Introduction

The Server Component Resilience Workstream within the Open Compute Project (OCP) is releasing a request for research proposals on detecting and preventing Silent Data Corruption (SDC).

Defects and process excursions can occur during manufacturing of advanced integrated circuits using modern process technologies. Components that form part of these integrated circuits can also age over time in the field. These manufacturing aberrations can result in a multitude of behaviors, one of which is known as Silent Data Corruption. SDC's are computing errors that are not detected by either hardware or software. The computational errors can range from highly repeatable to intermittent to dependence on environmental conditions.

The need to rapidly identify Silent Data Corruptions and their root causes in compute systems is driven by modern computing SoCs having 10's of billions of transistors, cloud deployments of millions of nodes, 24/7/365 operation, all combined with applications that have high consequences associated with computation errors. Moving into the future, the problem will get exponentially worse unless innovations in the ability to detect and remedy are developed.

As part of the Open Compute Project (OCP), in October 2022 a Server Component Resilience working group was formed (under the Hardware Management Project) focused on Silent Data Corruption to address hardware-induced SDC challenges. The scope of this initiative includes CPUs, GPUs, and other hardware accelerators. The members of this working group are AMD, ARM, Google, Intel, Meta, Microsoft, and NVIDIA.

The high-level goals of the working group include:

1. Drive solutions and best practices that identify root causes and prevent SDCs.

2. Create awareness about SDC challenges across the computing community.

3. Partner & engage with the academic community to address SDC challenges.

4. Harvest research results in the form of technology transfers back to industry.

The research vectors below are provided purely as guidance to the principal investigators to focus the research tasks. They are by no means exhaustive.

Potential Research Areas for addressing SDCs

1. **Tests that can efficiently screen a large proportion of SDCs**
    a. New tests that provide expanded SDC detection for manufacturing or in the field. These can include solutions that fit into an open-source framework (ex: Open DCDIAG).
    b. Techniques to reduce SDC-detecting test times to a minimum while retaining reproducibility of failure.

c. Tests that can provide instruction/data level diagnosis such as instrumentation with additional code.
d. Techniques that leverage hardware knobs such as voltage, temperature, frequency etc.
e. Tests that can be executed concurrently with mission-mode operation (perhaps on other cores).
f. Tests which provide external instrumentation for SDC detection which can augment end-user application programs.
g. Tests which re-use DFT-based manufacturing test approaches in the field, including scan, BIST, and functional content preloaded into on-chip memories. Such approaches could be enhanced with new fault models, new pattern generation algorithms, and adaptive scheduling.

2. **Techniques in the software stack to detect/correct errors**
    a. Compiler level options
    b. Application-level coding techniques, operating system techniques, etc.
    c. Leverage microservices encapsulation, containers, virtual machines, etc. for detection and recovery.
3. **Pre-silicon and post-silicon coverage/susceptibility assessment and improvement**.
    a. Grading designs: Define and assess susceptibility of various architectural regions to SDC (example: Architectural Vulnerability Factor - AVF). Assessments to provide feedback to the designer of the susceptibility and point to the susceptibility sources.
    b. Grading tests: Define and assess the quality of tests (applied at system-level or manufacturing automated test equipment) based on uniqueness and accuracy of coverage, reproducibility, efficiency of the method, and cost. Assessments to provide feedback to the test creators/implementers.
4. **Hardware techniques for SDC detection and resiliency**
    a. Low-cost hardware design methods for detecting SDC in random logic (beyond arithmetic and logic units that are addressed with parity and residue-checking).
    b. Circuits or techniques that monitor and provide hints that an SDC-causal mechanism might exist in the circuitry.
    c. New design-for-test/diagnosis apparatus that can improve SDC detection (example: alternatives/improvements to scan).

Proposals should be submitted by April 3, 2024 via the form. Any questions should be sent to scr-feedback@ocproject.net.

The following provides a sample proposal outline form:

Cover Page (1 page)

- Title of the proposal: Provide a descriptive title of the proposed research.

- Names and contact information of all the university faculty (Principal Investigators) to be funded by this proposal
    - Names of the PIs and their contact information
    - Contact information of the sponsored programs office at the university. If the proposal involves a multi-university team, include the information of all the participating institutions.
- Names of any member company liaisons for the proposed research: List who Participants (if any) may serve as a contact point for the proposed work.
- The total funding amount requested
- Total cost-sharing amount (if any): This is the amount of non-OCP cost-sharing provided if the proposal were to be funded.
- In the spirit of OCP's tenet of openness, the research intends to be broadly usable by the community. How will you enable this? Where and when they intend to publish the research and/or collateral.

Executive Summary (1 page)

- Clearly describe the focus of the research and its expected impact.

Detailed Research Proposal (3 pages maximum)

- Describe the proposed research, its novelty, and expected impact on the state-of-the-art. Include relevant citations to publications.
- Include preliminary research results from the university team, as appropriate, to articulate the research plan and highlight its progress.

Statement of Work (2 pages maximum)

- Clearly describe the tasks in the proposed research, their schedule, and any deliverables.

Communication and Collaboration (2 pages maximum)

- Include details on how the university team plans to communicate their results with the OCP Resilience Workstream.
- If the proposed work is planned to be a collaboration with one or more companies, please include details on the collaboration plan. The collaboration plan will need to describe the specific roles of all the personnel involved (in the university team and industry), and how the Technology Track will be managed and coordinated.
- Please disclose any relevant research and ongoing projects and funding.

Proposal Personnel (2 pages maximum per CV)

- Include NSF-style two-page CVs of each university PI in the proposal.

Citations (no page limit)

Budget (no page limit)

- Include a cost proposal that breaks down the total budget across key line items, such as the number of students/post-docs to be funded, equipment, and travel.
- What type of hardware will be required to complete the research and is this research contingent on access to resources from working group members that your institution does not currently have.

Example References

S. -B. Park and S. Mitra, "IFRA: Post-silicon bug localization in processors," *2009 IEEE International High Level Design Validation and Test Workshop*, San Francisco, CA, USA, 2009, pp. 154-159, doi: 10.1109/HLDVT.2009.5340160.

Kuvaiskii, Dmitrii, et al. "HAFT: Hardware-assisted fault tolerance." *Proceedings of the Eleventh European Conference on Computer Systems*. 2016.

"SDCInfer: Inference of silent data corruption causing instructions," 2015 6th IEEE International Conference on Software Engineering and Service Science (ICSESS), Beijing, China, 2015, pp. 228-232, doi: 10.1109/ICSESS.2015.7339043.

"CARP: Handling Silent Data Errors and Site Failures in an Integrated Program and Storage Replication Mechanism," *2009 29th IEEE International Conference on Distributed Computing Systems*, Montreal, QC, Canada, 2009, pp. 385-394, doi: 10.1109/ICDCS.2009.58.

"Optimization of Tests for Managing Silicon Defects in Data Centers," 2022 IEEE International Test Conference (ITC), Anaheim, CA, USA, 2022, pp. 578-582, doi: 10.1109/ITC50671.2022.00076.

"Evaluating the Viability of Using Compression to Mitigate Silent Corruption of Read-Mostly Application Data," 2017 IEEE International Conference on Cluster Computing (CLUSTER), Honolulu, HI, USA, 2017, pp. 603-607, doi: 10.1109/CLUSTER.2017.99.

"Towards Transparent Debugging," IEEE Transactions on Dependable and Secure Computing, vol. 15, no. 2, pp. 321-335, 1 March-April 2018, doi: 10.1109/TDSC.2016.2545671.