Agency Announcement (BAA) for a potential Ubiquitous High Performance Computing (UHPC) program.

Specifically, this RFI is looking for feedback from the computing community on the proposed program structure and schedule as outlined below. DARPA is not interested in responses that propose specific computer system solutions. This RFI is being issued for planning purposes and for the purpose of gathering information that could be used to increase the effectiveness of any potential UHPC BAA. All responses to this notice will be treated as information only.

INTRODUCTION

The goal of the envisioned UHPC program is to provide the revolutionary technology needed to meet the steadily increasing demands of DoD applications – from embedded to command center and expandable to high performance computing systems. This may be accomplished by developing highly parallel processing systems with significantly increased power efficiency, enabling ease of programming application development for the user, and resilient execution through all modes of system failure.

Current processing systems are grossly power-inefficient and typically deliver only a small fraction of peak performance. Until recently, advances in Commercial Off-The-Shelf (COTS) systems performances were enabled by increases in clock speed, decreases in supply voltage, and growth in transistor count. These technology trends have reached a performance wall where increasing clock speed results in unacceptably large power increases, and decreasing voltage causes increasing susceptibility to transient and permanent errors. Only increasing transistor count continues to drive performance increases, with value only if we can minimize energy while optimizing our ability to efficiently utilize available concurrency. Further, increasing density has not helped reduce the energy costs of data transport either across a chip, between neighboring chips, or between chips on disparate boards. Current interconnect protocols are beginning to require energy and power budgets that rival or dwarf the cost of doing computation.

To address these concerns the UHPC program will pursue, but will not be limited to: 1) co-development and optimization of ExtremeScale architectures and applications, programming models, and tools – the critical co-design of hardware and software; 2) low-energy architectures and protocols for logic, memory, data access, and data transport; 3) concurrency management and efficient use of massively parallel resources; 4) locality-aware architectures to reduce data movement; 5) self-aware and

learning capabilities to manage real-time performance and resource optimization; 6) coordinated development of resiliency techniques (including detection and correction, fail-in-place self-healing, and learning); 7) security; and 8) physical packaging and thermal issues. UHPC is seeking solutions to break through current performance limitations and increase achievable performance, power efficiency, and system reliability for in field embedded through advanced intelligence analysis applications.

The UHPC program is seeking solutions to develop radically new computer systems that overcome energy efficiency and programmability challenges. These solutions include:

- New *system-wide* technology approaches specifically including hardware and software co-design to minimize energy dissipation per operation and maximize energy efficiency, with a 50GFLOPS per watt goal, without sacrificing scalability to ultra-high performance DoD applications - **efficiency.**
- Develop new technologies and execution models that do not require application programmers to explicitly manage system complexity, in terms of architectural attributes with respect to data locality and concurrency, to achieve their performance and time to solution goals - **programmability.**
- Develop technology that will manage hardware and software concurrency, minimizing overhead for thousand- to billion-way parallelism for the system-level programmer.
- Develop a system-wide approach to achieve reliability and security through fault management techniques enabling an application to execute correctly through both failures and attacks.

The UHPC Program is seeking solutions that will explore the technologies and architectures required to enable the development of revolutionary computing architectures and systems and overcome “business as usual” advances. This can only be achieved via dedicated investment, hardware-software co-design, integrated design techniques, and continuous innovation. Since a high degree of innovation will be a critical element throughout the UHPC program and heavily weighted in all program evaluations, it should be an important consideration in team formation. The goals of the program are to explore, develop, and prototype a UHPC system, applicable to all DoD computer systems. It is assumed that architectural designs developed under this program will be applicable to systems ranging from embedded terascale systems up through at least single cabinet petascale configurations.

TERMINOLOGY

The system components and definitions terminology is introduced as convenience for the future descriptions found within this document.

General Definitions

UHPC System Design: A complete integrated system design that includes hardware, software, execution model, operating system and prototype complier.

ExtremeScale System: A highly productive system that is a thousand times more capable than a current comparable system, with the same power and physical footprint.

Highly Efficient Systems: Systems that significantly advance energy efficiency, effective use of concurrency, and resiliency, relative to the currently projected technology path that computer vendors are expected to follow.

High Performance Systems: Systems that significantly improve the performance characteristics for a broad spectrum of applications, relative to the currently projected technology path that computer vendors are expected to follow.

High Programmability: A system characteristic that does not require programmers to manage the complexity of the system to achieve their performance or "time to solution" goals.

Execution Model: A paradigm for organizing and carrying out computation across all levels of the computer system. It provides the conceptual scaffolding for deriving system elements in the context of and consistent with all of the others. It is a coherent abstract schema that permits co-design and operation of such multiple layers.

Non-UHPC Technology: Technology that is not funded by or developed under the UHPC program.

pJ: 1/(1,000,000,000,000) (pico) joules

GFLOPS: 1,000,000,000 (giga) floating point operations per second

TFLOPS: 1,000,000,000,000 (tera) floating point operations per second

PFLOPS: 1,000,000,000,000,000 (peta) floating point operations per second

TOPS: 1,000,000,000,000 (tera) operations per second

B: 1 byte

GB: 1,000,000,000 (giga) bytes

TB: 1,000,000,000,000 (tera) bytes

System Components and Definitions *(Figure 2 illustrates the relationship between these system components).*

Processor Module: includes processing resources and internal memory.

Processor Node: includes one or more processor modules and memory resources in a micro/multi chip module (MCM) style configuration.

System Node: includes one or more processor nodes and memory, most likely on a printed circuit (pc) board.

Interconnection Network: a high performance fabric that connects processing modules.

I/O System: provides the high performance subsystem capable of streaming input or output data of various types

Storage System: designed to retain data for archival or scratch space.

BACKGROUND

The architectural advances that are required to build ExtremeScale computers (see Terminology section) have many significant hurdles to overcome. The technological advances that are required to build these systems was investigated and identified in the DARPA ExtremeScale Study.[1] To achieve the goal of building these ExtremeScale computer systems the challenges of power, concurrency, memory/data density access and placement, and resiliency must be concurrently addressed.

1. The Energy and Power Challenge is the most pervasive of the four. A key observation is that it will be easier to solve the power problem associated with base computation than to reduce the problem of transporting data from one site to another - on the same chip, between closely coupled chips in a common package, between different racks on opposite sides of a large machine room, or on storing and accessing data in the aggregate memory hierarchy.
2. The Memory and Storage Challenge concerns the lack of currently available technology to retain data at high enough capacities (and access it at high enough rates) to support the desired application suites at the desired computational rate and still fit within an acceptable power envelope. This information storage challenge lies in both main memory (DRAM today) and in secondary storage (rotating disks today).
3. The Concurrency and Locality Challenge likewise grows out of the flattening of silicon clock rates and the end of increasing single thread performance, which has left explicit, largely programmer visible, parallelism as the only mechanism in silicon to increase overall system performance. ExtremeScale systems may have to support upwards of a billion separate threads.
4. A Resiliency Challenge that deals with the ability of a system to continue operation in the presence of either faults or performance fluctuations. This concern grew out of not only the explosive growth in component count for the larger classes of systems, but also out of the need to use advanced technology, at lower voltage levels, where individual devices and circuits become more and more sensitive to local operating environments, and new classes of aging effects become significant.

These challenges cannot be pursed independently at the component level such as processor, memory and network switches; they must be addressed as an integrated solution. Co-design of the system hardware and software that is driven by processing requirements for selected application domains is essential. Solving individual challenges will not result in viable system solutions for the DoD. The driving force behind all of these challenges is the need to minimize the metric, pJ/*op*, where "*op*" is any operation that must be performed to complete the execution of an application. This is a daunting effort, but one well worthwhile to enable the progress and advancement of

[1] ExaScale Computing Studies: see http://users.ece.gatech.edu/~mrichard/ExascaleComputingStudyReports/ECS_reports.htm

future computer systems. A high degree of innovation throughout the life of this program is essential to success.

One of the fundamental problems with current High Performance Computing (HPC) systems is that they are based on sequential models of computation that cannot utilize parallelism. A new model of computation or an execution model is needed that enables the programmer to perceive the system as a unified and naturally parallel computer system, not as a collection of microprocessors and an interconnection network. The execution model provides the conceptual scaffolding for deriving system elements in the context of and consistent with all of the others. Ideally, the execution model implements a decision chain where each layer contributes to the optimum determination of when, where, and how data placement, data movement, and operation of a computation is performed. Current execution models do not emphasize nor manage the specific characteristics critical to a system. This leads to inefficient use of system resources and premature saturation of system efficiency, as measured in terms of metrics such as GFLOPS per watt (GFLOPS/W).

An execution model has the following characteristics:

- provides the governing principles for system design, operation, management, and application;
- impacts the design across all layers of the system architecture stack;
- provides a conceptual framework for the co-design of the layers of the system architecture;
- supports the notion of the operation "decision chain," i.e.,
  - when, where, and why each operation is performed and
  - every layer contributes to each decision of actions; and
- permits reasoning and design decisions in addressing critical efficiency factors.

An execution model is *not*:

- a programming language although it may strongly influence the underlying programming model semantics of which the language is a representation;
- a computer architecture although it establishes the needs for low-level mechanisms that the architectures must support and provides the governing principles that guide the structures and actions of computer architecture in the performance of a computation; nor is it
- a virtual machine isolating the abstractions above it from the implementation details below because it crosscuts all layers from programming language to architecture influencing all aspects of the operation of all system layers in concert.

To realize the potential performance of a system for a specific application, all the system resources (both hardware and software) must be effectively utilized.

Optimization and execution decisions should be made based both on the instantaneous system state and global knowledge about the application behavior. Current operating systems have pre-programmed behaviors that are based on estimates of resource

availability. They are, therefore, ill-suited to large-scale computers based on complex multicore processors and result in sub-optimal performance and potential system failure in the face of changing conditions.

In contrast, desirable Operating System (OS) and run-time solutions will to behave like a self-aware system that “learns” to address a particular problem by building self-performance models, responding to user goals, and adapting to changing goals, resources, models, operating conditions, and even to failures with the following characteristics:

- Introspective - it observes itself, reflects on its behavior, and learns;
- Goal-oriented - ideally, the system’s client only specifies the goal, and it is the system’s job to figure out how to get there;
- Adaptive - the system analyzes the observations, computing the delta between the goal and observed state, and takes actions to optimize its behavior;
- Self-healing - the system continues to function through faults and degrades gracefully; and
- Approximate - the system does not expend any more effort than necessary to meet goals

A notional Self Aware OS (SAOS) is shown in Figure 1.

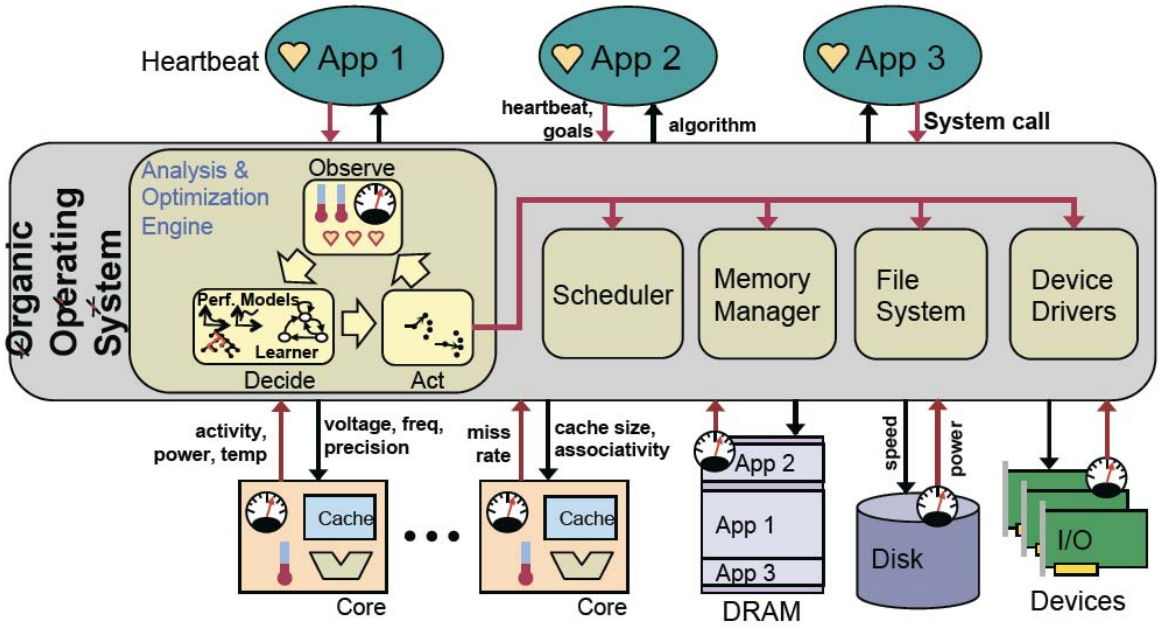

Figure 1: Notional Self Aware Operating System

A separation of concerns is achieved between the application and the SAOS where the application communicates goals and options to the SAOS, and the SAOS uses observation and models of component performance to decide how best to meet

application performance goals given system resources, actual observed performance, and dynamic system constraints.[2]

To reach the goals of the UHPC program, an integrated, multi-level, multi-discipline innovative and complete system approach is essential. Evolutionary advances or the combination of individual technology advances will be insufficient. The UHPC program must provide revolutionary approaches including energy, massive resource concurrency, data location and movement, resiliency, and self-aware operation and the ability to program and effectively utilize system resources as one, integrated system design. To reach the energy goals alone will require reducing energy per operation from thousands of pJ/Op (representative of current processors), to tens of pJ/Op. As an example, achieving 50 GFLOPS/watt is equivalent to expending only 20 pJ per floating point operation – a budget that must encompass far more than just the flop: leakage losses, operand accesses, and operand transport, along with instruction issue and concurrency control. This will require energy-optimized solutions from the basic functional elements through subsystems and systems. Approaches for the efficient use of massive parallel resources must be created, both for the programmer and for actual run-time application performance. This requires the system to be easily programmable by the application developer. The system software stack must have a global perspective of an application execution. In addition, the system itself must be cognizant of resource availability and performance and efficiently use resources at all system levels to optimize performance. Performance in this case includes the ability to both optimally perform the application as well as to perform through failures and dynamic mission variations. New approaches to the system execution model must be developed that recognize and utilize system resources and capabilities. The execution model and operating system developed must incorporate the ability to optimize the use of energy, concurrency, data locality, and resiliency to enable the optimized performance of an application. Without the integrated solution of all of the above, the UHPC program cannot achieve its goals. This will require a high degree of innovation at all system levels and throughout the program by a creative and interwoven team of industry and academic participants. By achieving the goals of the UHPC program, future, computationally demanding DoD missions will become achievable.

PROGRAM DESCRIPTION AND STRUCTURE

The UHPC program is seeking to develop and prototype highly efficient and easily programmable ExtremeScale systems. The anticipated time frame for the availability of ExtremeScale systems is 2017. The transition targets for this program are DoD applications that depend on high-performance, power efficient, physically constrained computing resources. This program targets DoD computationally challenging problems that require computing systems capable of delivering sustained performance approaching 10^{15} operations per second (peta*ops*) on real DoD applications that

[2] For a further description of these concepts, please reference: Organic Computing at http://groups.csail.mit.edu/cag/raw/documents/Agarwal-Harrod-organic-2006.pdf.

consume large amounts of memory, and/or work with very large data sets. There are two challenge problems in the UHPC program; first is a massive streaming sensor data problem resulting in actionable knowledge. The second challenge problem is very large graph-based problems. It is anticipated the challenge problem specifications will be used to aid in the design of each performer's system and would be provided after award. The UHPC systems do not need to provide high performance for legacy applications.

UHPC Goals

The goals of the UHPC system are as follows:

- one PFLOPS, air-cooled, single 19-inch cabinet ExtremeScale system. The power budget for the cabinet is 57 kW, including cooling.
- achieve 50 GFLOPS/W for the High-Performance Linpack (HPL)[3] benchmark.
- The system design should provide high performance for scientific and engineering applications.
- The processor node should be capable of being used within terascale embedded and multiple cabinet systems.
- The system should be a highly programmable system that does not require the application developer to directly manage the complexity of the system to achieve high performance.
- The system must explicitly show a high degree of innovation and software and hardware co-design throughout the life of the program.
- Additional program goals and targets are provided in Table 1 below.

The scope of the UHPC software effort spans the spectrum of operating systems; runtimes for scheduling, memory management, communication, performance monitoring, power management, and resiliency; computational libraries; and compilers. The three key challenges for ExtremeScale software are concurrency, energy efficiency, and resiliency.

A significant problem is managing parallelism and locality. OS-related challenges include parallel scalability, spatial partitioning of OS and application functionality, and direct hardware access for inter-processor communication, asynchronous rather than interrupt-driven messages, and fault isolation. There are additional challenges in runtime systems for scheduling, memory management, communication, performance monitoring, power management, and resiliency, all of which will be built on future ExtremeScale operating systems. The OS should be a self-aware system that "learns" to address user goals and adapting to changing goals, resources, models, operating conditions, and even to failures.

[3] http://www.netlib.org/benchmark/hpl/

The UHPC system must provide ease of programming by the application developer without requiring extensive system expertise. The time required to develop a high performance application code by an application domain expert should not be substantially greater than the time required by a system specific expert programmer. The application developer should be implementing parallel algorithms not developing sequential algorithms that are parallelized by using communication functions. An application development environment should allow an application developer to express all of the known parallelism and data locality characteristics for a particular application code. The expression of the parallelism should be independent of the number of cores or other architectural elements of the processor module/system. An application development environment will extract additional available parallelism. The system itself must then be capable of selecting and configuring the appropriate computational elements. The system configuration must be capable of being dynamically modified during the execution of the code. The UHPC system software must support these capabilities.

The full development of an application development environment is not within the scope of the UHPC program. However, the system hardware and operating system must enable the features specified in the previous paragraph. Demonstrating these capabilities is a UHPC program requirement. The proposed system design must include a description of a programming model for the UHPC system.

To overcome the ExtremeScale challenges and to achieve the aggressive UHPC goals, a new execution model must be developed. Current execution models and system designs won't work at ExtremeScale because of their sequential execution foundations and inherent energy inefficiencies. The challenge of determining a new execution model that drives the development of a UHPC system that achieves the requirements is a significant research effort and the foundation of the UHPC program. In addition, attempting to use current execution models at ExtremeScale will result in prohibitively large software development costs. Recent trends in High Productivity Computing Systems (HPCS) have demonstrated reductions in the human effort involved in developing high-productivity software for current Petascale systems. This does not address the challenges of ExtremeScale architectures such as energy-constrained, many-core parallelism and heterogeneous processors. It is anticipated that the results of the UHPC program will reinvent how computers are utilized and operate. This program does not promote any particular system architecture. Fundamentally, a system consists of processing and memory resources interconnected by a network fabric.

Table 1: Hardware Goals and Targets

System Element	Design Goals	Design Target
Cabinet		
Energy Efficiency	50 GFLOPS/W	
Form Factor	Standard 19” rack such as EIA 310-D standard	
Maximum Cabinet Power	57 kW including cooling	
Cooling	Air cooled	
Address Space		Globally shared
I/O Capability	Support of massive streaming sensor data	
Processor Module		
Energy Efficiency		80 GFLOPS/W
Numeric Format – Floating Point	IEEE754 single and double precision	5 - 10 TFLOPS double precision floating point
Numeric Format – Fixed Point	16, 32, and 64-bit	5- 10 TOPS 64-bit fixed point
Internal Memory		> 32 GB
On-Module Memory Bandwidth		1 B/FLOP
Off-Module Memory Bandwidth		> 1 TB/s
System Node Memory		
Other	Sufficient to support the specified application domains	> 512 GB per Processor Module
Interconnection Network		
Description	High performance computational environment supporting a shared global address space and overall system performance.	
Other	Support high performance interconnects at all levels of the system: inter-module, intra-module on a node, within a cabinet, and between cabinets.	
Storage		
Description	Sufficient to support the specified application domains, including check-pointing, scratch-space, and archival. This system could be comprised of non-volatile memory and/or disk drives.	10 B/FLOP

Notes on Hardware Goals and Targets

1. Meeting the design goals are essential to the success of this program.

2. Targets are values to be pursued in the proposed UHPC program .

3. System resiliency should be addressed in the design. This shall include the concepts of execution through failure, graceful degradation, and repair and redundancy.

4. Multi-cabinet system operation should be supported.

5. The architectures and technologies proposed and developed are intended to eventually be manufactured and sold as a product or as a component within a product.

6. The storage system must be contained within the cabinet.

PROGRAM STRUCTURE

The envisioned UHPC program will have five phases. DARPA currently anticipates that the five phases will include three solicitations - one for Phases 1 through 3 (both Technical Areas), one for Phase 4 and 5 Technical Area 1 and one for Phases 4 and 5 Technical Area 2. At the conclusion of Phase 3, subject to projected UHPC system characteristics, funding, and other program considerations, the other two solicitations may be issued to seek proposals for the fourth and fifth phases of work.

In Phase 1, the focus will be on the initial execution models, conceptual designs and an analytical analysis of the proposed system that will be used to show the offeror's path forward, and initial metrics. Phase 2 will deliver a preliminary execution model, initial system design, exploration and initial development of critical hardware and software technologies working towards an initial simulation of the system, and methodologies for metric evaluation. Phase 3 will deliver the execution model, a preliminary system design, preliminary hardware and software technologies demonstrations, a full system simulation, and evaluation of metrics that show progression toward meeting the program goals, and the results of performance models for the two challenge problems and two DoD applications. Phase 4 will complete the proposed design, build and deliver an operational system, develop benchmark applications. Phase 5 will continue to refine the UHPC systems, including completion of the operating system and prototype compiler, and evaluate the systems by running DoD relevant applications.

The UHPC program also has two technical areas that are described in detail below. DARPA anticipates multiple teams for Technical Area 1, but only one team for Technical Area 2. DARPA's decision whether to authorize Phases 2 and 3, as well as the number of teams for each phase, will be based on an overall assessment of Phase 1 performance. The Government reserves the right not to proceed with Phases 2 or 3 if no technically viable program exists or funding is not available.

Technical Area 1 [TA1] **UHPC System Development**: Team(s) for Phases 1 through 3 will be responsible for the development of the preliminary design(s), and also for demonstrating the capabilities of critical technologies selected to show a viable path toward reaching the program goals. During Phases 4 and 5, teams will be responsible

for the development of systems that reach the UHPC program goals. Across all phases, each team will develop increasingly refined execution models including power efficiency, data locality and usage, operating system and self-aware capabilities. Capabilities will be demonstrated using the selected two challenge problems and two DoD applications.

***Technical Area 2 [TA2]* UHPC System Test and Evaluation:** During Phases 1 through 3, this team will aid DARPA with review of all execution models, system designs, and preliminary critical technologies to ensure the TA1 teams have a viable path toward building a UHPC system and that they will be able to attain the final program goals. During Phases 1 through 3, the team will develop metrics, kernel benchmarks, and performance models for four selected applications (approved by DARPA) for evaluating the systems developed in TA1. During Phase 4, the team will develop benchmark codes to be used to evaluate the TA1 systems. During Phase 5, the team will evaluate final systems using the developed benchmarks. The team will also be responsible for the actual test and evaluation of the final system and compare the results against the program goals.

UHPC Program Metrics

The primary metrics that will be used through the life of the program are:

- Energy efficiency: 50 GFLOPS/W for the HPL benchmark
- System Performance: 1 PFLOPS for the HPL benchmark
- Programmability
- Cabinet Power Requirement

It is anticipated that additional primary metrics will be selected for the final two phases of the UHPC program.

The TA2 team will be responsible for determining additional metrics that can be used to evaluate programmability and the energy, concurrency efficiencies of the proposed systems. They will also generate metrics based on the two challenge problems and the two DoD applications codes. Each TA1 team will generate a set of metrics that they will use to evaluate their proposed system. Near the end of Phase 3, a more comprehensive set of metrics will be selected by DARPA. These metrics will be used to specify the goals of the system that will be built in Phase 4 of the UHPC program.

DETAILED TECHNICAL AREA DESCRIPTIONS

Technical Area 1: UHPC System Development

DARPA seeks a single cabinet, air-cooled system design that will deliver 1 PFLOPS for the HPL benchmark and overcome the UHPC energy efficiency and programmability challenges. To overcome the energy challenge, a new system-wide approach must be developed to minimize energy dissipation per operation, with a 50GFLOPS per watt goal for the HPL benchmark, without sacrificing our ability to scale to powerful enough systems that will support ultra-high performance DoD applications. To overcome the programmability challenge, new technologies must be developed that do not require

application programmers to explicitly manage the system complexity, in terms of architectural attributes with respect to data locality and concurrency, of the system to achieve their performance and time to solution goals. Also the system designs must address a means to expose and manage hardware and software concurrency, minimizing overhead for thousand-way to billion-way parallelism for the system-level programmer. A system–wide approach must be developed to achieve reliability and security through fault management techniques enabling an application to execute through failures and attacks. A high degree of innovation at all levels and areas of development of a UHPC system is critical through all phases of the program,.

It is imperative that the TA1 teams maintain a high degree of innovation during their involvement in the UHPC program.

A high-level representation of one possible system is shown in Figure 2.

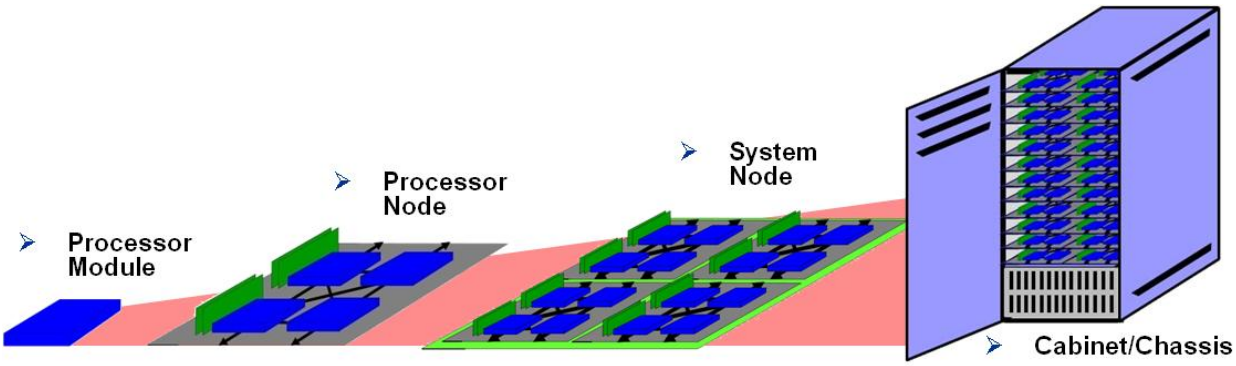

Figure 2: Notional Components of an UHPC system

In this example, a processing module, a processing node, and a system node would be produced and integrated into a full cabinet. It is envisioned that, from the processing module up, each level should have stand-alone capability and be usable in a variety of DoD systems. At each stage, teams will need to show that theirs is an innovative approach and show a viable path, through critical technology demonstrations, system level analysis, and sub-system and system level simulations, to prove that the team can reach the overall system goals. This will incorporate the development of the execution model, operating systems, and self-aware approaches.

Quarterly Program Review (QPR) Meetings will be held each year. The QPR location will be at team's facility or a location of the DARPA PM's choosing. Each team will be expected to provide periodic technical and programmatic updates through QPRs. The QPR will reflect an increasing level of design detail and provide overall program status as well as progress toward meeting the program goals. The QPR should also cover updates to technological exploration, hardware/software co-design, significant design changes, and any issues/concerns. Each team must present the design trade-offs and the impact on the overall goals. Each team will be expected to share in an open forum details of their proposed UHPC design (both hardware and software), including their execution model. Prior to the start of the program, the teams must identify all proprietary information that will be used within or developed for the proposed UHPC system. However, DARPA expects a sufficient level of detail of the UHPC system

design to be openly discussed so that other UHPC researchers can share in the research efforts.

Go/No-Go decisions and possible down-selects of TA1 teams will occur prior to the beginning of Phase 2. These will be based on factors that include degree of innovation, meeting the Go/No-Go criteria, overall design, and the availability of funds.

Phases, Metrics and Deliverables

Phase 1 Overview: Initial execution models, conceptual system design, analytical analysis of proposed system.

Phase 1 will deliver a conceptual system design document, including the initial execution model, for the proposed UHPC system. Innovative technical approaches achieving energy efficiency and programmability must be clearly shown within the deliverables. Each team shall provide an analysis of the proposed system that validates their ability to reach program goals. This includes critical technology assessments.

Phase 1 Deliverable:

The deliverable is a comprehensive report that includes the following documents. These documents will be presented and reviewed at the Conceptual Design Review (CoDR), which will be held one quarter before the end of Phase 1.

- [UHPC Design Document] *A conceptual system design document with sufficient detail to evaluate viability of the proposed system approach/design. This document must include:*
  - an analytical analysis of the proposed system, critical subsystems and components performance;
  - an initial execution model with sufficient technical detail to evaluate capabilities;
  - sufficient technical detail for DARPA and the TA2 team to be able to qualitatively evaluate the proposed system;
  - a justification for the proposed execution model;
  - a description of the full system architecture, including hardware and software, including a proposed programming model; and
  - detailed analytical information that demonstrates the ability of the proposed UHPC system to reach the final goals of the program.
- [non-UHPC Critical Technologies and Subsystems Document] *Details for all non-UHPC developed critical technologies and subsystems that will be utilized in the proposed UHPC system, including acquisition plans, risk assessment and risk reduction plans for each technology. This document must include:*
  - a detailed description of the fundamental critical technologies and subsystems that are proposed to be used in the UHPC system, which will not be developed under this program. An example of a critical technology is a non-volatile memory device that is proposed to be used in the system node.

An example of a critical subsystem is a storage subsystem that might be manufactured by a storage vendor in the UHPC system timeframe.

  - a road map on how the technologies will be developed and acquired; and
  - a discussion of the risk associated with each critical technology and the risk reduction plan.

- [UHPC Test Plan Document] *A plan for the analysis and evaluations to be performed within this phase.*

Phase 1 Go/No-Go Metrics:

There will be three qualitative Go/No-Go evaluations based on the proposed system analysis.

The first qualitative evaluation will be based on the proposed UHPC conceptual design (including the initial execution model) as outlined in the UHPC Design Document. The second qualitative evaluation will be based on the non-UHPC Critical Technologies and Subsystems document. The third qualitative evaluation will be the Conceptual Design Review (CoDR). DARPA will determine if a TA1 team has successfully completed the CoDR.

Phase 2 Overview: Preliminary execution model, system architecture design analysis, initial critical technology simulations, initial system design and simulation demonstrating viable path to program goals.

In Phase 2, each TA1 team will deliver updated versions of the UHPC Design, non-UHPC Critical Technologies and Subsystems, and UHPC Test Plan documents developed in Phase 1. The first document describes the initial system architecture, including the execution model. This document must include sufficient technical detail for DARPA and the TA2 team to be able to evaluate the proposed system. This shall include initial technology exploration and system simulation results. Simulations will incorporate critical technology explorations results, including design elements such as hardware/software interfaces and protocols. The degree of innovation and system impact, in terms of energy efficiency and programmability, must be clearly provided within the deliverables. Each team will provide an analysis and initial simulations/emulations of the proposed system that validates their ability to reach program goals. DARPA will use the TA2 team mentioned above to evaluate the innovativeness, strengths, and limitations of each UHPC TA1 team's overall design.

Phase 2 Deliverable:

The deliverable is a comprehensive report that includes the following documents:

- [UHPC Design Document, v2] *A system design document with sufficient detail to evaluate the viability of the proposed system approach/design. This document must include:*
  - initial simulations of critical technologies, subsystems, and system design;
  - a preliminary execution model of the UHPC system with sufficient technical detail to evaluate capabilities;

  - information on each team's ability to reach the final goals of the program;
  - a complete initial design and simulation/emulation results including hardware/software interfaces and protocols;
  - detailed information on critical sub-systems, component metrics, and the integrated system design; and
  - a description of the methodology of the performed simulations.
- [non-UHPC Critical Technologies and Subsystems Document, v2] *Details for all non-UHPC developed critical technologies and subsystems that will be utilized the proposed UHPC system, including acquisition plans, risk assessment and risk reduction plans for each technology.*
- [UHPC Test Plan Document, v2] *A plan for the analysis, initial simulations, and evaluations to be performed within this phase.*

Phase 2 Go/No-Go Metrics:

There will be two qualitative Go/No-Go evaluations based on proposed system analysis.

The first qualitative evaluation will be based on the system design documents described above, including the preliminary execution model. The second qualitative evaluation will be based on the non-UHPC Critical Technologies and Subsystems document described above. DARPA and the TA2 Team will evaluate these documents. DARPA will determine if a TA1 team has successfully completed its design review.

Phase 3 Overview: Preliminary system design, preliminary hardware and software technologies demonstrations, full system simulation, and evaluation of the UHPC PDR.

In Phase 3, each TA1 team will deliver a preliminary design and the validation of the proposed system that must include sufficient technical detail for DARPA and the TA2 team to evaluate the proposed system. This shall include preliminary technologies, critical subsystems, and system simulation/emulation and critical technologies demonstrations results. Simulations will incorporate critical technology results, including design elements such as hardware/software interfaces and protocols. Simulations/emulations must be capable of running UHPC prototype system software. Innovativeness and system impact, in terms of performance and programmability, must be clearly shown within the deliverables. Each team will provide an analysis and simulations/emulations of the proposed system that validates their ability to reach program goals. During this phase, the TA1 teams will provide a performance model of the UHPC Challenge Problems and the codes for the two additional DoD applications. DARPA, in conjunction with the TA2 team evaluations, will assess the innovativeness, strengths, and limitations of each UHPC TA1 team's overall design.

Phase 3 Deliverable:

The deliverable is a comprehensive report that includes demonstrations and updated versions of the following documents. These documents will be presented and reviewed at the Preliminary Design Review (PDR), which will be held two calendar quarters before the end of Phase 3.

- [UHPC Design Document, v3] *A comprehensive preliminary system design document with sufficient detail to evaluate the viability of the proposed system*

approach/design. This will be based on simulations of the proposed system, including critical subsystems, components, and technologies performance. This document must include

- the execution model of the UHPC system with sufficient technical detail to evaluate capabilities;
- the results of critical components and subsystems demonstrations and the impact on the system design;
- an identification of the high-risk components in the design and a risk reduction plan for each identified component;
- information on each team's ability to reach the final goals of the program;
- a complete design and simulation/emulation results including hardware/software interfaces and protocols; and
- detailed information on critical sub-systems, component metrics, and the integrated system design..

- [non-UHPC Critical Technologies and Subsystems Document, v3] *Details for all non-UHPC developed critical technologies and subsystems that will be utilized the proposed UHPC system, including acquisition plans, risk assessment and risk reduction plans for each technology.*
- [UHPC Demonstrations] *Preliminary hardware and software technologies demonstrations, including prototype system software components*
- [UHPC Test Plan Document, v3] *A plan for simulations, critical component and subsystem demonstrations and performance evaluations to be performed within this phase.*
- [UHPC Benchmark and Metric Document] *Provides results based on simulation for the UHPC benchmarks and Metrics and UHPC application performance models. This document must include the results of the performance modeling of the two UHPC Challenge Problems and the two additional DoD application codes and fully document the performance modeling methodology.*

Phase 3 Go/No-Go Metrics:

There will be four Go/No-Go evaluations based on proposed UHPC system analysis.

The first evaluation will be based on the UHPC preliminary system design documents including the execution model. The second evaluation will be based on the non-UHPC Critical Technologies and Subsystems document. The third is the documented results of the UHPC demonstrations. DARPA and the TA2 Team will evaluate these documents. The fourth evaluation metric will be based on the results of the simulation of the UHPC benchmarks and metrics to include those developed by the TA2 Team.

DARPA will determine if a TA1 team has successfully completed its design review.

The phase lengths shown in Table 2 below for Technical Area 1 are notional. Offerors should propose phase schedules that realistically match their proposed development efforts.

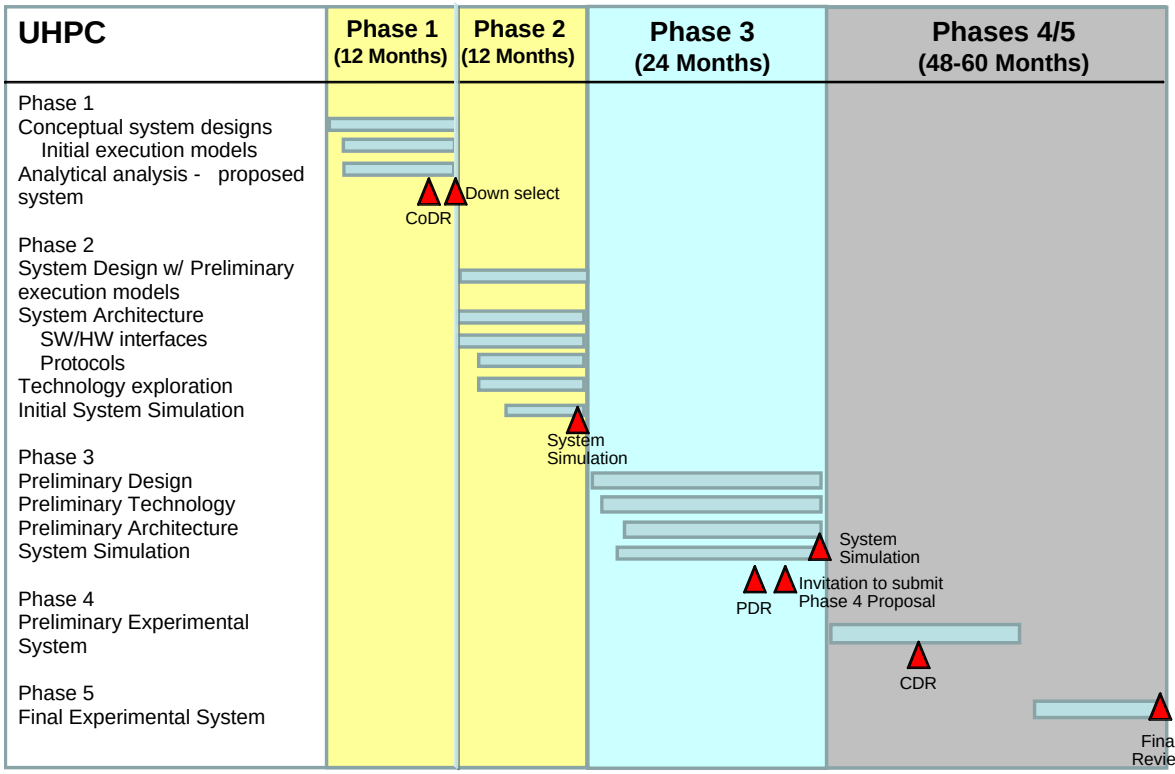

Table 2: UHPC Technical Area 1 Program Elements

Technical Area 2: Ubiquitous High Performance Computing (UHPC) System Test and Evaluation

A single UHPC Test and Evaluation Team will be selected to develop metrics and evaluate the deliverables of the TA1 teams. The TA2 team will review all system designs (including the execution models) and preliminary critical technologies to ensure the TA1 teams show a viable path toward their system and that they will be able to attain the final program goals. The TA2 team will propose, develop, and apply a set of metrics and benchmarks, as approved by DARPA, for evaluating the systems developed in TA1. The TA2 team will develop the specifications and performance models for the Streaming Sensor Data and Informatics Challenge problems and two additional significant DoD applications to be selected by DARPA. The TA2 team will develop performance evaluations of the proposed UHPC systems based on TA1 team simulations and modeling. The TA2 team will be responsible for the actual test and evaluation of the final TA1 systems and comparison of the results against program goals. The TA2 Team is expected to participate throughout the first three phases of this program. Continuation of this TA2 team will depend on such factors as meeting Go/No-Go criteria, and other factors determined by DARPA. Continuation of TA2 is contingent on the decision to proceed with ongoing TA1 activities among other programmatic considerations.

This team should consist of experts in modeling and simulation, and have experience in the development of large-scale applications, computer performance evaluation, and computer architectures. The TA2 team should have expertise in streaming sensor data, informatics applications, and other representative DoD applications.

Quarterly Program Review (QPR) Meetings will be held each year. The QPR location will be at team's facility or a location of DARPA PM's choosing. Each team will be expected to provide periodic technical and programmatic updates through QPRs. The QPRs will reflect an increasing maturity of methodologies, metrics, challenge problem and DoD application specifications as well as reviewing TA1 design activities. The TA2 team will be expected to share details of their activities in an open forum details. DARPA expects a sufficient level of detail in the methodologies and challenge problems and DoD application specifications to enable the evaluation of UHPC system designs.

Phases, Metrics and Deliverables

Phase 1 Overview: Definition of initial metrics, challenge problems and DoD applications specifications, development of test and evaluation schedule, and evaluation of conceptual UHPC designs.

During Phase 1, the TA2 team will select two DoD applications (approved by DARPA) that, along with the two UHPC Challenge problems, will be used to develop specifications and provided to the TA1 teams. The specifications will have a sufficient level of detail for the TA1 teams to use the information to drive the design of their UHPC systems. These specifications will be provided to the TA1 teams six months after the initiation of the TA2 contract. The TA2 will propose a set of metrics, to be approved by DARPA, for evaluating the systems developed in TA1. Proposed metrics shall include energy, concurrency, resiliency, programmability, and other metrics as deemed necessary to fully evaluate proposed TA1 systems. The TA2 Team must provide a written report that examines the TA1 team proposed UHPC conceptual designs, including the initial execution models. This report shall cover in detail whether the execution model and conceptual design document delivered by each TA1 team shows a viable path toward meeting the program goals. It should also review the analytical analysis of the proposed system that is to be supplied. The TA2 team shall review all information pertaining to the sub-system and components to make sure they also align with the TA1 team meeting the final goals of the program.

The TA2 team will develop and provide a methodology and schedule for the complete test and evaluation process for the UHPC systems. The DARPA Program Manager will review and approve the proposed metrics and schedule.

Phase 1 Deliverables:

- [UHPC Challenge Problems and DoD Applications Specification Document] *This document will provide the specifications for the two UHPC challenge problems and the two DoD applications. The initial specifications will be delivered six months after contract initiation. An updated document will be delivered at the end of the Phase.*

- [TA1 Team UPHC Evaluation Document(s)] *An initial evaluation of each TA1 Team proposed conceptual UHPC system, including the execution model. There will be one document for each TA1 Team.*
- [UHPC Metric Document]. *This document will provide information concerning the proposed metrics for evaluation of the characteristics of the UHPC systems.*
- [UHPC Test and Evaluation Plan]*. This document will provide the methodology and schedule for the test and evaluation of UHPC systems through Phases 1 through 3.*

Phase 1 Go/No-Go Metric:

The Phase 1 Go/No-Go Metric for the TA2 Team is the completion of all Phase 1 deliverables.

Phase 2 Overview: Methodology for evaluating metrics, software and hardware technology evaluations, develop performance models and update specifications for the two challenge problems and two DoD applications, initial kernel benchmarks, and evaluation of proposed UHPC designs.

During Phase 2, the TA2 team will develop the initial methodology for evaluating the UHPC metrics (to be approved by DARPA) used for evaluating the systems developed in TA1. The TA2 team will develop the initial performance models and update specifications for the Streaming Sensor Data and Informatics Challenge problems and two additional significant DoD applications selected by DARPA. The team will develop initial kernel benchmarks to be used for performance evaluation of the TA1 team designs. TheTA2 team must provide a written report that examines each TA1 team proposed design, including the preliminary execution model, and initial proposed critical technologies. This report should cover in detail whether the execution model and design documents delivered by each TA1 team shows a viable path toward meeting the goals of the program. The TA2 team will evaluate the proposed UHPC systems based on TA1 team simulations and modeling. The TA2 team shall review all information pertaining to the sub-system and components to make sure they also align with the TA1 team meeting the final goals of the program.

Phase 2 Deliverables:

- [UHPC Challenge Problems and DoD Applications Specification Document v2] *This document will provide updated specifications and initial models for the two UHPC challenge problems and the two DoD applications.*
- [TA1 Team UHPC Challenge Problems and DoD Applications Performance Model Document(s) v1] *This document will provide the methodology for generating the performance model a TA1 team's proposed UHPC system, for the two UHPC challenge problems and the two DoD applications. There will be one document for each funded TA1 Team.*
- [UHPC Kernel Benchmarks v1] *This document will provide initial Kernel Benchmarks. These benchmarks will be provided to the TA1 teams.*

- [UHPC Metric Document v2]. *This document provides updated information concerning the proposed metrics and the methodology for evaluating the metrics of the UHPC systems.*
- [TA1 Team UPHC Evaluation Document v2] *This document provides an evaluation of a TA1 team's proposed design, including the preliminary execution model, and critical technologies.*
- [UHPC Test and Evaluation Plan Document v2]. *This document will provide an update of the methodology and schedule for the test and evaluation of the Phase 2 UHPC systems.*

Phase 2 Go/No-Go Metrics:

The Phase 2 Go/No-Go Metric for the TA2 team is the completion of all Phase 2 deliverables.

Phase 3 Overview: Preliminary design review of TA1 team designs, final performance modeling of the two challenge problems and two DoD applications, and benchmark kernels.

During Phase 3, the TA2 Team must provide a written report that examines the TA1 team preliminary design, including the execution model, and proposed critical technologies and demonstrations. The TA2 team will evaluate the performance of representative application code on the TA1 teams' simulation/emulations. The TA2 team will participate in the TA1 teams' PDRs. This report should cover in detail whether the execution model and design documents delivered by each TA1 team shows a viable path toward meeting the program goals. The TA2 team will perform performance evaluations of the proposed UHPC systems based on TA1 team simulations/emulations and modeling. The TA2 team shall review all information pertaining to the sub-system and components to make sure they also align with the TA1 team meeting the final program goals. The TA2 will finalize and deliver the methodology for the UHPC metrics, to be approved by DARPA, and evaluate the systems developed in TA1. The TA2 team will continue the refinement of the specifications and deliver the final performance models for the Streaming Sensor Data and Informatics Challenge problems and two additional significant DoD applications selected by DARPA. The TA2 team will continue development of and deliver benchmark kernels.

Phase 3 Deliverables:

- [TA1 Team UPHC Evaluation Document v3] *This document provides an evaluation of a TA1 team's preliminary design, including the final execution model, and critical technologies.*
- [UHPC Test and Evaluation Plan Document v3]. *This document will provide an update of the methodology and schedule for the test and evaluation of the Phase 3 UHPC systems*.
- [UHPC Challenge Problems and DoD Applications Specification Document v3] *This document will provide an update of the specifications and final models for the two UHPC challenge problems and the two DoD applications.*

- [TA1 Team UHPC Challenge Problems and DoD Applications Performance Model Document v2] *This document will provide an update of the methodology for generating the performance model of a TA1 team's proposed UHPC system, for the two UHPC challenge problems and the two DoD applications. There will be one document for each funded TA1 Team.*
- [UHPC Kernel Benchmarks v2] *This document will provide Kernel Benchmarks. These benchmarks will be provided as open source to the community.*

Phase 3 Go/No-Go Metrics:

The Phase 3 Go/No-Go Metric for the TA2 team is the completion of all Phase 3 deliverables.

TA2 team selection for Phases 4/5 will be based on an open solicitation.

The phase lengths shown in Table 3 below for Technical Area 2 are notional. Offerors should propose phase schedules that realistically match their proposed evaluation efforts.

Table 3: UHPC Technical Area 2 Program Elements

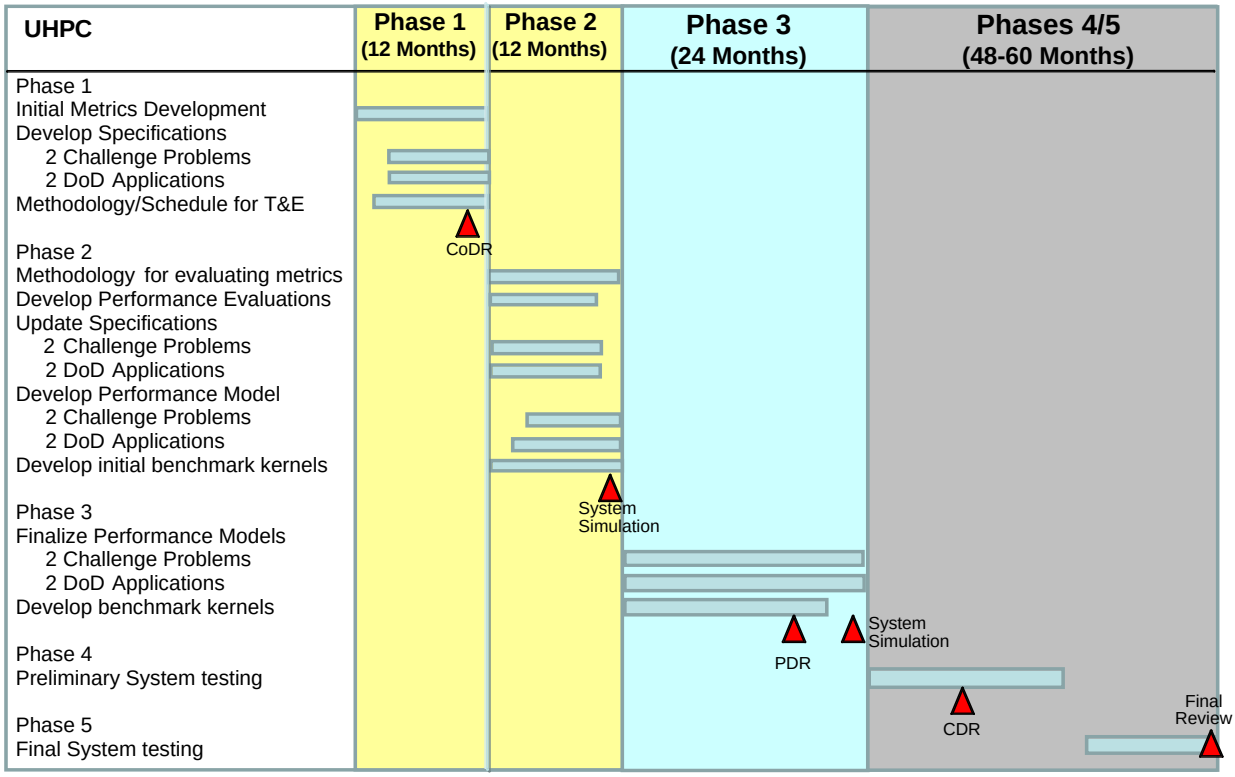

TEAMING & COLLABORATION

DARPA strongly encourages teams that fully address the set of technologies required to overcome all challenge areas discussed above and accomplish UHPC Program goals. These teams must provide demonstrated experience in the appropriate technology

areas. TA1 team expertise should include: energy efficiency, massive resource concurrency, data location and movement, resiliency, self-aware operations, execution models, operating system, compiler design, simulation/emulation development, the ability to program and effectively utilize system resources as one integrated system design, and relevant DOD application experience. TA2 team expertise should include: DoD applications, streaming and informatics applications, component to system level test and evaluation, metrics, and benchmark development and evaluation. Team participants should be drawn from both academic and industrial communities. It is expected that both communities will be involved in all aspects of the UHPC program.

The goal of multi-discipline teaming is to achieve faster progress by creating a critical mass of innovative, relevant expertise. While DARPA expects strong, multidisciplinary teams, each team should have a single identified lead designated as the primary point of contact (POC) with DARPA. DARPA expects each team to submit a single, unified proposal. Subcontractors should not submit separate proposals.

In order for the program to make maximum progress, all performers will be required to share non-proprietary technical information and results with other performers. All proprietary information must be disclosed within each team's proposal.

Team Structure Transition Strategy

DARPA envisions Phases 1 through 3 of the proposed UHPC program under one solicitation and Phases 4 and 5 will be addressed by two separate solicitations. If Phases 4 and 5 are warranted, DARPA anticipates a limited solicitation for the TA1 teams and an open solicitation for the TA2 team. For TA1 teams, Phases 4 and 5 will be based on the preliminary designs that are developed in Phase 3. Changes or improvements to these designs must be justified using an appropriate level of analysis that will include results that show the impact of these changes on the UHPC program goals. DARPA acknowledges that the composition of TA1 teams who submit proposals for Phases 4 and 5 may be different than the teams that developed the preliminary designs in Phases 1, 2 and 3; however, a potential critical evaluation criterion will be the ability of the proposed Phase 4/5 team to complete the proposed design and build the system. Maintaining a high degree of innovation through the life of this program is a requirement.

SUBMISSION INSTRUCTIONS AND FORMAT

Responses must be submitted per the instructions here no later than 1200 ET on 27 July 2009. DARPA appreciates responses from all sources. DARPA will employ an electronic upload submission system for responses to this RFI. To respond to this RFI, interested parties must complete an online RFI Cover Sheet for each response, which will include the submitter's name, organization and contact information. Cover Sheet submissions are made by going to www.csc-ballston.com/rfi/rfiindex.asp?RFIid=09-46. After finalizing the Cover Sheet, a Confirmation screen will appear, along with instructions for uploading your response. Submissions must be in Microsoft Word or Adobe PDF format. **Since respondents may encounter heavy traffic on the web server, they SHOULD NOT wait until the day the submissions are due to fill out a coversheet and upload their response.**

Failure to comply with these submission procedures may result in the response not being read.

Responses should not exceed 5 pages and should be as succinct as possible while at the same time providing actionable insight. Format specifications include 12 point font, single-spaced, single-sided, 8.5 by 11 inches paper, with 1-inch margins in either Microsoft Word or Adobe PDF format. **NO CLASSIFIED OR PROPRIETARY INFORMATION SHOULD BE INCLUDED IN THE RFI RESPONSE.**

1. Cover Page (1 page): As referenced above, this cover page will include the respondent's technical and administrative points of contact (names, organizations, addresses, phones and fax numbers, and email addresses)

2. Technical Content (up to 4 pages): Provide comments on the technical information provided above, specifically on the proposed UHPC program structure and schedule.

ELIGIBILITY

DARPA invites participation from all those engaged in related research activities and appreciates responses from all capable and qualified sources including, but not limited to, universities, university-affiliated research centers, Federally-Funded Research and Development Centers (FFRDC), private or public companies and Government research laboratories.

DISCLAIMERS AND IMPORTANT NOTES

This is an RFI issued solely for information and program planning purposes; this RFI does not constitute a formal solicitation for proposals or abstracts. Your response to this notice will be treated as information only. In accordance with FAR 15.201(e), responses to this notice are not offers and cannot be accepted by the Government to form a binding contract. DARPA will not provide reimbursement for costs incurred in responding to this RFI. Responses to this RFI are not required to propose to subsequent Broad Agency Announcements or research announcements (if any) on this topic. **NO CLASSIFIED OR PROPRIETARY INFORMATION SHOULD BE INCLUDED IN THE RFI RESPONSE.** Respondents are advised that DARPA is under no obligation to acknowledge receipt of the information received or provide feedback to respondents with respect to any information submitted under this RFI. Responses to this RFI do not bind DARPA to any further actions related to this topic including requesting follow-on proposals from vendors responding to this RFI. Submissions may be reviewed by: the Government (DARPA and partners); Federally Funded R&D Centers (such as MIT Lincoln Laboratory); and Systems Engineering and Technical Assistance (SETA) contractors (such as Schafer Corporation, Science and Technology Associates, CACI International, and System Analysis, Inc.).

POINT OF CONTACT

Dr. William Harrod, IPTO Program Manager, DARPA. ANY INQUIRIES ON THIS RFI MUST BE SUBMITTED TO DARPA-SN-09-46@darpa.mil. NO TELEPHONE INQUIRIES WILL BE ACCEPTED.