Welcome to the NSF Webinar on our newest solicitation on hardware security.

I am Keith Marzullo, the division director of the Computer Network Systems division at NSF. And seated next to me is Celia Merzbacher, vice president of innovative partnerships at the Semiconductor Research Corporation (SRC). Also joining me at NSF is Jeremy Epstein, lead program officer for the Secure and Trustworthy Cyberspace (SaTC) program, and by phone Ralph Wachter and Angelos Keromytis from the SaTC program who are also part of the STARSS program.

We are very pleased to announce this joint solicitation between NSF and the SRC, titled Secure, Trustworthy, Assured and Resilient Semiconductors and Systems program, which is also known as STARSS.

NSF and SRC have worked together before this, but this solicitation is the first time NSF and SRC have joined together on a research program in cyber security. It recognizes that partnering together we can advance and accelerate research in this vital area.

Also listed on the slide are several other program directors at NSF for STARSS who will manage the review panel process. The selection of awards will be a joint decision by NSF and SRC.

Before I get into the slides, I'll mention that these slides will be posted to the NSF website after the conclusion of this webinar. We're going to try to keep this to an hour, but can stay longer if needed for Q&A.

We'll begin with a discussion of this unique partnership and the motivations behind it.

We'll follow with the major theme of the program, which can be summarized as "Design for assurance". We will also discuss the specific topic areas of interest.

Because SRC will be an active participant in this process, we will carefully highlight the special aspects of the program administration and proposal selection. In particular how it differs from proposal administration for the nominal SaTC program.

We'll close out our prepared remarks on questions that we think the community might ask about the program and then open the discussion to anyone on the Webinar.

As I alluded to at the beginning of this presentation, NSF and SRC have a successful history of collaboration. Previous joint solicitations have focused, for example, on Failure Resistant Systems and Multicore Architectures. These collaborations provide a number of benefits to our researchers, including the chance to have deeper interactions between university researchers and students with industry technologists. Such interactions can offer insights on the needs and capabilities of industry, provide mentoring of students interested in industry careers, and facilitate technology transfer. The interactions also offer the chance to influence and educate industry on innovative approaches and results.

This year, we are fortunate to be partnering with SRC on the Secure, Trustworthy, Assured and Resilient Semiconductor and Systems – or STARSS - program, which is a new activity aimed at making

semiconductors and hardware systems more reliable, trustworthy and secure. This program expands NSF's Secure and Trustworthy Cyberspace program, known as SaTC, which is a cross-disciplinary, cross-directorate program devoted to all aspects of cybersecurity.

At this point, I'll turn it over to Celia Merzbacher from SRC.

Thank you, Keith. For those who are not familiar with SRC and how it operates…

SRC was established in 1982 by a group of semiconductor companies to help each member to be more competitive by collectively funding precompetitive university research. Through SRC's engagement with and funding of universities:

- Ideas are generated to help industry address technology challenges,
- Fruitful interaction take place among industry and academic experts, and
- A robust pool of talent (for ongoing research and future employment) is sustained

Through a variety of activities and processes, SRC aims to maximize the value of its research investment.

A key to SRC's longevity is that it provides benefits to all participants—whether in industry, government or academia. Benefits to university researchers who receive SRC funding include:

- Guidance on industry needs
- Funding, often in partnership with government, as in the case of STARSS
- Input and feedback during research
- Industry contacts and relationships, and
- Pathways to practical application and technology transfer

Since its inception, SRC has invested roughly $2B of industry funding, creating and sustaining a university research enterprise with expertise that is critical to the ongoing advancement in semiconductor science and technology.

STARSS is part of a new SRC thrust focused on Trustworthy and Secure Semiconductors and Systems, or T3S with initial funding from AMD, Freescale, Intel and Mentor Graphics.

Several trends are impacting the security and trustworthiness of semiconductor-based hardware.

Semiconductors, or integrated circuits, are increasingly pervasive. They are part of the devices and systems that we rely on—individually and as a society and a nation. They are key to our economy, our security, our well-being, and our quality of life.

In addition to the increasing prevalence of semiconductors, they are more networked and interconnected.

Moreover, products--and the design and manufacturing processes for making them-- are becoming more complex. (A single processor can have billions of transistors and multiple functions—e.g. video processing, communication, computation, storage, sensing, etc.) Many design or IP blocks used in a design come from various third parties.

Growing complexity is paralleled by a more fragmented—or lengthy-- and global supply chain, with many components produced and processing steps performed by businesses located around the world.

These trends are leading to:

- *increased vulnerabilities*
- *greater impact if a chip fails, and*
- *More --and more attractive-- targets for attack*
- This graphic shows the path from an idea for a functional processor to the actual chip.
- There are numerous steps from the initial description developed by “architects”, through increasingly detailed and specific representations and eventually conversion from an abstract expression to the physical layout that is ultimately manufactured.
- The entire design and manufacture chain involves many individuals, at many firms, typically located in globally distributed locations.
- Between each step the design is verified to insure that the more detailed and eventually the physical version does what it was designed to do.
- To provide assurance regarding the final product, the processes from end to end must be secure. From the architecture and specification through verification and finally the ability to authenticate. And in order to make progress, we need to be able to measure security and trustworthiness. These capabilities underlie the priorities for STARSS.
- In addition to possible errors, bugs, and design weaknesses that can lead to security vulnerabilities—such as side channels or backdoors, hardware is vulnerable to counterfeiting and tampering—both of which can cause a chip to not perform as intended.
- This schematic shows the same flow as in the previous slide—but oriented vertically. In addition points in the process where tampering or counterfeiting can occur are highlighted. Some of these attacks are more likely than others.
- It should be noted that a number of measures are in use today to thwart these threats. But the growing number of “targets” for attack and the increasing sophistication of attackers is creating a demand for tools and strategies that provide even greater assurance.

Threats and challenges to security and trustworthiness of hardware can take many forms, including those shown here, and can arise from various potential points of weakness or attack.

- Weak specification, design or implementation can lead to unintended functionality or unauthorized access or control
- Side channels can result in data exfiltration or the ability to hamper proper performance
- In addition to unintentional weaknesses or bugs, tampering by an adversary—e.g. inserting a hardware Trojan—during design or manufacture can impact hardware performance
- Interfaces between modules or IP blocks may provide points of attack
- Design IP and tools provided by third party suppliers are not generally verifiable for security properties
- Authentication techniques that are easy to implement and assess but hard to counterfeit are needed as well as methods for tracking provenance of a product in the supply chain

STARSS seeks to address these threats and challenges with new technology-based strategies and techniques.

With these challenges, how can we be assured that a chip does what it is supposed to do… and nothing else?

Our goal is to develop strategies and tools that enable the design & manufacture of semiconductors and systems that are secure, trustworthy, assured and resistant to attack or counterfeiting.

The Objectives are to make semiconductors and systems not only functional, but also secure and trustworthy. That is to:

Avoid unintended or unwanted behavior, access or control

Increase resistance to tampering or counterfeiting, and

Improve the ability to authenticate– whether it is 3rd party IP or a final product.

We are looking at these problems from a science and technology perspective and with a view towards ultimately being able to transition the results to practice. SRC member companies are interested in novel approaches for managing the dynamic threats going forward.

I will now turn it over to Jeremy Epstein to describe the solicitation in more detail.

STARSS is interested in receiving proposals on any idea that addresses the Goals & Objectives.

Some topics of <particular> interest to the STARSS program include:

New architecture and design approaches and frameworks for specifying and reasoning about the security of hardware.

New security-specific properties and principles, as well as metrics to measure security.

Threat assessment to identify, classify, analyze and share information about security threats in hardware.

Hardware specific security verification and analysis techniques and tools.

Tools and frameworks to develop a semiconductor security development model to guide semiconductor design and manufacture.

New models for authentication and attestation.

In each area, we are looking for new approaches that offer more general solutions rather than ad hoc or narrowly applicable concepts.

I will elaborate on each of these areas in more detail on the next slides. Note that the topics are inter-related.

In the space of architecture and design, we seek new approaches, and models and frameworks for specifying, measuring and reasoning about hardware-specific security.

Also of interest are methods for ensuring that a security-specific IP block is secure, and also for ensuring that there are no security-related vulnerabilities resulting from unintended behavior or side effects related to any other IP blocks.

Research on novel design or specification languages that are “security aware”, for example through quantifiable security attributes, are also of interest.

Areas NOT of particular interest include “yet another design decryption or specification capture language”. This would be an example of an unduly narrow topic, as mentioned earlier.

.

An area of interest is properties, principles and metrics for security of hardware.

The program seeks new high level hardware design principles and semiconductor-specific properties that go beyond high-level security properties such as confidentiality, integrity and availability of security-sensitive assets and access mechanisms.

Also of importance and interest is the development of a knowledge base of concrete examples, scenarios, and other empirical evidence. Development of new metrics that provide a measure of the security of a particular design is of particular interest.

Another example of a topic of interest is a metric based on general modeling of threats, system vulnerability that correlated and applies to a variety of hardware and applications, such as embedded, IP block, etc. An example not of interest would be another metric for a specific type of hardware or application.

In order to assess the threats in the semiconductor community, one needs to be able to identify, classify, analyze and share information about security threats in hardware.

Research on designing a dynamic information base that captures this information from either unintended vulnerabilities or those included as part of a malicious design or fabrication is needed.

Research into taxonomies and representations of hardware-related security threats is also of interest.

Tools, techniques, and methodologies for verifying hardware-specific security properties and enforcing the security design principles are essential.

This includes techniques to ensure coverage and equivalency regarding security between various design, implementation, integration, and manufacturing phases.

These can be extensions and/or enhancements to existing tools and methodologies, intersecting existing design and verification process flows.

Research in this area includes novel abstraction and mathematical representation of hardware security, threats and vulnerabilities

Regression and other testing methodologies for security are also of interest.

There is need for the semiconductor design and manufacturing equivalent of leading software security engineering models and frameworks, such as IBM's Secure Engineering Framework.

A comprehensive semiconductor security development model would provide a framework for responding to vulnerabilities and would offer ways to measure organizational maturity with respect to security of products, as well as to assess product assurance over time.

Such a development model could guide academic and industrial curricula targeted at instructing architects, designers, and engineers.

Research in this area could be facilitated by access to industry practices.

New models for authentication and attestation are of great importance.

These could be models for the insertion of artifacts and/or design elements that are verifiable throughout the design and manufacture pathway.

Also of interest are techniques for dynamic verification in the field and non-destructive authentication.

Of great importance is a semiconductor provenance model and other related design artifacts including, for example, hardware fingerprinting and third party design element model checking.

Other issues, such as the generation, protection and establishment of trust models for hardware-implemented keys, are also of interest.

Let's move now into the logistics of the program.

Proposals to this joint program will be submitted to NSF with the following requirements

- All proposals will be up to 3 years with a total award value of $500K.
- The submission deadline is 5pm proposers local time 26 March 2014
- The title should start with the prefix “SaTC-colon-STARSS-colon”
- Only one proposal per PI/co-PI to STARSS is allowed
- This limit is separate and distinct from the SaTC program.
- No proposal containing or referencing classified material will be accepted.

Proposals must follow all the usual NSF policies, including inclusion of biographical sketches, budgets and budget justifications, lists of references, data management plans, etc. Proposals are submitted through Fastlane or grants.gov, as with any NSF proposal.

One requirement that is different from the usual submission to NSF is a statement of consent that gives NSF permission to share the proposal, reviews and any other related information with SRC. This statement of consent needs to be uploaded into Fastlane or grants.gov as a supplementary document. There is no particular format for this statement, but we expect it will be no more than about a page.

SRC will treat your proposal as proprietary.

The proposal review process will be conducted by NSF, according to the standard rules and procedures in place for NSF Small proposals.

The review and award recommendations will be coordinated by a Joint NSF and SRC Working Group of program officers from both NSF and SRC.

Projects that are selected for joint funding by both NSF and SRC will be funded under two funding instruments. That is, NSF support will be provided via an NSF grant and SRC support will be provided via an SRC contract.

Note that not all projects will necessarily be jointly funded, NSF or SRC may decide to fund certain projects separately.

The budget submitted with the proposal should include all necessary project funds without regard to the two funding organizations. The budget will be divided subsequent to selection.

All awards involving SRC funds shall be made under a contract that provides for non-exclusive, royalty free rights to all SRC members for any intellectual property generated as a result of the SRC funded research.

NSF and SRC will manage their respective awards or contracts according to their own procedures and guidelines.

All awardees must submit annual reports to the respective funding agency. If a project is co-funded by NSF and SRC, then the PI must submit reports to both NSF and SRC, which are similar in content.

One or more project representatives from an awarded STARSS project must attend the first PI meeting which is currently planned for the Fall of 2014.

And in years in which no SaTC PI meeting is held, SRC will hold a review of all SaTC:STARSS projects.

You will have a chance to ask questions shortly, but first we will review some of the questions we've received to date that we believe are of general interest.

Q: Can I submit the same proposal to SaTC and STARSS?

A: No, you must choose one program. You may submit related proposals to both, but not the same. See the NSF Grant Proposal Guide, and in particular note that you must list the related proposal, and describe how it differs.

Q: Does a submission to STARSS count towards the 2 proposal limit for SaTC?

A: No, you can submit 2 proposals to SaTC and 1 to STARSS.

Q: I just submitted the perfect STARSS proposal to SaTC, what should I do?

A: You can withdraw the proposal and submit to STARSS.

Q: Can a STARSS proposal include a Transition to Practice option, which allows for additional funds?

A: No. TTP options are for SaTC Small/Medium/Frontier proposals only. However, STARSS awardees will work closely with industry to enable transition. An NSF supplement to a STARSS grant could be used to fund this transition effort.

Q: How do I decide whether to submit to SATC or STARSS?

A: If you are interested in working closely with industry and/or in SRC co-funding then please submit to STARSS. Talk to a SaTC or SRC program director if you are unsure. Proposals on hardware security may be submitted to either SaTC and STARSS. That said, SaTC is a large, broad, interdisciplinary program that includes hardware, but also spans software, theory, networking, social sciences, and numerous other topics. STARSS, a joint effort between NSF and SRC, is more narrowly focused on security aspects of hardware and offers an opportunity to interact more with industry.

Q: Does the new STARSS program mean that SATC will be reducing funding for hardware security?

A: No. Hardware security is still a priority for SATC.

Q: Are only hardware proposals in scope for STARSS?

A: Yes. STARSS is hardware-oriented, but interfaces with software, including microcode, firmware, and software tools that are used to design circuits and systems.

Q: How many awards will be made?

A: We expect to make 6 STARSS awards, subject to the availability of funds.

Q: Is there a difference contractually between an NSF award and SRC award?

A: NSF awards "grants" whereas SRC awards "contracts". The standard SRC contract has been executed with hundreds of university and includes requirements for certain deliverables, including annual reports on research progress, copies of publications, and student information.

Q: If co-funded, can the same proposal be used by NSF & SRC

A: Yes. There's a single "statement of work" based on the technical description in the proposal, but there are two funding agreements. The proposal you submit should include all required funds; NSF and SRC will determine how to split the costs between the organizations and will provide guidance on revising the budget.

Q: Will SRC fund participants at non-US institutions (which are not funded by NSF)?

A: No. All funded participants must follow standard NSF eligibility requirements. Collaboration with non-US organizations is encouraged.

Q: Is this new money in addition to existing SaTC funds?

A: Both NSF and SRC are putting in funds. NSF's funds are mostly from the SaTC program; SRC's funds are a new investment.

Q: Are there any restrictions on intellectual property since SRC is involved

A: Under the standard SRC contract, the University retains ownership and SRC receives a non-exclusive royalty free license for its members to any IP developed.

Q: Can we get access to SRC member technology

A: This would be handled on a case by case basis between a member company and university. In the past, information has been made available. On occasion, an NDA has been executed to allow for sharing of company proprietary information.

Q: Do I have to attend the NSF biennial SaTC PI Meetings and SRC reviews?

A: Yes, but not in the same year. The SaTC PI meeting will alternate years with SRC reviews; in years when there's a SaTC meeting, an SRC review may be added on, perhaps as a supplemental day.

Q: Do I have to submit reports to both SRC and NSF on a STARSS award?

A: Yes, assuming both SRC and NSF jointly fund the award. Typically, the report required by SRC and NSF are similar in content.

Q: Will SRC be actively involved as collaborators and working on spin-off projects, or provide in-kind support?

A: SRC provides for interaction and engagement between technical experts at member companies and funded researchers, which can result in a variety of other collaborations.

Q: Will SRC act as an industrial advisory team or as part of a larger advisory team?

A: SRC coordinates industry input and feedback through a number of mechanisms, including webinars, annual reviews, and through an industry Liaison Program that connects individuals from member companies with university research projects.