Energy Efficient Computing: from Devices to Architectures (E2CDA)

Information Webcast

January 21, 2016

Solicitation: NSF 16-526

AGENDA

January 21, 2016

1:00 PM **Welcome and Introduction of Program Officers**

- Sankar Basu, NSF-CISE-CCF
- Dimitris Pavlidis, NSF-ENG-ECCS
- Thomas Theis, SRC-NRI
- Jon Candelaria, SRC-GRC

1:10 PM **Overview of E2CDA Program Solicitation**

1:40 PM **Questions & Answers**

Housekeeping Notes
AFTER THE MEETING

- After the meeting, a copy of the slides will be archived and available. Visit this website for more information: **http://www.nsf.gov/pubs/2016/nsf16526/nsf16526.htm**

- After the meeting, you can submit questions to:

Sankar Basu, sabasu@nsf.gov ; (703) 292-7843

Dimitris Pavlidis, dpavlidi@nsf.gov ; (703) 292-2216

Thomas Theis, thomas.theis@src.org ; (914)945-2244

Jonathan Candelaria, jon.candelaria@src.org ; (919) 941-9482

Purpose of Webcast

The purpose of this webcast is:

- **To inform the community about the goals of the E2CDA Program Solicitation and,**

- **To respond to questions from potential applicants.**

Directorate for Computer & Information Science & Engineering (CISE)

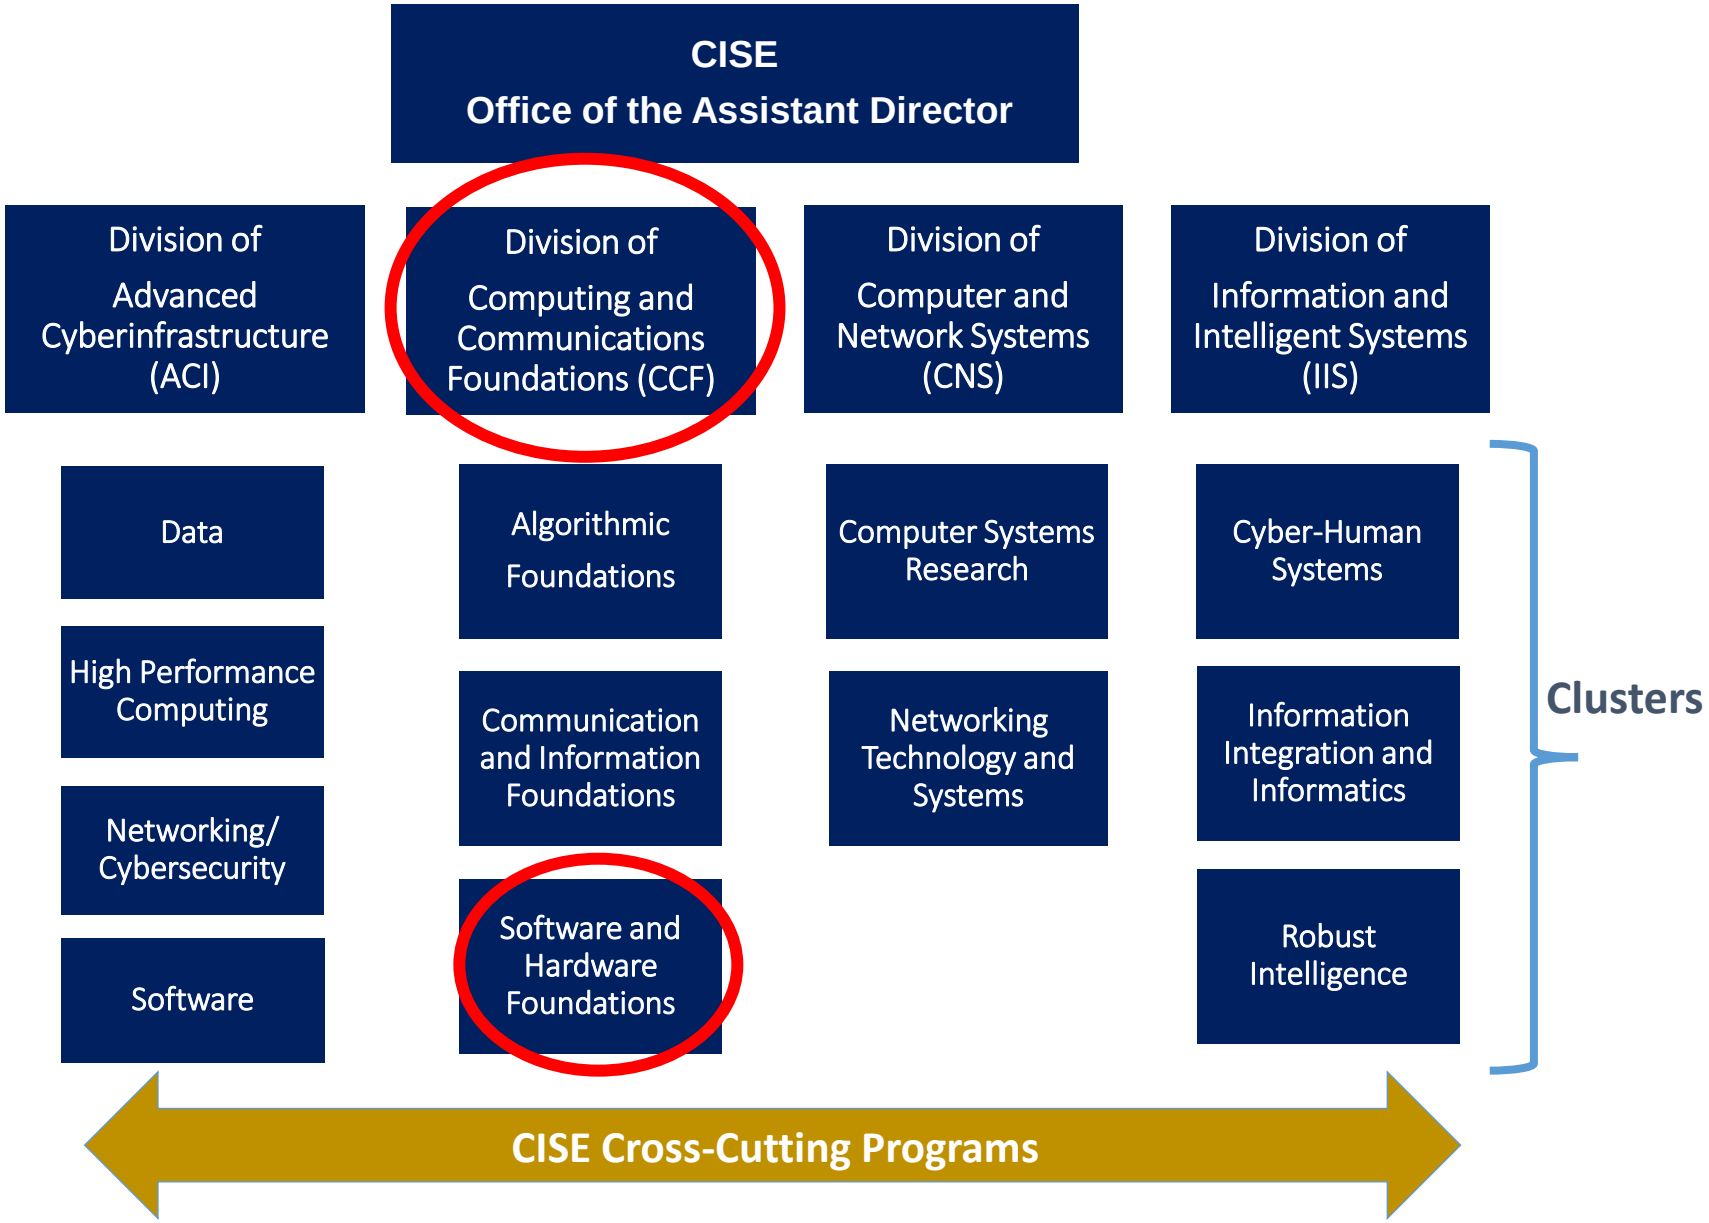

Directorate for Engineering (ENG)

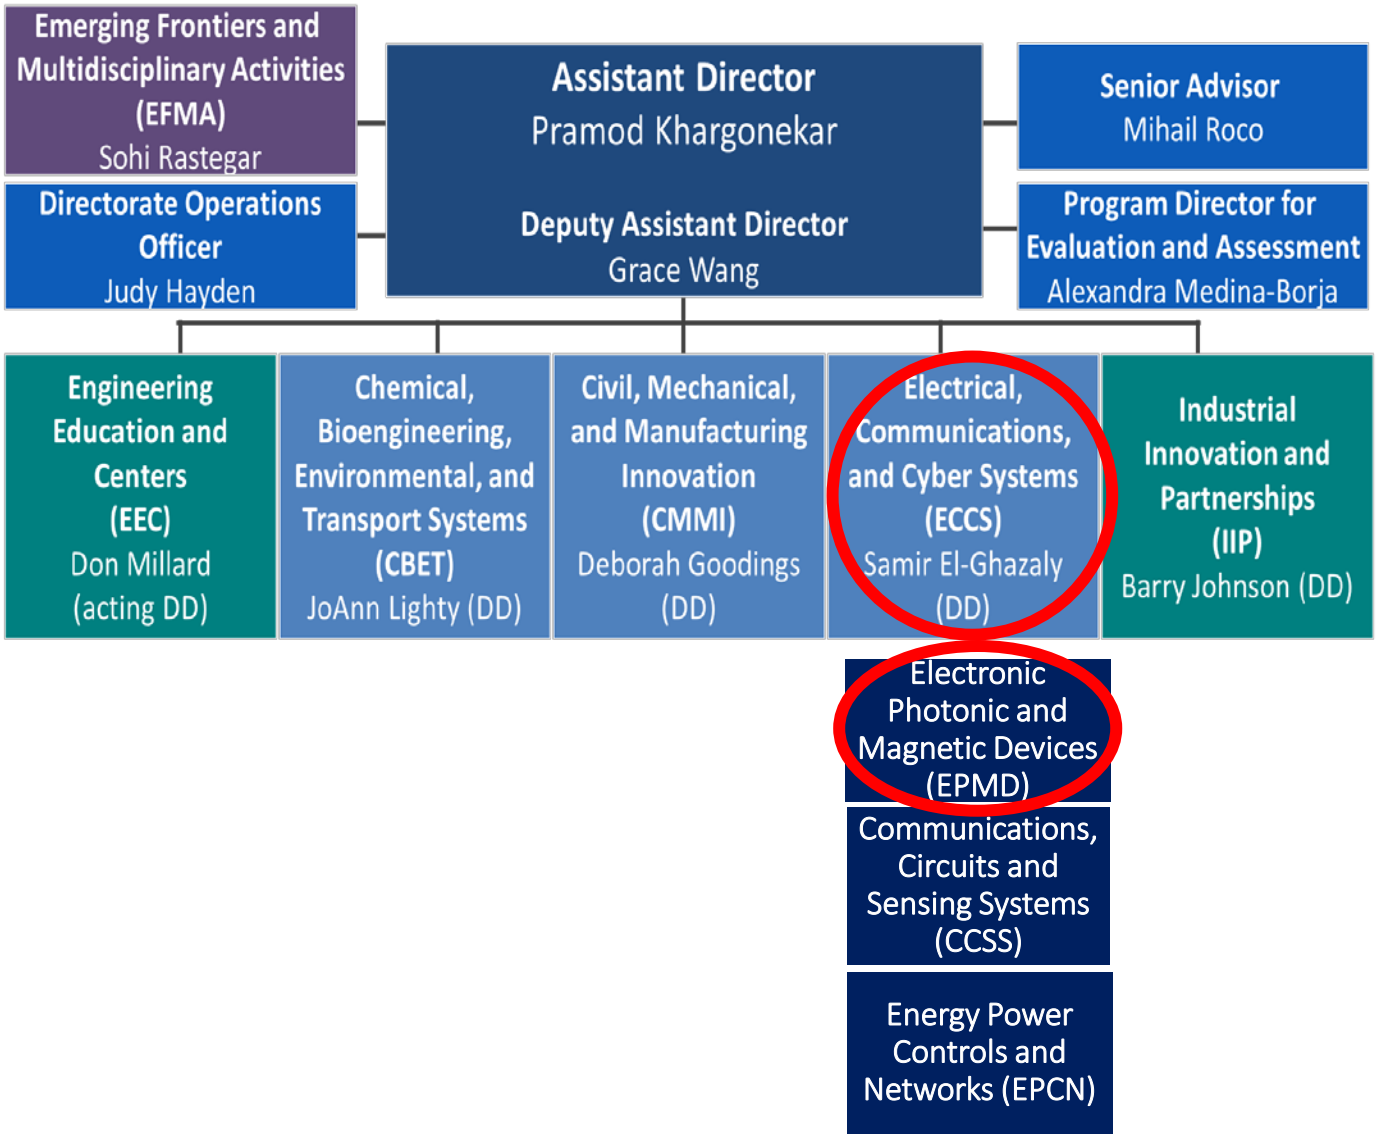

Semiconductor Research Corporation:
A Family of Distinct, Related Program Entities

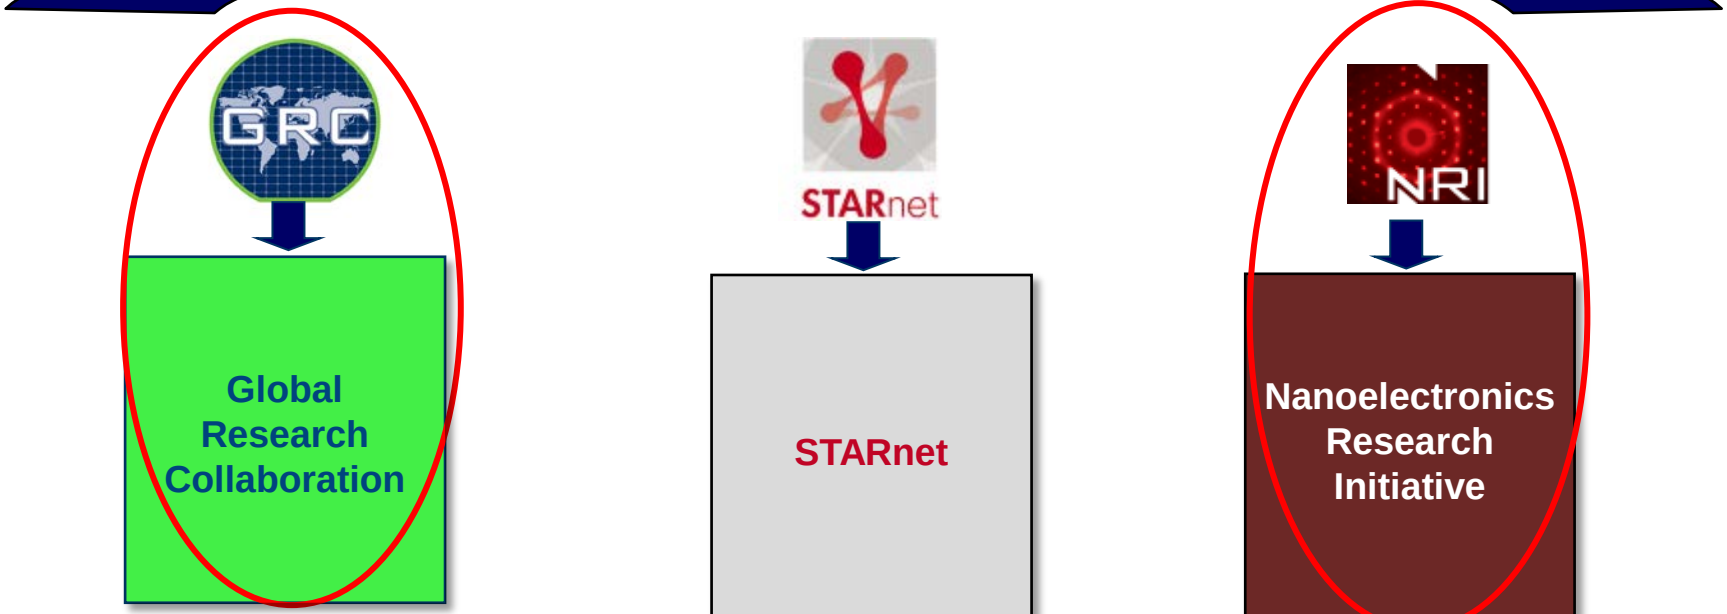

AMD
ARM
Analog Devices
Freescale
GLOBALFOUNDRIES
IBM
Intel
Mentor Graphics
Microsoft
Mubadala Tech
Qualcomm
RTI International
TI

Gov't Participants
State of Arizona
State of Texas
State of NY
NIST
NSF

Other Participants
SEMI
SIA
MOSIS

GLOBALFOUNDRIES
IBM
Intel
MICRON
Raytheon
TI
United Technologies

Government Participant
DARPA

IBM
Intel
MICRON
TI

Gov't Participants
NIST
NSF
State of Nebraska
State of NY
State of Texas

Broad Program Goal:

NSF and SRC are partnering to support new research to set the stage for **the next paradigm of computing, suitable for applications from mobile devices to data centers.**

Program Motivation: Slowing Progress in Computing

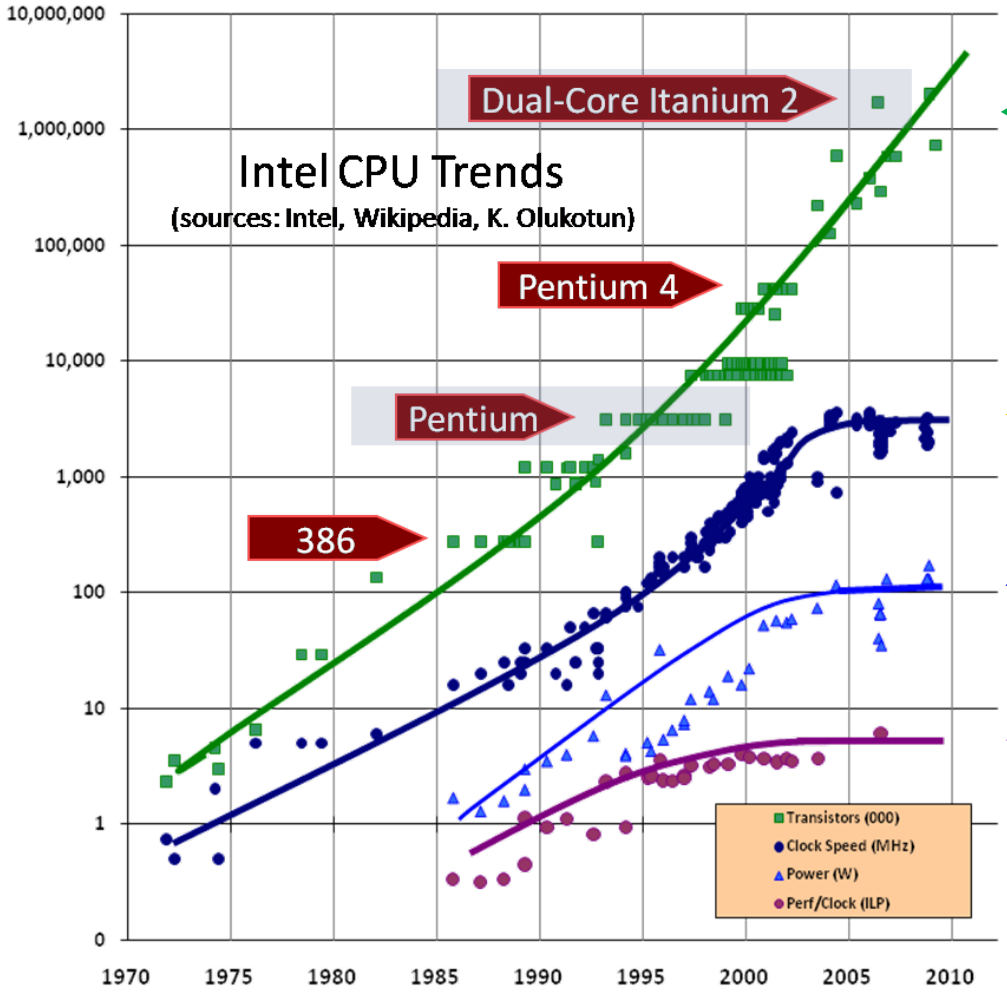

Transistor Count
(still growing)

Clock Frequency
(flat since 2003)

Power
(flat since 2003)

Architectural Innovation (ILP)
(also flagging!)

http://www.gotw.ca/publications/concurrency-ddj.htm

Program Motivation:
Briefly summarizing the last 10 years…

- To keep areal power density and total power within economically acceptable limits, industry froze clock speed and slowed the deployment of multiple cores.
- To escape this new status quo, there is an urgent need to explore devices with switching mechanisms that are fundamentally different from that of the conventional FET, and architectures than are fundamentally different from the von Neumann architecture
- The explosion of unstructured data and the urgent need for future computing systems to efficiently provide the 'insight' to transform raw data in to actionable information has created a need for architectures than are also fundamentally different from the von Neumann paradigm.

Proposals

Should address one or both of the following research paths:

1. Novel architectures and attendant devices and interconnect technology aimed at achieving the highest levels of system connectivity and energy efficiency, and integrates programmed arithmetic abilities with heuristic learning and predictive capabilities on the same physical platform.
2. New device concepts and attendant circuits and architectures for greatly reduced computational energy dissipation.
   See - https://www.src.org/nri/energy-efficient-computing-workshop.pdf

Should aim for solutions which are:

- disruptive, truly game-changing.
- scalable, from mobile platforms to data centers.
- extendable, capable of sustaining the long-term vitality of the IT revolution.

Characteristics of an Ideal Scalable, Extendable Platform

- Orders of magnitude higher levels of computational energy efficiency
- Substantially increased system-level functional density
- Workflow adaptive, to address the broadest range of applications
- Neuro-inspired cognitive learning and predictability, integrated on the same physical platform with ultra-energy efficient arithmetic computation
- Embodying co-design/co-optimization of materials & devices-to-system architecture through cross-disciplinary, collaborative research, guided by a common set of system-level, workflow requirements

Examples of Research (Path 1)

- Proposals utilizing revolutionary connectivity technologies such as plasmonic, photonic, terahertz or any other approach that can enable a dramatic lowering of overall system energy dissipation
- Interconnect technologies that enable functionality (such as embedded 'intelligent' routing, etc.) beyond point-to-point connectivity and the architectures that implement them
- Proposals merging heuristic learning and predictive functionality on the same physical platform as a programmable algorithmic capability, for example …
  - non-digital or hybrid digital-analog architectures,
  - brain-inspired architectures including reservoir computing and other neuro-inspired approaches,
  - stochastic, approximate, memory-centric or merged logic-memory, or other such fundamentally novel, heterogeneous architectures
- "Sprinting" architectures allowing portions of a system to briefly exceed time-averaged power dissipation limits, to accelerate operations that are critical to overall system performance.

Examples of Research (Path 2)

- New Device Concepts
  - Steep-slope (transistor-like) devices based on gated phase transitions or novel transduction principles, and more
  - Hysteretic (memory-like) devices based on antiferromagnetic or ferrimagnetic materials for *fast* magnetoelectric switching, and more
  - Oscillatory (energy-conserving) systems for Boolean and non-Boolean computation

- New architectures
  - Enabled by hysteretic devices combining non-volatile logic and memory (compute in memory, neuromorphic systems, reconfigurable logic, and more)
  - Enabled by oscillatory systems (neuromorphic architectures based on arrays of coupled oscillators; computation in the network, enabled by energy-conserving optical logic, and more)
  - Other areas of architectural research that leverage emerging devices (Stochastic and Approximate Computing, heterogeneous architectures, and more)

In summary: Potential research topics have been mentioned for illustrative purposes only.

What is important is that the proposed research address the following critical needs:

- **Revolutionary hybrid architectures** with programmed arithmetic logic integrated on the same physical platform with neuro-inspired adaptive learning and predictive capabilities
- **Revolutionary device concepts** and associated circuits and architectures which greatly extend the practical engineering limits of energy efficient computing
- **Novel materials for on-chip conductors,** scaled to dimensions below 10 nm where quantum scattering effects are dominant, with substantially reduced parasitics and potentially, embedded 'routing intelligence'

Prior Workshop References

This solicitation has been guided in-part by outcomes from the following and many other workshops:

- **Energy Efficient Computing:**
(https://www.src.org/nri/energy-efficient-computing-workshop.pdf)

- **Rebooting IT:**
(https://www.src.org/newsroom/rebooting-the-it-revolution.pdf)

- **Cross-layer Power Optimization (CPOM):**
(http://sportlab.usc.edu/cpom/final_report.pdf)

- **Emerging Technologies for Interconnects:**
(http://www.nsf.gov/awardsearch/showAward?AWD_ID=1148697)

Energy Efficient Computing: from Devices to Architectures (E2CDA)

Solicitation Requirements

Award Size and Limits

- 'Type I' Proposals:
  - Collaborative, multi-disciplinary teams
  - $800k-$1.6M/yr. for 3 years (approx. 2-4 awards)
- 'Type II' Proposals:
  - Individual projects
  - $100-$200k/yr. for 1-3 years (approx. 2-4 awards)
- Potential for additional 2 years of funding, depending on success of the projects and availability of funds
- Proposal Limits:
  - No limit per institution
  - 1 proposal per Principle (or Co-) Investigator

Organization Eligibility and Limit

- **Organization Guidance:**
  - E2CDA proposals may be submitted by a single organization or a group of organizations consisting of a lead organization in partnership with one or more partner organizations.
  - Only U.S. academic institutions which perform research and with degree-granting education programs in disciplines normally supported by NSF are eligible to be the lead organization.
  - Academic institutions are defined as universities and two- and four-year colleges (including community colleges) accredited in, and having a campus located in the United States, acting on behalf of their faculty members.
  - Principal investigators are allowed to form synergistic collaborations with government laboratories, industrial researchers, and scientists and engineers at foreign organizations where appropriate, though no NSF funds will be provided to government labs, industry, or foreign organizations.

Full Proposals **Due: March 28, 2016**

- Must meet the **requirements** of NSF Grant Proposal Guide (GPG) www.nsf.gov/publications/pub_summ.jsp?ods_key=gpg

- **Project Summary**
  - Proposals that do not separately address both intellectual merit and broader impacts will be returned without review

- **Project Description** (15 page limit) includes:
  - Vision and Goals
  - Approach and Methodology
  - Team Description and Capabilities
  - Impact Section (at end): Describe the transformative aspects of the project including: how the synergy of experts from different disciplines in the proposed research will achieve a significant advancement in fundamental knowledge and how it will have a strong potential for ultimately achieving the E2CDA goal.

- **References** Cited

REVIEW CRITERIA

- NSB-approved Merit Review Criteria
  - Intellectual Merit
  - Broader Impacts

- NSF Staff will give careful consideration to the following:
  - Integration of Research and Education
  - Integrating Diversity into NSF Programs, Projects and Activities

REVIEW AND AWARD PROCESS/SCHEDULE

- Full Proposals are due on **March 28, 2016**.
- Awards are expected to be made, pending availability of funds, by **July 2016**.
- Grantee Kick-off Symposium is tentatively planned for **Fall 2016**. (web-based, no travel necessary)
- PIs are advised to budget travel to annual Program Review meetings

Questions and Answers

Question: **[PI Limit]**

Can an investigator be PI on proposal and Co-PI on another proposal?

Answer:

No. Each person can only be a PI or Co-PI on one proposal

Q and A (cont'd)

Question: **[INTERDISCIPLINARY]**

Is it required to have a PI/co-PI from different academic departments?

Answer:

No, however for 'Type I' proposals, it is expected that 'multi-disciplinary teams' will include participants from multiple departments – perhaps also from multiple universities.

Q and A (cont'd)

Question: **[Type I proposals]**

Is there a minimum number of PIs expected for Type I proposals?

Answer:

No, as long as the scope of the work proposed is adequately addressed by diversity of expertise.

Q and A (cont'd)

Question: **[Industry co-PI?]**

Can a person from industry serve as co-PI?

Answer:

Yes, but only under the following conditions: When collaboration with industrial researchers is appropriate for the proposed research, the GOALI mechanism (Grant Opportunities for Academic Liaison with Industry) http://www.nsf.gov/publications/pub_summ.jsp?ods_key=nsf10580) may be used. Please pay particular attention to the format needed for GOALI to ensure that your proposal is eligible for review.

Q and A (cont'd)

Question: **[Participating Institutions/senior personnel]**

How many participating institutions/senior personnel are allowed on an E2CDA proposal?

Answer:

There is no limit to the number of participating institutions or senior personnel allowed.

Q and A (cont'd)

Question: **[Type I vs Type II proposals]**

Do the proposals need to be marked Type I or Type II at the time of submission?

Answer:

Yes. From the solicitation: "Proposal titles must start with **E2CDA** followed by a colon, then specify the class (**Type I** or **Type II**) followed by a colon, and then provide the title of the project."

Q and A (cont'd)

Question: **[Proposals from foreign institutions]**

Can a researcher from a non-US institution participate in a proposal?

Answer:

Yes, however, any proposals from foreign universities must have a U.S. collaborating university who in turn must be the entity submitting such joint proposals to this solicitation.

Should they be chosen for funding, NSF regulations preclude funding of universities outside the United States. However, proposers are encouraged to secure funding from other organizations, including SRC, for those portions of proposals which would be executed by foreign university collaborators.

Q and A (cont'd)

Question: **[Proposals from non-academic institutions]**

Will proposals from non-profit research institutions be eligible for awards?

Answer:

Only academic institutions are eligible for funding through this solicitation. However, individuals from non-academic institutions can participate as collaborators at the discretion of the proposing principle investigators through sub-contracts, etc.

Q and A (cont'd)

Question: **[Collaborations with Industry]**

Are collaborations with SRC member companies encouraged?

Answer:

Such collaborations are not required in generating a proposal. Mentors/liaisons from industry will be assigned if a proposal is funded.

Q and A (cont'd)

Question: **[Collaborations with National Labs]**

Will collaborations with National Labs be encouraged? If so, how will IP be handled?

Answer:

Collaborations which would increase the merit of a proposal are encouraged, but National Lab engagement must not alter SRC and NSF guidelines regarding IP rights, etc. as referenced in the solicitation text.

Q and A (cont'd)

Question: **[Metric targets]**

In the solicitation there was an example given of a metric for performance and energy efficiency. Can you give more details?

Answer:

The metric (> 1G MAC/s/nW) was suggested as one example of a possible targeted goal for that portion of a new system which could be measured in such a way. It is extremely aggressive ('orders of magnitude beyond any projections of existing system roadmaps'), but hopefully conveys the need to create truly novel paradigms and not simply evolve existing architectures and technologies within this program.

The solicitation also suggests that new metrics be developed for the (integrated) portions of the proposed systems which would be aimed at creating 'heuristic learning and predictive functionality'.

Q and A (cont'd)

Question: **[Intellectual Property]**

In the solicitation there was a requirement stated regarding disclosure of any potentially 'blocking background IP' in the proposals. Can you expand on that as well as what the requirements would be for handling any foreground IP?

Answer:

Background IP becomes 'blocking background IP' when it would be necessary to obtain a license to that IP in order to utilize IP created during the course of the proposed research. Proposers must disclose background IP. SRC will resolve issues regarding blocking IP prior to awarding an SRC contract. NSF funding will not be contingent upon resolution of any blocking IP, and funding by SRC or NSF will not create an obligation for the other organization to provide funds.

Foreground IP is new IP created through execution of awarded proposals. All research contracts for awarded proposals will be funded from pooled NSF-SRC resources and will grant limited IP rights to all SRC member companies for any IP created within the scope of the proposed research – i.e., 'non-exclusive, royalty-free, right-to-practice'. *All inventors and/or their institutions will retain ownership of any and all IP they disclose.*

Q and A (cont'd)

Question: **[Uniqueness of proposed research]**

Does an E2CDA proposal have to be fully orthogonal to research that is already funded?

Answer:

No, but proposers should clearly distinguish their proposed research from research that is already funded. We hope that some truly novel research ideas will be proposed and funded through E2CDA, but we recognize that even the most novel ideas have precursors and even the most original research builds on previous research.

Key website Address

- **Please refer to E2CDA website for up-to-date information:**

  http://www.nsf.gov/pubs/2016/nsf16526/nsf16526.htm