

Domain-Specific System on Chip (DSSoC)
Updated 10/23/2017

1. Q: For the Phase 0 schedule delivery using an existing SOC preferred. Is it possible to propose analyzing a portion of the SOC that would move forward to Phases 1 or 2 as a new heterogeneous architecture with supporting tools?

A: Continued evolution of an SoC to support the goals of DSSoC is possible.

2. Q: Can domain-specific languages be an alternative to “reverse-engineering” specialization code?

A: The DSSoC program is not seeking new programming languages. If it is a domain-specific framework, however, then that would represent a significant amount of knowledge about a domain that could be repurposed to the program.

3. Q: How does the DSSoC vision relate to OpenCL?

A: OpenCL is a useful tool to develop for heterogeneous processors. It only addresses part of the problem, but could be one potential way to support generic programming practices to target a number of different types of PEs.

4. Q: Is it correct that early acceptable alternatives for the software would be to use legacy code and or to create libraries?

A: DSSoC is looking to advance new hardware concepts of how complex SoCs are built. Code and code development practices should, when possible, come from previous development efforts as well as feed back to those efforts. A benefit of the software development of DSSoC is the ability to improve and optimize complex mathematics that already exist while also providing a new hardware processor that can use them better. The emphasis on existing code and existing development practices is a practical one to make sure that current developers and development efforts can easily adapt to and therefore adopt these new concepts.

5. Q: How does neuromorphic processing fit in this BAA? In what ways?

A: DSSoC does not aim to advance neuromorphic computing, but teams could make use of neuromorphic processing as part of the solution, for instance, as one of the accelerators.

However, the ontology would have to show that this type of specialization has utility to the selected domain.

6. Q: Should the DSSoC chip cater to applications across multiple domains or multiple applications within a given domain?

A: A given DSSoC should be able to efficiently execute multiple applications within a given domain.

7. Q: Phase 0 calls for using an existing SOC within the first 12 months. Yet your talk pointed out that applications development on existing chips often takes much longer. Could you clarify?

A: In DSSoC, application development should be taking place throughout all four phases. Each phase will produce a new instantiation of a SoC (starting with a commercially available one in Phase 0), and applications need to be shown running on each instantiation at the end of each phase. Nevertheless, one of the primary goals of DSSoC is to reduce application development time.

8. Q: How will this community accelerate the use of mechanical mathematics to enable reuse, to speed analysis, and to enable analytic verification?

A: Through the use of open source software and the method by which OSS is published and maintained, community access and the development of new communities interested in these topics can grow. By providing the mechanics of OSS code distribution and maintenance, structure, bug fixes, and overall improvement and verification become part of the development cycle.

9. Q. How will this community share their mathematics?

A: Through the use of open source software and publishing in conferences that are both intellectually interested in the area of mathematics as well as open source conferences to develop communities. Mathematicians in general are very good about sharing their ideas. And open source software is a good way to spread knowledge by encoding it in the software (i.e., code and documentation).

10. Q. Does DSSoC focus on monolithic solutions, or are SIP solutions are acceptable?

A: DSSoC does not prescribe the type of fabrication or packaging solutions for the final hardware deliverable. However, the expectation is the hardware deliverable will be a chip-level, as opposed to a board-level, solution.

11. Q. What are application drivers for runtime configuration? What is the time granularity of interest?

A: The driver for runtime configuration is twofold. First, the DSSoC program is looking to remove a large part of the burden of the current manual labor that goes into managing the complex processors already available that will become more complex with DSSoC. Second and more to the question, DSSoCs should run multiple simultaneous applications where the applications will change based on user or environmental input and conditions. Many military missions require a large number of and highly reactive radio systems that could be done in a single platform with DSSoC. Similar domains and use cases that respond to sensor input need similar resources. Multifunction systems share resources and so static scheduling becomes impossible under these conditions.

The time granularity of interest is going to be domain-driven. The metrics of the program require that, at the end, a new decision at runtime happen at a rate of a most 5 ns. Not all scenarios, applications, or domains will need to change behavior at this rate, but for optimizing the use of resources, the decisions are made at the time scale of the latency between moving data from one PE to another PE.

12. Q: From the presentation it seems that the work that you intend to fund is all software development (nothing physical). Is that correct?

A: DSSoC does require several hardware deliverables, but with DSSoC software and applications running on them. At the end of Phase 0, a commercial device should be selected and delivered. At the end of Phase 1, a hardware emulation should be shown running on discrete hardware components. At the end of Phase 2, a fabricated chip should be delivered. At the end of a Phase 3, a final version of DSSoC should be fabricated and delivered.

13. Q: Can you elaborate on the last objective, of actually building a DSSoC at the end when all previous tasks seem to be coding. Also, given the RF focus, is any interest in III-V Materials?

A: While the focus of DSSoC is on software, there are hardware deliverables as part of the program. Please see Question 12. Solutions making use of III-V materials are welcome, however, similar to Question 6, DSSoC does not aim to advance III-V research.

14. Q: What if a prime that is used to traditional DFARS CPFF contracts wants to have non-traditional subcontractors who intend to cost share a portion of the work? Is this completely at

the discretion of the prime? How do we create compliant Cost Proposals including subcontractor costs in this arrangement? How are the data and patent rights managed?

A: Yes, this is completely at the discretion of the prime. In such a scenario, the non-traditional subcontractor would have to comply with the required FAR/DFARS flow-down clauses (there is no such thing as an OT under a Procurement Contract). Also, in this scenario, the for-profit prime would have to give up (not propose) fixed fee since the cost share from the sub flips the instrument from CPFF to Cost Reimbursement with Cost Share. Regarding the non-traditional subcontractor cost proposal, the prime would need to ensure as much of the cost details that are being requested in the BAA are provided as possible for/from the non-traditional subcontractor. If the non-traditional subcontractor is a small business, and will be receiving under \$7,000,000, there is a certified cost and pricing data waiver in place (see the BAA) that offers some relief (but the primary point that prime needs to keep tabs on is providing enough cost information to ensure that what is being proposed is fully understood by the Government). It is assumed that in such scenario the non-traditional subcontractor would be seeking a firm fixed price subcontract due the fact they would not have an approved cost accounting system (and such an organization would likely be disinterested in dealing with cost accounting obligations anyway), so it would be important that the subcontractor also provide a payment milestone plan (similar to that discussed in the BAA for OTs) that maps to their SOW and cost proposal/build-up.

15. Q: Do applications have to be RF applications?

A: DSSoC has listed software radio as the preferred domain, but proposers may propose their own domain and applications. If the selected domain is software radio or an RF-related domain, then applications should be RF-related. Other domains should lead to the selection of other appropriate applications.

16. Q: Can a team be fully academic? Is industry optional?

A: So long as a proposal is fully responsive to all five program components of DSSoC, the make-up of the team can include any type and mix of partners. No one type of partner is required.

17. Q: Is the right approach to explore on CPUs and FPGAs and port to a DSSoC in later phases?

A: Phases 0 and 1 specifically ask for exploration on commercially available parts. These can be CPUs or FPGAs or other discrete hardware components. However, design and development work, particularly ontology development, of the DSSoC should begin at the start of the program.

18. Q: Is there a role for asynchronous architectures?

A: Solutions making use of asynchronous architectures are welcome, however, similar to Question 6, DSSoC does not aim to advance research in this area.

19. Q: Would in-kind cost sharing be considered as a cost sharing?

A: In-kind cost sharing is allowable when proposing/negotiating an OT as the award instrument (See, for TIAs, Title 32 - National Defense. Subtitle A - Department of Defense. CHAPTER I - OFFICE OF THE SECRETARY OF DEFENSE. SUBCHAPTER C - DoD GRANT AND AGREEMENT REGULATIONS §§ 37.525 through 37.555). The only allowable form of cost sharing under a procurement contract is cash (sharing in the incurred allowable costs/expenditures). See FAR 16.303.

20. Q: Does the ontology methodology need to be shared with the DSSoC community?

A: Yes at a conceptual level. The intent is to create a process to describe a domain and how the domain maps to the makeup of the DSSoC's set of processor elements.

21. Q: What is the purpose of the abstract?

A: The abstract process is designed so that teams can submit their proposal concept for direct feedback from the DSSoC program manager. Abstracts will receive feedback on whether they are encouraged or discouraged from submitting a full proposals.

22. How does DSSoC focus on computing on the edge?

A: In computing, "edge" is still not enough of a well-defined term, but the answer is likely "yes." DSSoC is focused on embedded systems and systems that sit close to the sensor.

23. Q: How does open source create a DoD advantage?

A: Open source software has already proven its effectiveness in commercial industry and the DoD has recognized its relevance and value to the DoD mission. OSS expands the abilities of any one company to focus on their areas of excellence while distributing the management and improvement of their software system with everyone else. Following this lead means that the government does not have to maintain specialized code by itself and move to a collaborative environment to manage the code and thereby allowing the DoD to focus more on the problems of interest on top of the fundamentals. Specialized code that speaks to DoD applications can be kept close-hold and even classified while still benefitting from the force multiplier that participating in open source development provides. And as such, the program is promoting proper engagement with the open source communities to make sure that the projects improve while also understanding that open source communities are willing to work harder and pay more attention to those that actively contribute back.

24. Q: What should the applications be?

A: Please see Question 15.

25. Q: Does a chip from DSSoC need to be open source hardware?

A: No. The program is pushing for open source software to improve development time and capabilities for areas that can be and should be shared across the entire development community. The hardware portion is not required to be open. For the PEs, this is true so long as it fits into the MAC layer that will be defined in phase 0 of the program such that the software written for it can use the PE without specialized access requirements.

26. Q: DSSoC requires proposals to be fully responsive to all five technical areas. If we have expertise in one area, how can we find partners to complete our team and be fully responsive?

A: ERI has been compiling a Teaming Partner List of organizations who are looking to partner on ERI proposals. You can be added to and have access to the list by registering at: <http://www.cvent.com/events/eri-teaming-partner-list/event-summary-b69b9f99098f48669d6e9eea8a382497.aspx>