

**Summary of Responses to
NSF/CISE
Request for Information on
Semiconductor Research and
Education**

**Sankar Basu, CISE/CCF
Erik Brunvand, CISE/CNS**

April 25,2022

Contents:	page
1. Executive summary	3
2. Background	5
3. Demographic information	7
4. Question 1 (Current and future research)	10
5. Question 2 (State of Education)	17
6. Question 3 (Missing facilities)	23
a. Question 3.1 (Specific technologies)	27
b. Question 3.2 (Overseas activities)	31
7. Question 4 (New research enabled)	34
8. Question 5 (Beyond facility access)	38
a. Question 5.1 (IP issues)	41
b. Question 5.2 (Industry connections)	42
9. Question 6 (Other considerations)	44
10. Endnote	48
11. Glossary of Acronyms	49

Executive summary

This document contains a summary of the information received in response to a Request for Information (RFI) sent out to the community supported by the NSF Computer Information Sciences and Engineering (CISE) Directorate. The questionnaire consisting of a set of six questions (with sub-questions) in the general area of semiconductor research and education, and topical impacts on the broader computing and information ecosystem. It was made available to the public including through a CISE *Dear Colleague Letter* to the community and the Federal Register. This summary is based on 254 responses received over a period of 81 days in fall of 2021. Most respondents - approximately 90 percent - were from academia, representing disciplines of electrical engineering, computer engineering, and computer science. This is reflective of NSF's academically centered stakeholder community. The questions spanned current and future research areas, such as the need for infrastructure, education and workforce development related to semiconductors. A list of what we consider the primary takeaways follows:

- **Novel architectures:** While the survey responses identify many research topics, the majority emphasized a variety of applications involving novel architectures, and their co-design and integration in substrates spanning silicon to emerging non-silicon nanotechnologies. Such applications include topics of national importance (e.g., artificial intelligence, machine learning, 5G/6G applications, and quantum computing). In the educational arena, the respondents highlighted the need for increased student enrollment as well as curriculum development and support for intellectual property (IP) and student fabrication and experiential learning at all levels for future workforce.
- **Affordable access to semiconductor fabrication:** A strong message in the responses is that an impediment in advancing on goals in both research and education is the lack of affordable and uncomplicated researcher, educator, and student access to semiconductor foundries. The survey indicates that a range of fabrication technologies are needed, ranging from mature legacy nodes to leading edge silicon CMOS processes and including their integration in non-silicon emerging technologies. Equally important are affordable electronic design automation (EDA) tools and software, process design kits, circuit IP, and easing NDA and export control issues. The diversity of these concerns could imply a need for national and regional fabrication facilities, and/or centralized tool and data repositories for semiconductor education and research. The survey responses also pointed to certain foreign countries that already provide access to such facilities, and recommend that emulating (e.g., IMEC, Belgium or CEA-LETI, France could be a strong way forward).
- **Future technologies:** Responses to the RFI's questions on future technologies resoundingly indicated that all areas of innovation in semiconductor research would be hugely impacted by the access to the fabrication facilities mentioned above. They involve all layers of the computing stack from materials, devices, and architectures to applications. In-house and/or shared facilities feature prominently among such needs.
- **Education and workforce development:** Semiconductor and microelectronics issues are fundamentally important to America's technical and economic future, and attention

must be placed on advancing and reinvigorating curricular and experiential learning for these topic areas. This includes bolstering the development of modern courses, framing their relevance to AI/ML in a way that connects to large enrollments in other topic areas, and complementing classroom learning with experiential fab-based projects.

- **Broad relevance:** Technology trends in semiconductors and microelectronics have reverberations into all aspects of computer systems design. RFI responses reflected the relevance of a broad range of application domains, architectural and design approaches, and technologies.
- **Other Considerations:** While issues related to semiconductor fabrication are seen in most of the responses, other considerations can also be observed. Such considerations include the need for scholarships and fellowships and industrial partnerships, student recruitment that includes underrepresented groups, tech transfer, startups, interdisciplinary approach to solving problems, and potential adverse effects of proprietary information.

Background:

On August 6, 2021, NSF published [Dear Colleague Letter \(DCL\) NSF 21-112: CISE RFI on Semiconductor Research and Education](#), and an RFI was initiated to get input from research and industrial communities on issues related to semiconductor research and education. The DCL was also sent on various program-specific mailing lists from within the Computer and Network Systems (CNS) and Computing and Communication Foundations (CCF) divisions, and shared in the CISE September 2021 monthly newsletter. The RFI was also announced in the Federal Register (document number 86 FR 35367: <https://www.federalregister.gov/documents/2021/07/02/2021-14159/request-for-information-national-science-foundations-directorate-for-computer-information-science>) on July 2, 2021. It was originally meant to be active until September 30, 2021, but the deadline was later extended to October 2021. The text of the DCL was as follows:

Semiconductor-related research, including underlying supply-chain, business, and economic impacts, are increasingly important to the Nation's long-term competitiveness and security. Through this Request for Information (RFI), the National Science Foundation's (NSF) Directorate for Computer and Information Science and Engineering (CISE) seeks input from those who are directly engaged in, or might potentially benefit from, CISE-related research and education in semiconductor and micro- and nano-electronics.

The computing stack has traditionally been viewed as a hierarchy of layers with devices and circuits comprising the lowest layers, and architectures, software, algorithms, and applications constituting progressively higher layers. Lower layers of the stack (e.g., devices, circuits, architectures) more directly involve semiconductor technologies to the extent that researchers may interact with large-scale fabrication facilities, but all levels of the stack are influenced by microelectronic advances to varying degrees. Thus, although in its entirety CISE research may not directly involve research on semiconductors, per se, the entire computing stack, from circuit design to architectures and on to software and applications such as sensor networks including the Internet of Things (IoT), embedded computing, next-generation wireless systems, large-scale data analytics, artificial intelligence (AI), edge and cloud computing, and high-performance computing, heavily depends on advances in this space.

As a result, much of the CISE directorate's portfolio is dependent upon advances in semiconductor technologies. For one example, tomorrow's AI innovations offer transformative societal impacts, but require advanced hardware capabilities that leverage newer semiconductor technologies. Conversely, the hardware design problem is a large, multi-objective, multiscale optimization problem that stands to benefit from the application of modern AI techniques.

On December 14-20, 2020, CISE funded a workshop focusing on the lowest levels of the computing stack. This workshop considered the scientific frontiers for semiconductor and microelectronics technologies as well as the needs for access to semiconductor foundries.

The workshop report is available at:

https://nsfedaworkshop.nd.edu/assets/429148/nsf20_foundry_meeting_report.

Building upon that workshop and report and given the diverse interests of the CISE directorate and community, the intent of this RFI is broader. Specifically, NSF/CISE seeks to:

- Gauge the extent to which the community's research and educational agenda would be enabled by the availability of new or different resources, or the re-introduction of resources that were available in the past. By this, NSF/CISE asks that respondents not restrict their answers to issues related to funding, but rather also consider issues related to infrastructure, facilities, access to tools/intellectual property/data, legal issues, etc., that support their research and educational agenda in the broader area of semiconductors.
- Understand what specific activities the research community would pursue and how that activity would impact societal and national interests, if the impediments mentioned in the first category above are removed. NSF/CISE asks respondents to be specific in making projections about new technologies potentially enabled by advances in semiconductor and microelectronics technologies within the 5-, 10-, or 15-year horizons, or longer. Also, if a respondent's research directly involves use of hardware fabrication, NSF is interested in learning specifics, as outlined in the questions below.

This RFI is issued solely for information-gathering purposes. NSF/CISE's intent is to analyze the responses received from this RFI for internal needs and for potentially formulating future program solicitations. NSF/CISE may make anonymized versions of the responses available to the public. This RFI does not constitute a formal solicitation for proposals.

Responses were collected to each question as unformatted text to give the respondents an opportunity to tell us what was most important to them in that area. We did not want to pre-bias the answers by using multiple choice or radio button type responses, although those would have been easier to collate. After we received the responses, we split up the questions and extracted information from each set of questions by hand.

Demographic information:

The data reported in the following pages was the result of a survey that remained active to the community during the period 08/06/2021 to 10/26/2021. A total of 254 responses were received. Not all responses included answers to every question requested. The number of responses to each question will be mentioned in the context of the respective questions.

The overall demographics of the 254 responses are shown in Figure 1. They consist of 90 percent university faculty, 4 percent industry professional, and an even smaller number 2 percent of nonfaculty university staff.

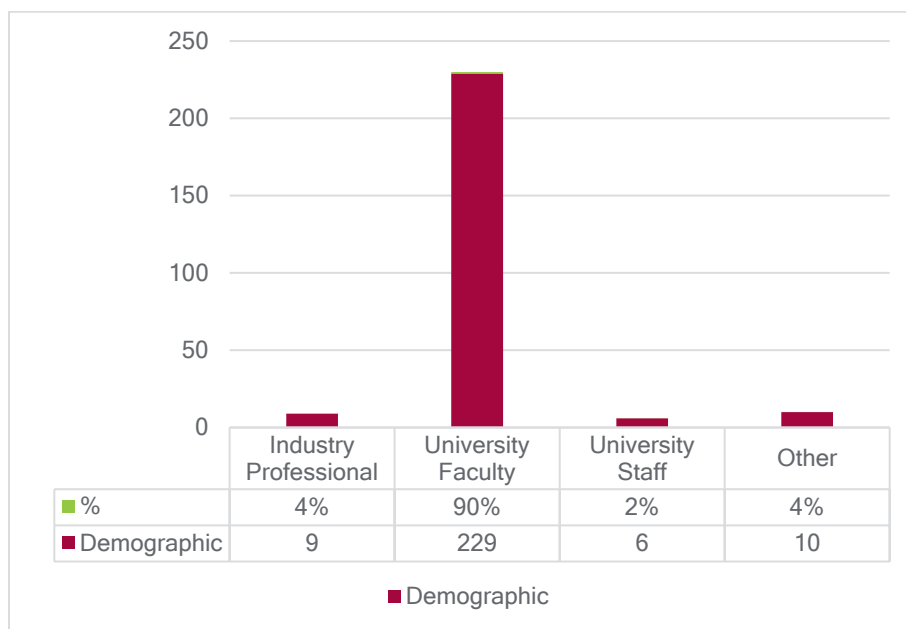


Figure 1: Breakdown of the responses according to the organization type.

University faculty respondents belong to academic departments of computer science, electrical and computer engineering, electrical engineering, material science and to a lesser extent other physical sciences as well. This breakdown is visually displayed in Figure 2.

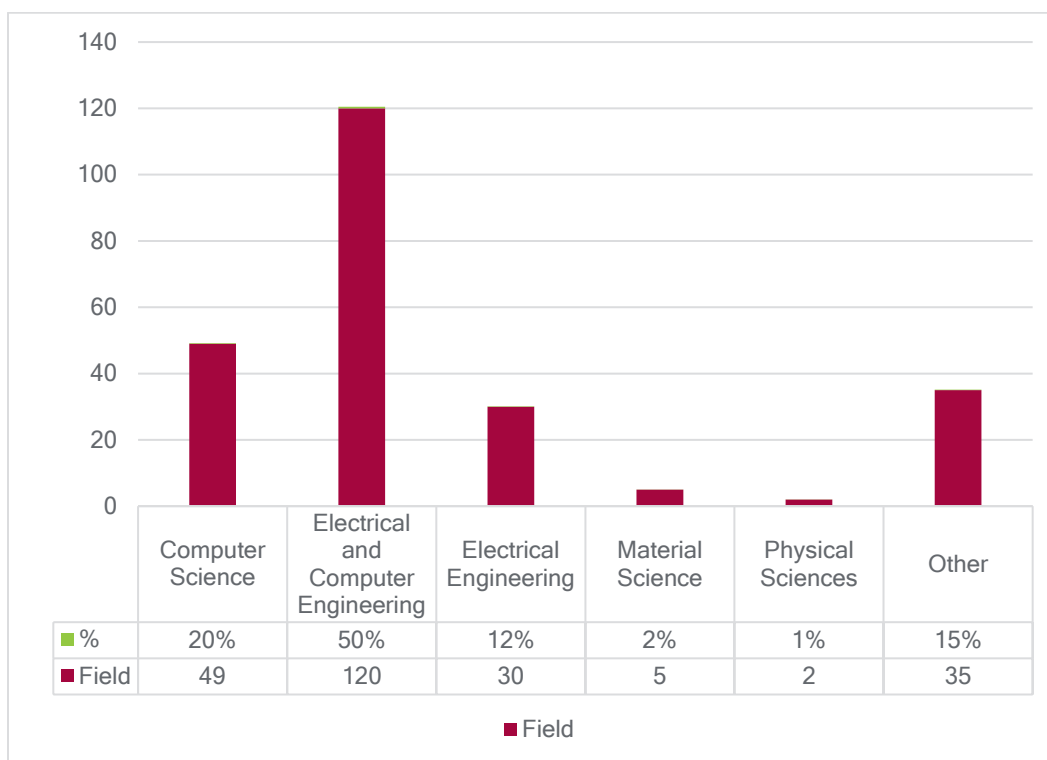


Figure 2: Breakdown of academic responses according to the disciplines.

A total of 127 academic responses came from 34 US states in varying numbers. More than 40 (42 to be exact) of these responses came from only three states (California, New York and Texas), whereas more than 50 percent (67 to be exact) of the responses were received from only six states (Massachusetts, Pennsylvania and Virginia, in addition). This suggests a heavy regional concentration of interest in the field. The distribution of the number of responses from various states is shown in Figure 3.

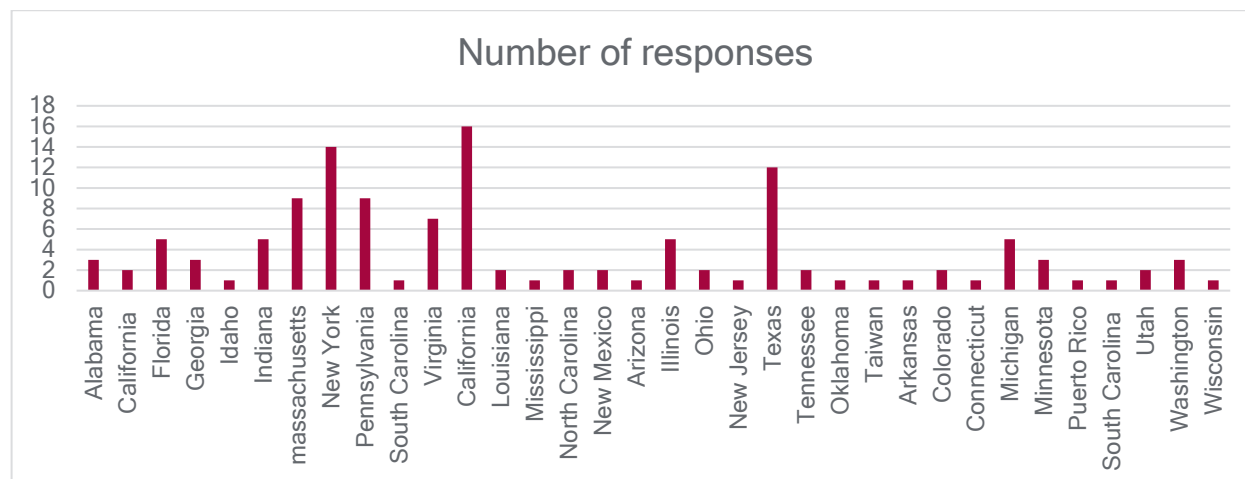


Figure 3: Distribution of the academic responses according to the US States and territories.

The survey requested answers to a total of six somewhat open-ended questions numbered 1 through 6. Two of the questions - Question 3 and Question 5 - each had a set of two further questions of detailed nature embedded in them in the form of Questions 3.1, 3.2 and Questions 5.1, 5.2 respectively. Not all questions were answered by all 254 respondents. A

breakdown of the number of answers received to different questions are as follows, which shows a total of 757 answers.

Question Number	Number of responses
Question 1	118
Question 2	115
Question 3	96
Question 3.1	74
Question 3.2	64
Question 4	80
Question 5	61
Question 5.1	55
Question 5.2	38
Question 6	56
Total:	757

Table 1: Breakdown of number of responses received.

Analysis of the data retrieved from the answers to each of the questions along with excerpts from the verbatim responses are presented in the sections to follow.

Question 1 (Current and future research):

Semiconductor and microelectronic research: Describe current or emerging research enabled by semiconductor or microelectronic technology, providing context in terms of recent research activities, and standing questions in the field. NSF/CISE is particularly interested in cross-disciplinary challenges that will drive requirements in a variety of research fields, but that are rooted in, or rely upon, advances in semiconductor and microelectronic technology.

This question was at the core of the RFI and requested information on current and emerging research topics of interest to the community. Correspondingly, questions on classroom education and workforce development were asked in Question 2. One hundred and twenty-three responses to this somewhat open-ended question, were received. The expectation was to assess the topical engagement/interest of the research community at the current time.

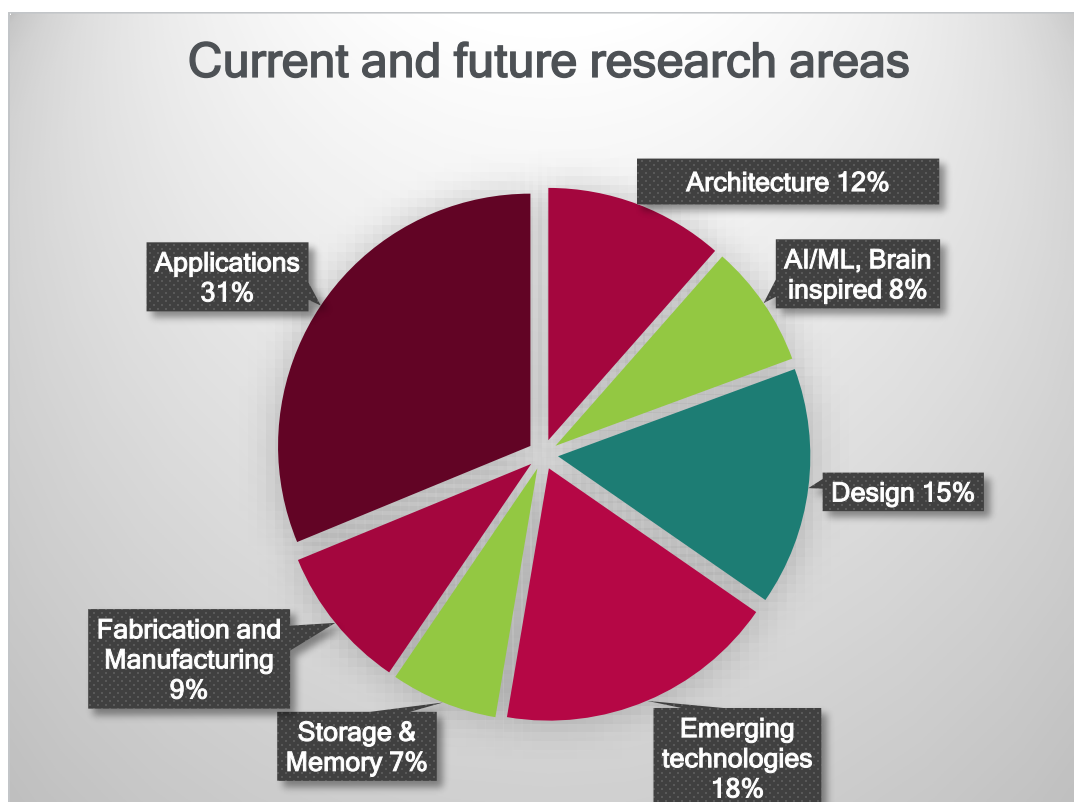


Figure 4: Current and future research categorized into areas of emphasis based on 123 responses to Question 1.

All 123 responses to Question 1 were categorized according to the topical areas addressed by them, which is shown in the pie chart in Figure 4. If a response addressed more than one topical area, then that area was counted in each of the categories it addressed (*this counting strategy was adopted for all responses to Question 1*). The intent here was a methodological categorization rather than listing of fine-grained topics, which are shown in subsequent figures for some of the larger categories. The latter categories included Computer Architecture, (micro- and nano-electronic) Design, Emerging technologies, Storage and

Memory, and Applications. Artificial Intelligence, Machine Learning (i.e., AI/ML), Neuromorphic and Brain-inspired approaches could have been partly included in Architecture or in Design (and to a lesser extent in Applications), but since they constituted a large segment, it is shown as a separate category to emphasize its important role. Approximately 20 percent of the responses (breakdown not shown in the pie chart) in the AI/ML category dealt with neuromorphic/Brain inspired approaches. More fine-grained topics in other categories just mentioned are shown in subsequent charts in this section. Notably a significant segment of the responses (9 percent) included fabrications and manufacturing as part of their current and future needs.




Figure 5: Current and future research; Application sub-categories from a set of 182 cited topics.

The chart in Figure 5 shows sub-categories of application from a set of 182 cited topics in the research work that the researchers are currently pursuing or plans to do so. The most prominent sub-categories of applications include security, along with robotics, vision, imaging, and sensing, as well as in Internet-of-Things at the edge, and autonomous vehicles. Wireless next-G including MIMO antenna, and medical applications also occupy a significant segment of this space. The latter included flexible/wearable electronics, prosthetic/diagnostic tools as well as organ replacement. Relevant hardware aspects of High-Performance Computing (HPC) and Quantum computing were also seen as applications of semiconductor electronics for the purpose of this survey. Applications in sustainability (carbon footprint), cloud computing, data science and analytics, power-grid also featured in some responses, but to a lesser extent.

Selected excerpts from the responses relevant to applications are as follows:

<https://www.nsf.gov/dir/index.jsp?org=CISE>

- Autonomous driving cars and drones are commonly discussed, but medical application such as genetic analysis based customized medicines, and vaccines are also on the verge of breakout.
- Sensing and computing are merging. As a result, semiconductor and microelectronic technology are important to enable such merging.
- Transformative technologies will emerge through expanded university engagement in semiconductors and microelectronics research and developments including: 6G wireless communication network; distributed edge intelligence (intelligent autonomous sensors, edge computing and machine learning); extended reality and immersing teleconferencing (holographic telepresence, multisensory and tactile communications); 3D hyperspectral high-res imaging; 3D indoor positioning and navigation; ubiquitous energy harvesting
- Reliability- and security-aware IC design for safety- and mission-critical applications, ASICs for IoT, Edge, and Autonomy.
- Classical analog computing can be used to supplement digital supercomputing for scientific applications. Quantum analog computing has useful applications, obviously in cryptography, but also in optimization problems, physics simulations etc. There is a growing need to design and fabricated cryogenic computing circuits for such applications.

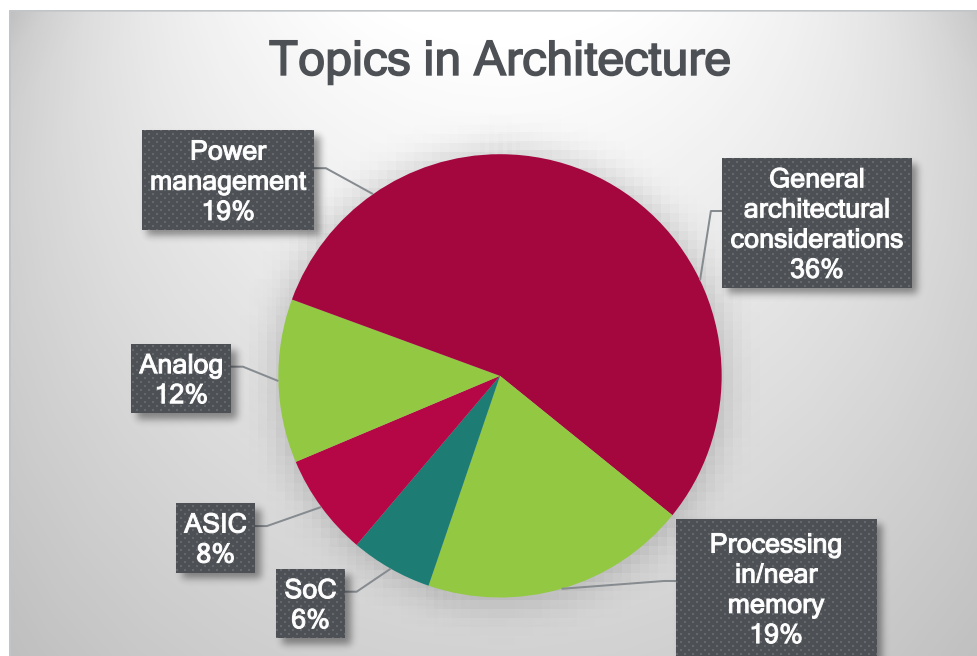


Figure 6: Current and future research; Architecture sub-categories from a set of 67 cited topics.

Figure 6 shows a finer sub-categorization of the Architecture category shown in Figure 4. While this category may have topical overlap (but not overlap in counting) with the (low-level) Design aspects, those are dealt with later in Figure 7. All architectural considerations, such as parallel computing, non-von-Neumann, and others, are lumped together in General architecture considerations category. Among specific techniques processor-in-memory (PIM) and related was by far the winner and is thus separated out in one category. Likewise, power management in the architecture context featured equally strongly and lumped in a separate

category. Among other categories Analog, Systems-on-Chips, ASICs, and domain specific architectures are worthy of separate mentions in terms of the counts.

Selected excerpts from the responses on computer Architecture enabled by microelectronic semiconductor research is included next.

- Research on designing *novel computer architectures* to increase the single-threaded performance of general purpose computers rely on advancements in semiconductor technology and microelectronic research.
- Degrees of integration that permit highly novel *parallel architectures*, allowing for exploration of how structure is required in achieving true machine intelligence.
- Next generation *CMOS in 3D*. Combined device, circuit, fabric architecture, manufacturing aspects.
- *Domain specific architectures* for machine learning are heavily influenced by software/architecture/VLSI/circuit co-design. Technology advancements in available memory devices, compute devices, packaging approaches, and EDA automation tools can have dramatic impact on viability of different machine learning approaches.
- As cloud applications scale their data and instruction footprints performance improvements require an increasing number of low-power switching components, logic elements, memory elements and interconnection network components.
- Pathways to Beyond CMOS: scale up of non-conventional materials, *ultra-low power electronics*, *logic-in memory* architectures using ferroelectrics and multiferroics.
- Computer systems are affected by the emerging *non-volatile memory* and *analog* computing.

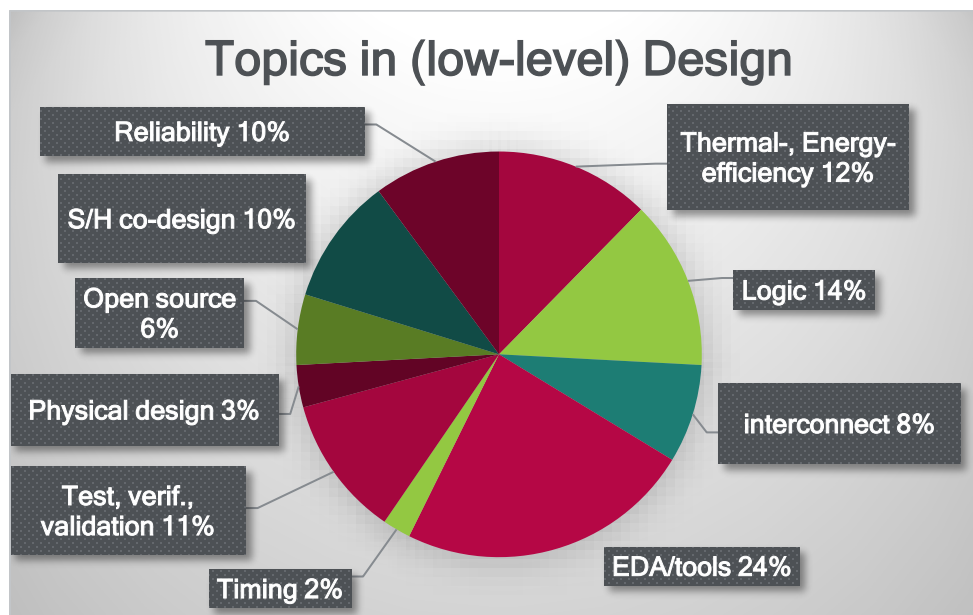


Figure 7: Current and future research; Low-level Design sub-categories from a set of 89 cited topics

Figure 7 shows a finer sub-categorization of the Low-level Design category shown in Figure 4. Perhaps not unexpectedly, EDA/tools take up a major share of this category, followed by research on logic and interconnects, and research involving test, verification and validation, reliability, and thermal/energy efficiency (not counting the topic of power-management included in the architecture category reported in Figure 6). The need for open-source design and physical design feature in this list as well.

Selected excerpts from the responses on topics enabled in the low-level design area are presented next.

- *EDA for emerging technologies* such as optical interconnects on chip, 3D/2.5D stacking, monolithic integration, and others.
- There is a significant effort from the research community to create an *open-source* infrastructure around EDA. This includes projects like OpenROAD, LiveHD, Chisel etc.
- Addressing the unique challenges of ML at the edge will require specialization, *hardware/software co-design*, and integration of new advances in semiconductor and microelectronic technology. The open flows allow teaching and exploring new ideas in a more representative/realistic setup, also ideal for training future industry leaders.
- The cross-disciplinary challenges include (i) co-design and co-validation of hardware, firmware and software at various abstraction levels using an effective combination of simulation-based validation, formal methods, and side-channel analysis; ii) System-on-Chip Trust Validation using Security. iii) Automatic Implementation of Secure Silicon; iv) Post-Silicon Security Validation; and v) Securing On-Chip Communication Architecture for Designing Trustworthy Systems.

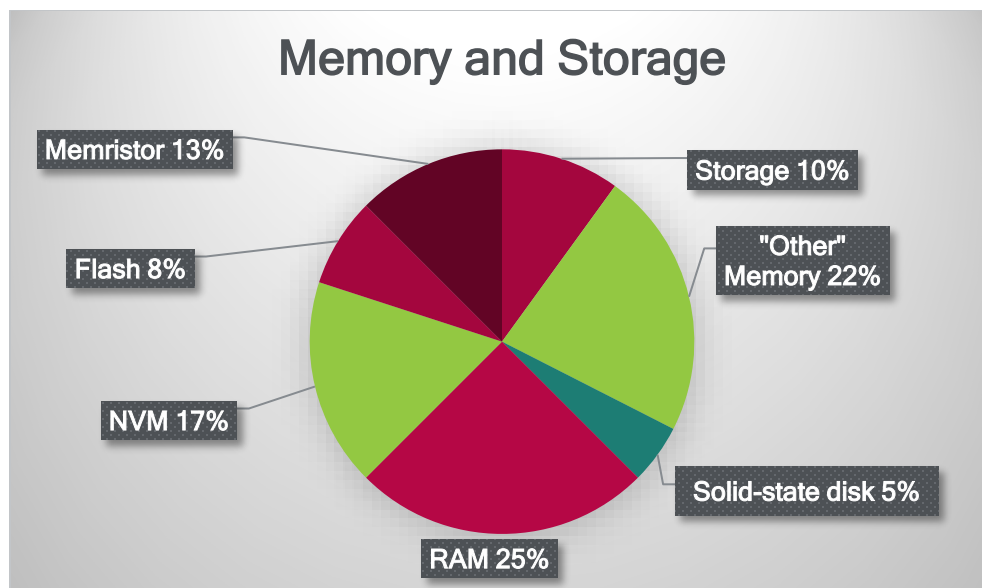


Figure 8: Current and future research; Memory and Storage sub-categories from a set of 40 cited topics

Figure 8 shows a finer sub-categorization of the Memory and Storage category shown in Figure 4. All sorts of random-access memory (e.g., DRAM, resistive RAM, PCRAM etc.) are combined under the category RAM. Explicit mention of memristors is categorized here as well, as are research on non-volatile memory (NVM), and flash memories. All other explicit mention of memory technologies is categorized under the category of “Other” memory. Mention of storage and solid-state drives are separately accounted for. Note again that some of these categorizations may have technical overlap, but the numbers of responses are reported here for each of the respective categories mentioned.

Selected excerpts from the responses on Storage and Memory research are:

- Investigation into memory density to provide the necessary amount of memory required to store the datasets that are being created today. We are quickly running out of storage space.
- New advances in semiconductors and microelectronics will dramatically reshape our approach to computing. It will also help to have many of these new technologies available to a broader audience, such as electrochemical memories, or resistive RAM in an available foundry process.
- The non-volatile memory technologies fabricated using semiconductors also report having lower power consumption and cost (as much as 2-32 times as reported in the literature) compared to traditional system memory. Thus, the use of NVM in computing and storage devices is increasing rapidly for computing systems with large capacity and high bandwidth memory and storage system requirements.
- One of our current research projects enabled by microelectronic technology is solid-state disks (SSDs), in which cache stores frequently accessed data to shorten user-I/O response time. We are focusing on leading-edge techniques to reduce the number of read/write operations in flash memory so that SSD performance and lifetime can be enhanced.
- The lack of highly efficient, large data processing in the sensing devices or in edge computing platforms results in significant energy waste, both in moving data between memory and the CPU in these systems, and in transmitting raw data across the network. To tackle this energy waste, we are exploring the potential to perform processing directly using memory cells of emerging device technologies, such as MRAM and ReRAM

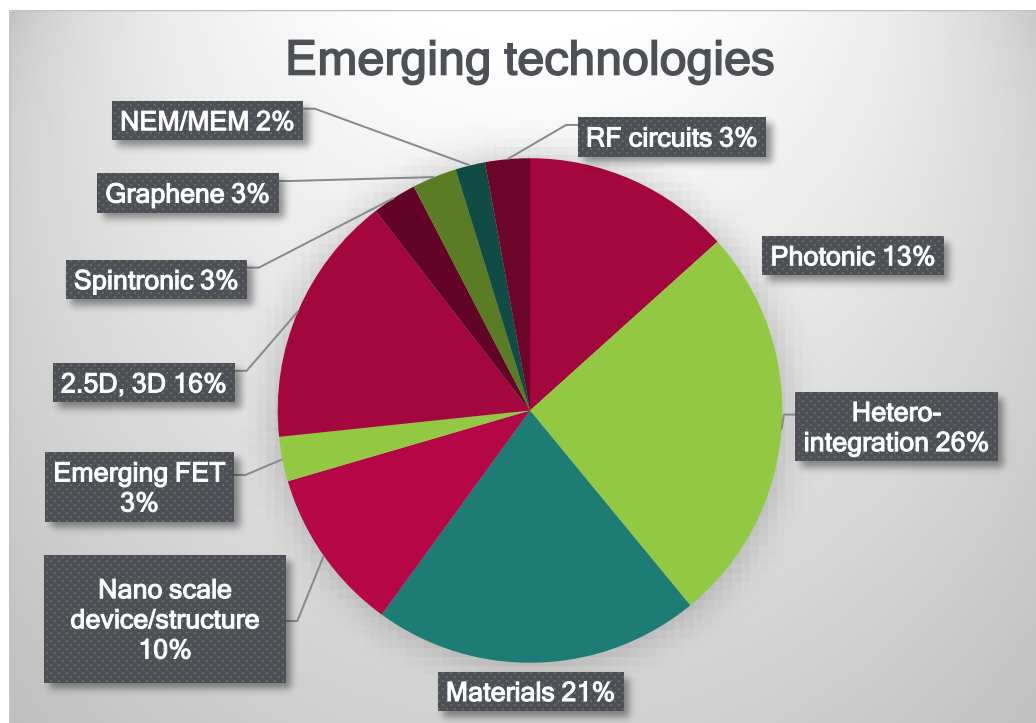


Figure 9: Current and future research; Emerging technologies sub-categories from a set of 105 cited topics

Figure 9 shows a breakdown of the responses received in the finer sub-category of the Memory and Storage shown in Figure 4. This constitutes the largest of the categories after Applications in Figure 4, and itself contains some elements of architecture, materials inspired by nanotechnology research of recent decades, and heterogeneous integration of various technologies. As a single technology, optics/photonics (including silicon photonics) features prominently here. The Materials sub-category includes all forms of semiconducting materials and their alloys e.g., 2D materials, III-IV, and other varieties of semiconducting materials. Miscellaneous Devices/structures at Nanoscale are included in a separate category, whereas a few prominent ones are specifically called out as separate items. The latter categories include e.g., Spintronics, NEMS/MEMS, Graphene, and 3D structures using carbon nanotubes.

Selected excerpts from the responses on research enabled by a variety of Emerging technologies are listed below:

- *Optical-domain* brain-inspired computing and High-performance computing in integrated photonics platforms. Progress in this field requires multidisciplinary and interdisciplinary work in the areas of integrated photonics, high-performance computing, network and simulation, optimization, design automation and tools, fabrication, and computer architecture.
- Current research challenges: exploring compute-in-memory for data intensive applications, design technology circuit co-design of emerging technologies, *heterogeneous integration* for sustaining Moore's law.
- Some of the outstanding questions include how yield/cost characteristics for *mono3D* ICs will evolve in the next 5, 10 years and whether the mono3D technology will be able to build more than two device tiers.
- Results show the promise of semiconducting magnets in *2D materials*-based *spin-orbitronic* devices.
- Remarkable advances in nano-devices and molecular technologies are occurring these days. Microelectromechanical systems (MEMS), *spintronics*, and nanoelectronics could play an important role in future electronic systems.

Question 2 (State of Education):

Semiconductor and microelectronic education: Describe your current educational situation relative to semiconductor, micro- and nano-electronics. What classes does your department/organization teach? What are the challenges moving forward in terms of education (e.g., enrollment or diversity challenges, access to fabrication facilities, detailed simulation data, design kits, IP, EDA tools, data sets, etc.)?

The intent of this question was to get information about the current state of education at the respondent's institution, and what are the major challenges that are faced. This question asked for two different answers: one about the courses that are taught at that institution, and another about the challenges seen relative to those courses.

Because we did not have a pre-defined set of courses for the respondents to choose from, we had several different ways of describing what are, we believe, essentially similar courses. For example, an Analog VLSI course might also be described as an Analog CMOS course, an Analog Circuits course, or an Analog Semiconductor course. Based on the answers, the courses mentioned are seen in Figure 10.

By far the most common courses mentioned were some flavors of Analog VLSI and Digital VLSI courses. Other courses commonly mentioned include some flavor of fabrication-based courses, Computer Architecture or Computer Organization, CAD/EDA, FPGA-based courses, and RF circuits courses. The accumulated data for this portion of the question is:

106 total responses
Analog VLSI: 50
Digital VLSI: 43
Fabrication/cleanroom: 22
Computer architecture/organization: 19
CAD/EDA: 14
FPGA-based: 8
RF circuits/VLSI: 8
Digital design: 7
Photonics: 4
Non-Von Neumann architecture: 4

Other courses receiving just one or two mentions include MEMS, packaging, storage systems, high-level synthesis, microwave, power systems, cybersecurity, quantum computing, and embedded systems.




Figure 10: Courses Mentioned

The second part of this question asks about current issues related to teaching those courses. The most mentioned issue was affordable access to fabrication of semiconductors (33 out of 106 responses or 24 percent). The related access comments included access to electronic design automation (EDA) tools (14 percent) and access to process design kits (PDKs) that contain detailed information about semiconductor process information (12 percent). Some representative comments from this group include:

- Challenges: Fabrication facilities, design kits, EDA tools, and data.
- Students taking coursework only can only fabricate if they take the fab courses and that is limited to a few transistors. There is no foundry access to students taking coursework.
- Challenges: Access to advanced technology nodes & PDKs, access to affordable fabrication facilities, small number of female and minority students.
- A key challenge in instruction is the lack of access to the more advanced technology nodes both in terms of affordable fabrication and process design kits.
- Modern tools are complicated --a challenge to maintain, keep up to date, and for students to navigate. Students are used to instant gratification from software, making it harder to keep them engaged with high-overhead and slow EDA tools.

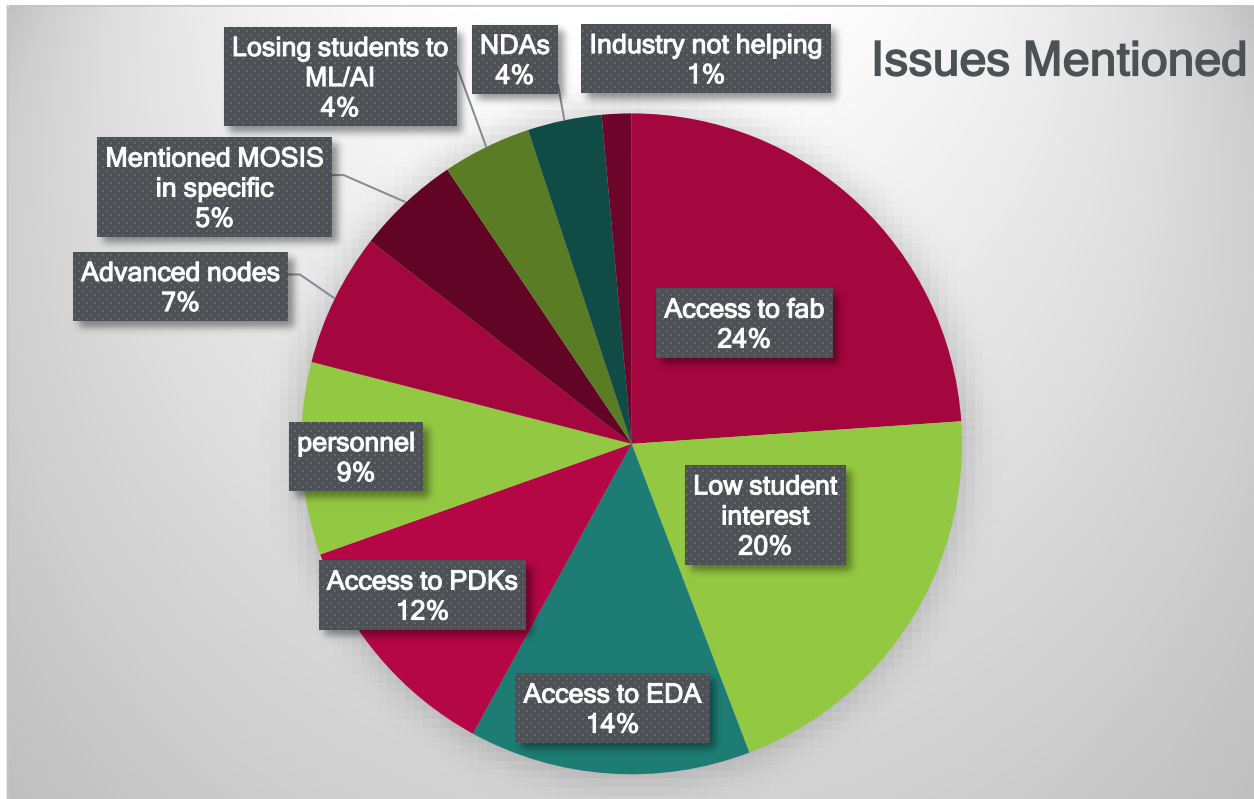


Figure 11: Issues related to delivering the courses mentioned in part 1.

- The tools are available from Cadence and Synopsis but the constant training to keep up to date and offer state-of-the art tools in the classroom is time consuming.
- Access to design kits, IP, EDA tools, data sets etc. has always been a problem.
- The lack of PDK kits for cutting edge technology node in research lab or industry companies, are creating the gap between the technology transferring from research to industry.
- The access to process design kits (PDKs) is problematic. A few educational solutions exist, but the inability to fabricate any of these designs is a disadvantage.
- At the beginning, I thought that old technologies are sufficient for educational purposes. But that is not true anymore since chip design process in 0.35 um and 10 nm technologies are vastly different.
- The challenges are: 1. Teaching for undergraduate and graduate students - tape-out chance such as 130 nm, 180 nm, and 90 nm technologies. 2. Research for graduate students based on traditional CMOS technologies - tape-out chip chance beyond 65 nm technologies.
- Access to advanced technology node PDKs and IPs without encountering very long delays and limitations due to NDAs
- The biggest problem we have is the lack of open-source technology files, and the fact that advanced technologies are not just NDA's, but export controlled.
- EDA tools for current/future technologies continue to be a challenge. In particular, detailed models of next generation devices, including NVMs, are not available to academics.
- Design kits, EDA tools are not a big concern. Primary issue has been access to affordable and quick turnaround fabrication for both education and research.

- Access to modern CAD tools and design kits is actually very good. What is missing is the resources to pay for fabrication in educational settings, particularly at more advanced nodes.

Related to access to fabrication and EDA tools, 5 percent of the total responses mentioned MOSIS by name. Some representative responses related to MOSIS include the following. We should note that these are verbatim comments from respondents to our RFI and do not necessarily reflect any official position of NSF towards the MOSIS program at ISI.

- MOSIS used to offer free fabrication for instructional programs but stopped this program. This is highly unfortunate. For this purpose, we do not need access to the latest technologies and mature platforms would suffice.
- We feel that a significant portion of this is related to the lack of free access to fabricate through MOSIS for educational purposes. All chip fabrication now comes from research funding.
- Some challenges moving forward include access to fabrication programs like the one formerly provided by MOSIS, connecting to the existing curriculum (<state> is pushing all public universities to 120 credit hour degrees), and appropriate tutorials.
- We need to fill the void that MOSIS left and provide greater access to state-of-the-art semiconductor fabrication facilities to all students as we had 30 years ago, which drove a revolution in EDA and microsystems innovation. Google-Skywater-enables investments are valuable, but we need much more, and coordinated at the national/federal level.
- Access to fabrication facilities for education is nearly nonexistent. In the past, I used free MOSIS MWP service for education (for senior design projects). But it is more and more limited every year and the available technologies are just too old.
- Access to fab is very limited especially after MOSIS has curtailed the access to design kits and funding supplements.

Another challenge that was mentioned in 20 percent of the responses was related to low or declining student interest in semiconductor issues. An additional 4 percent of the responses were even more specific in mentioning artificial intelligence and machine learning as being much more interesting to today's students. Some representative responses in this area include:

- We have noticed almost all the students who come to <my university> do not consider working in semiconductors.
- Challenges include fewer students who are interested in pursuing careers in Microelectronics due to the various reasons.
- I see a trend that less and less students are in fundamental science and engineering. Rather, they are going to business, politics, medical, and liberal art.
- We continue to have challenges for enrollment as more students are moving toward computer engineering. In other words, there is a trend away from hardware toward design and software.
- The main challenges in recent years are mainly the low enrollments and the lack of access to fabrication facilities. Lack of funding in EDA and semiconductor research also reduces the interests of faculty in teaching courses in them.
- Enrollment in hardware classes is significantly falling behind compared to CS/software.

- Students interest in VLSI has rapidly declined over the last several years.
- Modern tools are complicated -- a challenge to maintain, keep up to date, and for students to navigate. Students are used to instant gratification from software, making it harder to keep them engaged with high-overhead and slow EDA tools.
- The students are not fully prepared with enough background information to understand the main design concepts and find it hard to link them with some other courses they have (e.g., programming or machine learning).
- One main challenge is that the enrollment is relatively small in the Computer Science department comparing other courses such as machine learning.
- To a significant extent, this shortage of students in the fields of electronic and photonic materials and devices is a result of most students taking up software and software-related pursuits in their graduate programs. If this trend continues, then the center of gravity of semiconductor-based innovation and commercialization will shift even more strongly to countries in the Far East. Focused and sustained investment in experimental infrastructure is the best strategy to counter this trend.
- At the undergraduate level, students are increasingly having the tendency to avoid lower layers of the computing stack (such as circuits and architectures).
- The largest current challenges are the reluctance of faculty to develop and teach design-intensive courses involving physical systems and the reluctance of students to take courses on the foundations of computing systems. Most students want to write apps but don't want to learn about transistors or figure out how architectures must change due to changing device technologies.
- We find that we lose students to machine learning, and that circuit designers may become an endangered species. We need better access to free fabrication technology for educational purposes, notably with simplified bureaucracy, making it easy for an instructor to assign actual circuit design projects.
- A MAJOR problem is that there is diminishing interest in semiconductors and hardware related classes among incoming students. Most students are attracted towards AI/ML/Data Science classes for a variety of reasons.

Another theme in the responses was that lack of personnel to help both with teaching and with tool and lab support was an important challenge. For example:

- Proposed Needed Actions: Invest in human infrastructure (faculty, instructors, technicians, teaching assistants) to support educational and workforce development initiatives.
- Perhaps the most critical two pieces would be lab staff (typically more in number and skill level than is typical in an educational program) and access to foundry CMOS and supporting services like packaging.
- The tools are available from Cadence and Synopsis but the constant training to keep up to date and offer state-of-the art tools in the classroom is time consuming.
- The facility is fortunate to have a large team of full-time technical staff members responsible for equipment maintenance, user instruction, process support, and user support. CNF staff serve one function—supporting the user program.

Finally, although we did not specifically ask for the respondents to provide solutions to these issues, some respondents did offer some suggestions. For example:

- Proposed Needed Actions:
 - Create a nationwide university/industry/government program to develop and test educational content for broad dissemination and to support outreach initiatives designed to expose and attract diverse high-school students and first-year college students to microelectronics disciplines.
 - Invest and support the maintenance of educational facilities and programs at universities designed to foster hands-on, project-based, design-oriented, multidisciplinary research and educational experiences for undergraduates.
- I would aspire to maintain a nano-maker facility where these tools are maintained, curated, have many tutorial examples, etc. Like a regular maker-space, the key is in the expertise and availability of trained staff, relatively turn-key and user-friendly systems (software and hardware), lots of example and a culture of doing that includes visible publication of results and activities (Instagram, TikTok, etc.).
- We should also develop new courses with a cross-cutting learning focus where the traditional boundaries of device, circuit, and architecture-oriented courses are blurred, and the students are exposed to the entire design stack to holistically explore optimal computing paradigms.
- The absence of a centralized resource to allow IC design computing (e.g., the Canadian model) where centrally maintained computing resources can be efficiently provided to universities at-cost and enable increasingly strapped departments from having to maintain Linux workstations/servers, and deal with a range of CAD tools. Centralizing this will be a huge efficiency multiplier and provide economies of scale.
- Students sometime lost interests in learning signal processing, semiconductor devices, which are mostly on book and detached from reality. ECE curriculum need to be updated to reflect more interesting developments. Also, we need to provide students more opportunities to experiment devices so that they are more engaged.
- Google-Skywater-eFabless investments are valuable, but we need much more, and coordinated at the national/federal level.
- We need to excite students with hands-on labs, quick turn fabrication facilities, cheaper access to chip design tools and simulators. Increased funding for research in the space of semiconductors will enable these teaching innovations as well.

Question 3 (Missing facilities):

Infrastructure and capabilities needed for research: Describe any limitations or absence of existing semiconductor and microelectronic facilities, or capabilities and services and/or specific technical and capacity advances needed in this area (e.g., access to fabrication facilities, detailed simulation data, design kits, intellectual property, access to advanced technologies, electronic-design-automation tools, data sets, etc.) that would be required or desired to advance this research.

The intent of this question was to understand what capabilities and/or facilities were missing in the current research environment.

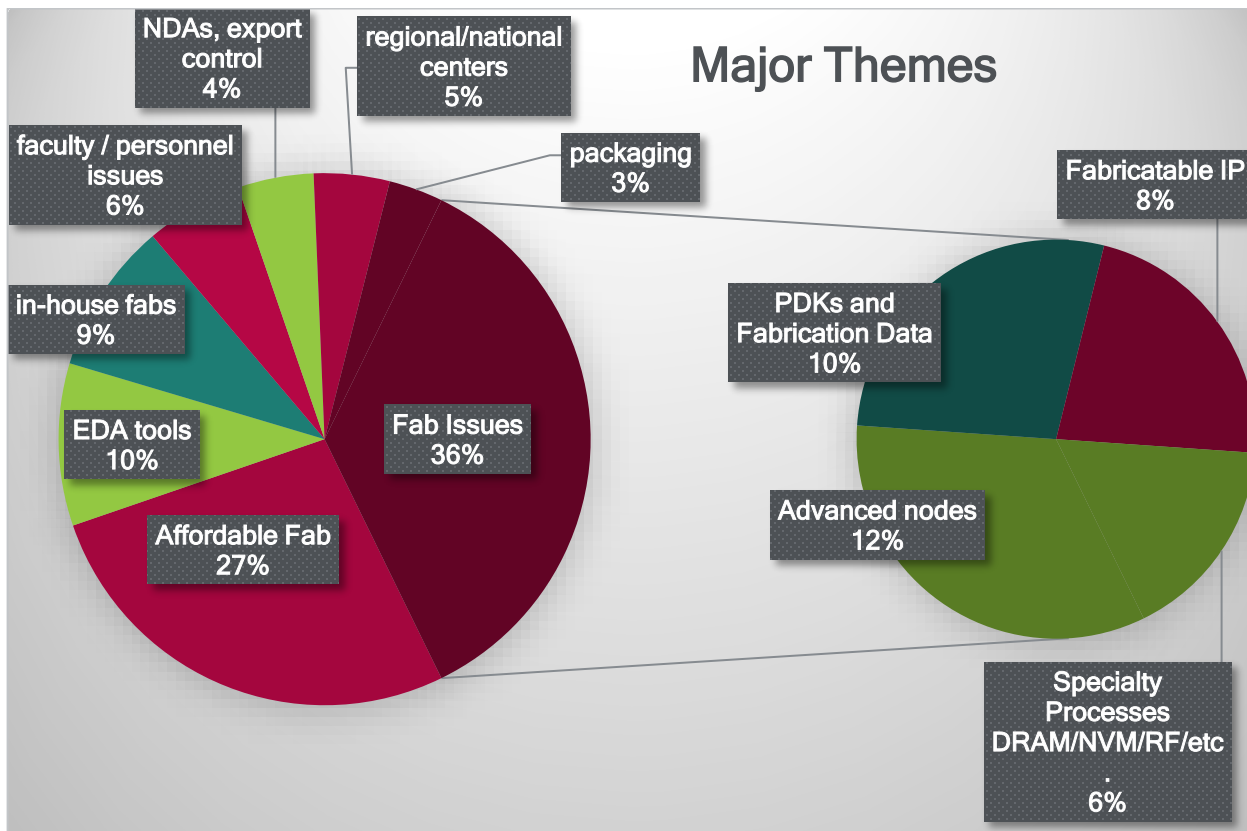


Figure 12: Major themes mentioned in relation to what capabilities and/or facilities were missing in the current research environment

The accumulated data for this portion of the question is:

91 total responses:

- Miscellaneous fabrication issues (PDKs, IP, advanced nodes, etc.): 54
- Affordable fabrication: 41
- EDA tools: 15
- In-house fab facilities: 14
- Faculty and other personnel issues: 9
- NDAs and export control issues: 7
- Regional/national centers: 7
- Packaging: 5

Some representative comments on miscellaneous fabrication issues include:

<https://www.nsf.gov/dir/index.jsp?org=CISE>

- In order to be competitive in venues like ISSCC we need to fabricate at 5-7nm CMOS node, and this is virtually impossible with NSF funding as the costs are too high.
- We lack the facilities or access to the proper tools and facilities to fabricate and evaluate emerging nano-CMOS technologies.
- Integrated Circuits require access to state-of-the-art deep sub-micron processes offered nowadays mostly by foundries in Asia.
- Access to advanced nodes for fabricated parts has become increasingly difficult, almost impossible.
- Access to sub-20nm CMOS IC technologies limit ability to prove out analog, RF, and mixed signal design concepts in state-of-the-art processes. Infrastructure support for large dies of 28nm-130nm CMOS technologies stopped at planar upper metallization level in preparation for “more than Moore” novel devices to be added on top of the CMOS and connected directly to the CMOS interconnect stack.
- A large hurdle for my research is the lack of access to advanced process technologies, particularly those for emerging memory technologies.
- Access to more advanced technologies, or at least parts of it, as well as IP libraries.
- Lack of access to fabrication facilities, especially those for advanced nodes, is a big drawback for advanced research in EDA.
- It would be nice to have full access to advanced process nodes, but I understand that there are cost barriers to that due to the high cost of fabrication. There are though partial solutions that would cost very little but are hampered instead by NDA/IP protection issues. PDKs, IP blocks, simulation models would cost very little in terms of dissemination but are not available because of the desire for secrecy of the (few) major semiconductor companies.
- Access to advanced technologies, electronic-design-automation tools, data sets would be great, but industry is understandably protective of their IP
- Mainly access to state-of-the-art process technologies, both in terms of PDKs for EDA, and access to affordable fabrication.
- The existing US-based foundries lack the advanced fabrication capabilities.
- For academic research to have high impact and be viable for influencing commercial state-of-the-art companies, universities need access to prototype designs in state-of-the-art technologies. For standard digital circuits, this could just be access to state-of-the-art commercial EDA tools and design kits.
- <my university> lacks cost-effective manufacturing opportunities, especially for modern technology nodes.
- Intellectual property is a big challenge. In particular, the latest circuit and architecture models (14nm or smaller), standard cells, and validated models.
- There is very limited access to fabrication facilities for universities. This is especially true for new technologies (e.g., 3D integration or designing with silicon photonics).

Related to these issues, there were a significant number of responses that specifically mentioned cost or affordability in their issues related to fabrication. For example:

- Access to affordable fabrication facilities, design automation tools, and interdisciplinary training for students.
- Access to commercial fab services at reasonable costs is problematic.
- Access to CMOS in form factors that allow post processing with reasonable cost

- We are fortunate to have outstanding fabrication facilities at our institution, but these facilities are also very expensive, and thus chip/wafer/device processing expenses are high.
- It would be nice to have full access to advanced process nodes, but I understand that there are cost barriers to that due to the high cost of fabrication.
- Some of this is due to the high cost of fabrication -- if support were provided to startups to offset fabrication, IP, and EDA this would enable a huge growth in technology in this space.
- Currently, the design and tape-out experience still cost a lot, even with support from grants.
- Access to microelectronics facilities that are priced so that we can teach classes and conduct research.
- We have tools but nowhere to fabricate cheaply and quickly.
- Getting more MPW chip fabrication opportunities for school will really be helpful. Currently, it is expensive and has lots of restrictions.
- Fabrication facilities is one of the most significant. Ever since MOSIS changed their policies for chip fabrication for students, universities have been scrambling for meaningful ways to get students meaningful opportunities to fabricate their designs.
- In chip tape-out projects so far, we used MOSIS and MUSE, a multi-project service. The costs for a tape-out are still quite expensive, and easily top \$40K for a newer technology node. That level of funding is only available in the context of larger projects or industry sponsored projects.
- We have no access other than the Google MPW which is not a permanent solution. We do not have access to other commercial PDKs. I was unable to negotiate an NDA with CMP for 28nm.
- I am on the design side of the business. So, for me access to fabrication (not in house) at reasonable cost (or free!) would be great.
- The main barrier is lack of funding and NOT availability of technology.
- The main things that are missing are: Cheap access to full versions of design tools, cheap access to fabrication, well tested and supported open-source hardware designs to build upon.

Access and affordability of EDA tools was another theme. For example:

- We lack the facilities or access to the proper tools and facilities to fabricate and evaluate emerging nano-CMOS technologies.
- While in the future, I would like to fabricate a test chip for end-to-end performance and energy numbers, the state of EDA tools is such that, I do not have the person-hours to fabricate a test chip while the main focus of my research is at a higher level.
- The main challenges are as follows: Tool chains that are curated and maintained (fabrication, CAD, simulation)
- Honestly, the impact of open-source EDA is so far behind commercial EDA that there is no way a startup can use open-source EDA to develop FinFET circuits. (Our startup has taped out 12nm chips and are working on a 5nm design. I don't see openEDA supporting this any time soon). The one place this has been successful is in the processor space where RISC-V has been able to find a commercial foothold. We need to enable many more research technologies to follow a similar path.
- Lack of open-source drivers and tool suites for the most advanced parts is limiting.

- There lacks an open-source toolchain and design flow like the LLVM project for compilers.
- Open EDA tools and library.
- Difficulty in accessing EDA tools (tools are expensive and come with complex legal haggling; we are able to obtain access to some logic libraries after legal hoops, but not DRAM libraries)
- Access to foundries is important, but the main challenge to evolve the EDA flow is to have open flows where designers can try/fix a part of a whole flow. This is difficult with close source tools.

Other responses that mention different issues include:

- Proposed needed actions:
 - Make large, sustained investments in updating fabrication and metrology equipment in university research facilities with emphasis on outfitting a few with flexible, production-class but research-oriented, 200 mm tools.
 - Establish programs to provide sustained support for operational costs of the national university tool base.
 - Invest in a nationwide program to underwrite the creation of new faculty positions, to provide flexible career-initiation grants to junior faculty, and to engage industry researchers in university activities.
 - Foster the creation of regional networks to create and manage research, educational programs, startup support, outreach, and internship programs with a regional dimension and that are designed to facilitate the involvement of educational institutions previously on the sidelines of the national microelectronics enterprise.
- Currently, the barriers to multi-institution NDAs are such that is extremely difficult for US institution to collaborate. National Laboratories could potentially play a role in creating a framework to access technologies for consortia of research institutions, following the example of European institutions like CERN.
- Increasingly companies produce advanced, programmable ASICs (e.g., P4-switches) with highly restrictive access to tools and specifications that limit their exploitation in academic settings.
- There are no IC fab processes run at the University. We have to contract out to MPW vendors for fabrication services and design kits. The IP and liability issues are problematic, when working between University lawyers and US companies.
- We also do not have faculty with experience so usually the entire design process is under one faculty member.
- Relying on PhD students for fabrication is no longer sufficient to meet the nation's chip manufacturing needs and hasn't been for a long time.
- The current access to technology is severely constraint from a legal perspective. At my former school (a public school), it was effectively impossible to perform a chip tape-out because the school was unwilling to sign an NDA with the foundry. The point of dispute, in this case, was that the NDA stipulates those legal disagreements have to be evaluated in an out-of-state court, which is a non-starter for a public, state-sponsored school.
- We need domestic fabrication facilities, including clean rooms, staff, and training of next generation of students.

- Missing is national/regional/local infrastructure for teaching and exploring computational science (in contrast to computer engineering or computer science) on emerging technologies. The gaps span hardware, software, documentation, expert personnel.
- if NSF can create a facility where academic users can experience the ASIC fabless model from all sides it would significantly improve their preparedness when they are ready to take jobs.

Question 3.1 (Specific technologies):

If your answers involve specific semiconductor or microelectronic technologies, please let us know what those are, and be as specific as possible (e.g., semiconductor technology, feature size, die size, processing requirements, packaging, specific foundries, turn-around time, expected annual cost, etc.). If your needs are more for leading-edge semiconductor fabs, or more for CMOS+X (beyond-Moore's-Law technologies), list the specific ("X") technologies in which you may be interested.

This question was trying to understand the specifics of exactly what technologies our respondents were interested in. We received 66 responses to this portion of the question. Figure 13 shows the tally of specific CMOS process nodes mentioned either directly, or in a range (e.g., "12nm and below"). There seems to be a bit of a cluster from 28nm to 7nm, with 7nm being the node most frequently named, although there were measurable responses at larger nodes up to 130nm and larger.




Figure 13: CMOS process nodes mentioned

Some representative comments related to specific CMOS process nodes include:

- Some subsidized program to allow fabrication at 7nm CMOS using standard cell ASIC approach would be a great start.
- But for most of the researchers in academia, access to design tools including EDA and TCAD tools, computational materials tools, and similarly newer technologies at 10 nm or below are lacking.
- I need access to sub-10 nm CMOS, MRAM, and a variety of exotic technologies currently in development.
- Intel 14nm FinFET technology, TSMC or GF 28nm technology, 45nm - 130nm mature CMOS technologies
- CMOS at the 180 nm or 130 nm node for prototyping and debugging, planarized after fabrication (cheaper and more tolerant development vehicle) CMOS at the 28 nm node for showing performance
- Most of our research happens in 65nm CMOS ... Most advanced products are fabricated in 7nm. We should close that gap but in an affordable way.
- We are interested in accessing advanced nodes for MPW runs, including 14nm and smaller.
- For the fabrication facilities, we are mainly talking about 7nm and below.
- We have built a complete 6mmx6mm 12nm (TSMC N12) design and are working on a 5nm (TSMC N5) chip. Such FinFET nodes are important for both research and commercialization.
- Access to FinFET SOI technologies (22FDX as an example).
- MOSIS had 0.25-micron shuttle runs which are out of date. Skywater has 130 nm and will come out with 90 nm soon which is fine for proof-of-concept but still not state-of-the-art
- I would be very interest in having access to a relatively recent CMOS technology node, e.g., 65 nm CMOS standard cells, the easy access to a design kit for such a technology node. ... I would not be willing to pay extra for newer technology nodes: it already takes tremendous effort and patience to train students in newer semiconductor technology, and they will directly benefit industry when they graduate.
- Lack of access to modern fabrication (beyond 22nm) technologies is a significant drawback. I was able to tape out at 65 nm TSMC node via Europractice.
- 1. Teaching for undergraduate and graduate students - tape-out chance such as 130 nm, 180 nm, and 90 mm technologies. 2. Research for graduate students based on traditional CMOS technologies - tape-out chip chance beyond 65 nm technologies.
- For academic research to move the needle, it's often quite important for it to be in close to state-of-the-art technologies - e.g., sub 16nm FinFET at TSMC or Samsung with flip chip packaging for good power delivery.
- 3 nm devices and beyond with TSV and other means of having 3D IC
- Primarily sub 10-nm CMOS.
- In research, we are using 28nm, 22nm CMOS processes, as well as the 130nm process. In education, we are using ASAP7 (educational) process.
- Obviously, newer processes are necessary for publishing (at least 28nm?) but older ones may be sufficient for education (45nm?).

Technologies Mentioned




Figure 14: Other technologies mentioned in question 3.1

In addition to standard CMOS processes, other technologies were mentioned with emerging non-volatile memory and CMOS+X processing being the most frequently mentioned.

Representative comments include:

- CMOS+X; X-> NVRAM.
- CMOS + X, e.g., 3D or adding to the back end is a very expensive research enterprise. My most recent experience was at 28 nm. It cost \$250k to get onto an MPW, then \$100k to get ten wafers. \$350k before I even started the +x
- Semiconductor sensor fabrication facilities with processes for Silicon, Germanium, GaAs, GaN, and compound semiconductors like CdZnTe, CdTe - CMOS Image Sensor technologies with customization options (e.g., TowerJazz, LFoundry) ... CMOS technologies for cryogenic operation down to 4k (e.g., FDSOI) - Heterogeneous integration (e.g., Silicon Photonics) - 3D vertical integration and advanced packaging - BEOL integration of devices for in-memory computing
- Transistor-level research on graphene nanoribbons, carbon nanotubes, low-temperature / BEOL compatible devices, and beyond-CMOS devices high-performance low-power devices needs to be stepped up. Ditto on memory elements that can be viable disruptive alternative to SRAM and DRAM, e.g., SOT-MRAM, Fe-RAM, etc.
- Superconductive circuit technologies such as SFQ are also of increasing importance.
- Ideally, we would like the ability to fabricate MRAM and ReRAM arrays where we can assemble individual memory cells as we please, to connect them and assert access

voltages with our custom designs (e.g., to turn on multiple columns of cells simultaneously to perform logic operations).

- CMOS at the 28 nm node with fabrication stopped at Metal 2 or 3 to allow integration or post CMOS devices (like RRAM) closer to the transistors CMOS (any node) + RRAM where we can really mess with the internals of the cells and how they are connected
- Beyond Moore-law technology = RRAM, STT-MRAM, SOT-MRAM, FeFET, Access to custom eDRAM and 3DI FLASH and UTBB-SOI FET technology, Chemical mechanical polishing technology and help with backend integration.
- CMOS+X (Beyond CMOS): Logic in memory with ferroelectrics and multiferroics; spin-based logic; ultralow power electronics.
- In my opinion the CMOS+X research is not something that should be done at universities since it is mostly "backward looking" and uses crippled toolboxes that industry decides how much to release, and the students are doomed to dream with their hands tied from the get-go.
- DRAM. As noted, even an older node would suffice, if it were compatible with state-of-the-art packaging.
- Our group is focusing on the new materials and new physics for resistive switching memory and self-selected devices i.e., oxide-based self-selected memristor in multilayer stacks structure,
- Access to affordable RRAM/MRAM devices; DRAM technology; Quantum devices
- Our specific interests that stand to benefit from NSF-led leadership in semiconductor technology initiatives include, hetero-material epitaxy, hetero-material device integration, heterogenous systems combining conventional CMOS circuitry with novel materials and devices to develop highly innovative sensors (e.g., chip-scale micro-channel plate-based infrared imagers, super-sensitive narrowband bolometers, hyperspectral sensors etc.)
- access to technology such as resistive random-access memory in a foundry process would be extremely helpful.
- 1. Packaging of CMOS+X is currently a custom experimentation by each PI and is not available at reasonable prices. 2. Quantum technology requires a comprehensive process design kit for cryoelectronic temperatures that is not being shared by the leading industry players to the research community (domestic/international).
- We are interested in CMOS+X relating to memory technologies.
- The main thing is access to memory IPs, including compiled SRAMs, DRAM controllers, embedded non-volatile memories (RRAM, MRAM, PCM).
- Our need is for detailed models and not fabs: 14nm CMOS or better and NVM technologies (STTRAM, PCM, ReRAM, SOTT ram) with hospice details as well as reliability data.
- I am interested in nanomagnetic and spintronic technologies. They are intrinsically non-volatile and therefore lend themselves to more powerful non-von-Neumann architectures.
- Thin film facilities are extremely ad-hoc and challenging to maintain and are the core of many new processes. Processes for recrystallized Si, organic semiconductors, IGZO, and related technologies would be especially important to support.
- My focus is on more-than-Moore and advanced packaging for bioelectronics and power electronics applications. I do CMOS+magnetic, CMOS+nanotubes, CMOS+nanofluidics, CMOS+ piezoelectrics, CMOS+optoelectronics, thinned CMOS for implantables, integration of CMOS with flexible packages for implantable.

- Currently there are no facilities that academics in the US can use to test DRAM or Flash components. This would be valuable!

Question 3.2 (Overseas activities):

If you have experience with international collaborations specifically involving semiconductor or microelectronic technologies, please let us know what those are, and how they are different or unique from domestic partnerships or facilities. What are the challenges and/or benefits of international collaborations in this space?

This question was designed to see if there were other models of support in other countries that could serve as inspiration for domestic efforts. We received 50 responses to this portion of the question. Many responses mentioned Europe in general, or specific European countries. Some called out specific European programs such as those at IMEC, CEA-LETI, and ST-Micro.

Other non-European countries mentioned more than once include Canada, China, and Japan. Countries that received a single mention include Singapore, Korea, Israel, Mexico, Russia, and India.




Figure 15: Countries and specific programs mentioned in Question 3.1

Some representative comments from this portion of the question include:

- We collaborate with Canadian researchers a lot. Strangely, NSF does not seem to collaborate with NSERC at any significant level. It seems obvious that there is much

to gain for both USA and Canada if NSERC and NSF have an extensive joint funding program.

- We have had collaborations in the UK, Switzerland, and other European nations. There are three types of difficulties in collaborations: One is the time difference for long term collaborations; Two is the funding and agreements needed seem to be different depending on the country and the institutions; Three are that travel budgets for these collaborations tend to increase the overall cost of the proposals.
- Generally, in Europe, science and engineering projects have been long term and sustainable. For example, many of the software used for condensed matter theory have their roots in Europe or UK. In addition, the budgets largely recognize the cost of developing fundamental capabilities and sustaining them for a longer term.
- CERN has been able to create a framework for accessing technologies for large collaborations under a common NDA. This allows participants to share microelectronics designs and collaborate.
- IMEC, labs and fab-like advanced fabrication facilities. Benefits, novel materials, and process exploration not possible on production lines, shared costs, etc.
- I work closely in the semiconductor area with academic institutions in both Israel and Singapore. The primary difference I've noticed is that the local semiconductor companies tend to be more supportive, almost paternalistic, in that these companies need well trained employees and recognize everyone is on the same boat.
- I have interacted with IMEC, and they have an interesting model that is different from either academia or industry and help form a bridge.
- IMEC- Belgium, CEA-LETI uniqueness: Measurement data from specialty technology Challenges: IP issues, physical presence required for the student
- I had collaborations with China but stopped a few years ago because the tensions started rising between US and China.
- Collaborations are fed by funding. There are plenty of opportunities for transatlantic, transpacific, or trans American collaborations, but unless there are funding programs that specifically support such collaborations (in a substantial way), they do not happen.
- The US is really lagging behind other regions such as Europe and Canada in terms of support for academia. Canada has CMC, Europe has CMP/Europractice, US has lost even the meager support from MOSIS.
- There are several holes in the semiconductor space in the US, but packaging and test is possibly the largest issue. We have been required to pay IMEC for substrate design, packaging, and test.
- I have worked with CMP in Grenoble France to manufacture prototype designs. There were strong restrictions on the internet access of the design machine. The benefits were access to additional technologies and lower prices.
- Closely integrated group of countries, such as the European Union (EU), have many such programs because of their shared financial and strategic goals. Even in this restrictive scenario, international collaborations between entities such as US universities and organizations, like IMEC, can be extremely beneficial.
- The Cornell Nanoscale Science and Technology Facility (CNF) is an open-access, nanofabrication facility with global scope and unmatched tools and capabilities. ... The CNF does not restrict access to the facility based on citizenship status. However, access to controlled materials and technology is only permitted in accordance with current U.S. export control laws and regulations.

- In the educational space, I find it is much easier to convince international collaborators of the importance of VLSI/EDA/CMOS+X education.
- We had to go to China to have our RRAM-CMOS chips fabricated. It would have been nice to have these capabilities available here, on US soil.
- I worked at TSMC <in this capacity>. I have many contacts there.
- I have experience working with Semiconductor researchers from Univ of British Columbia and IMEC, Leuven. both these places have strong supplementary support of EDA tools and discounted price for IC fabrication for advanced technology nodes.
- ITAR restrictions are a concern with international collaborations.
- We work closely with IMEC, giving us access to leading edge fabrication research.
- We fabricate research chips through TSMC. The main challenge is the limited area we get through the donation and export controls.
- TSMC has been one of the most important partners for these efforts because they have state-of-the-art manufacturing facilities; but they have limits to their horizon for investment and have become focused of late on advanced CMOS nodes and packaging at the expense of more-than-Moore activities. I have other collaborators in Europe that benefit from IMEC, which has a particularly effective model for more-than-Moore efforts.

Question 4 (New research enabled):

New research agenda: Please describe what new research would be enabled if you had access to the leading-edge infrastructure described in Question 3.

Out of a total of 254 respondents only 80 respondents addressed this question, which was prospective, seeking future potential impact of leading-edge infrastructure was sought.

Responses were naturally categorized into four broad categories: beginning from research using conventional as well as emerging technologies, applications, education to impact on research. The breakdown shows largest interest in research with conventional (45 percent) technologies, followed by emerging technologies (25 percent) and applications (25 percent), and educational endeavors (5 percent).

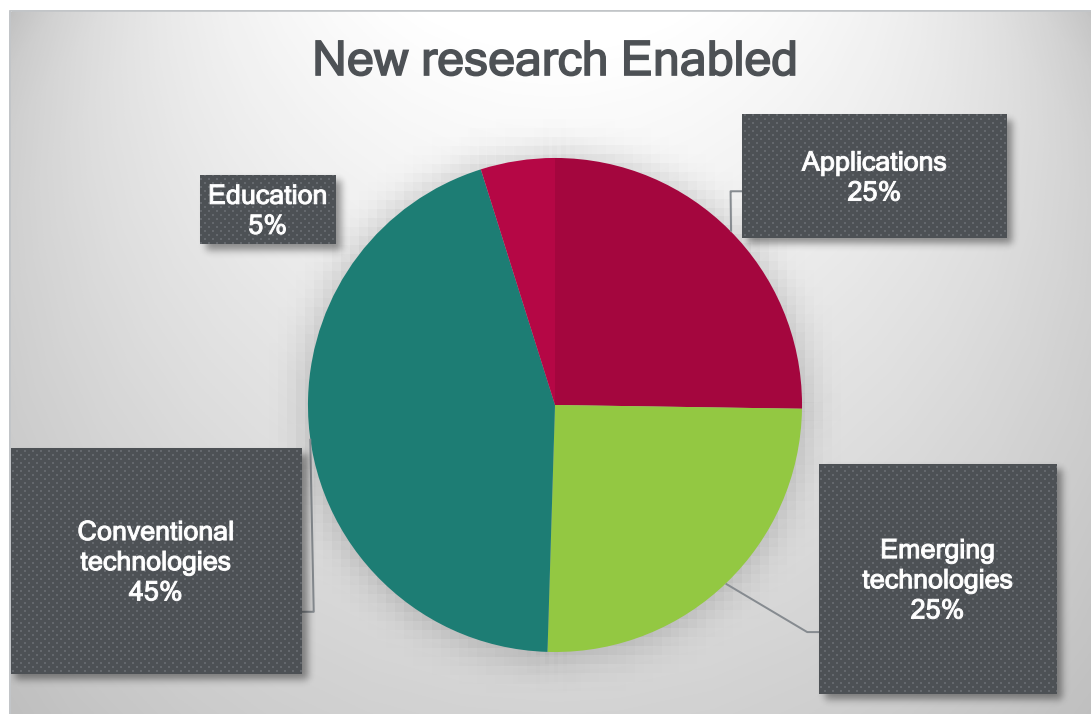


Figure 16: New research enabled by access to facilities

The responses in the first three categories can each be further broken down into specific subareas. The following pie chart suggests the following breakdown, among the 46 responses dealing with research using conventional technologies with highest emphasis on prototyping and validation followed by Low power, memory, security and purely design considerations.

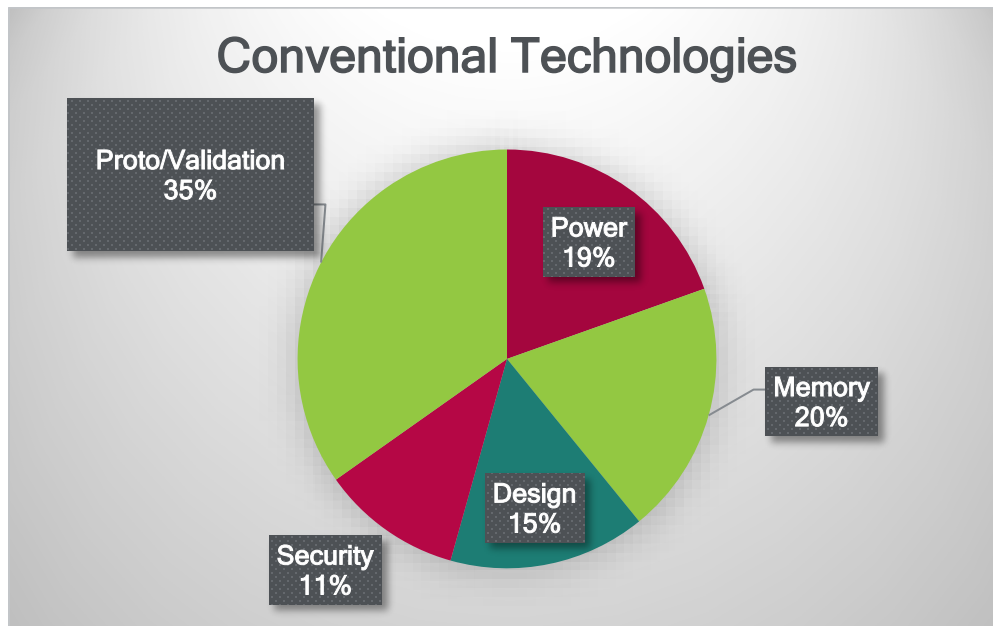


Figure 17: Further breakdown of new work enabled - Conventional technologies

The pie chart below suggests the following breakdown, among the 26 responses dealing with research using emerging technologies with highest emphasis on heterointegration in CMOS+X technologies, followed by potential use in quantum and neuromorphic technologies, and sub-nanometer design and fabrication.

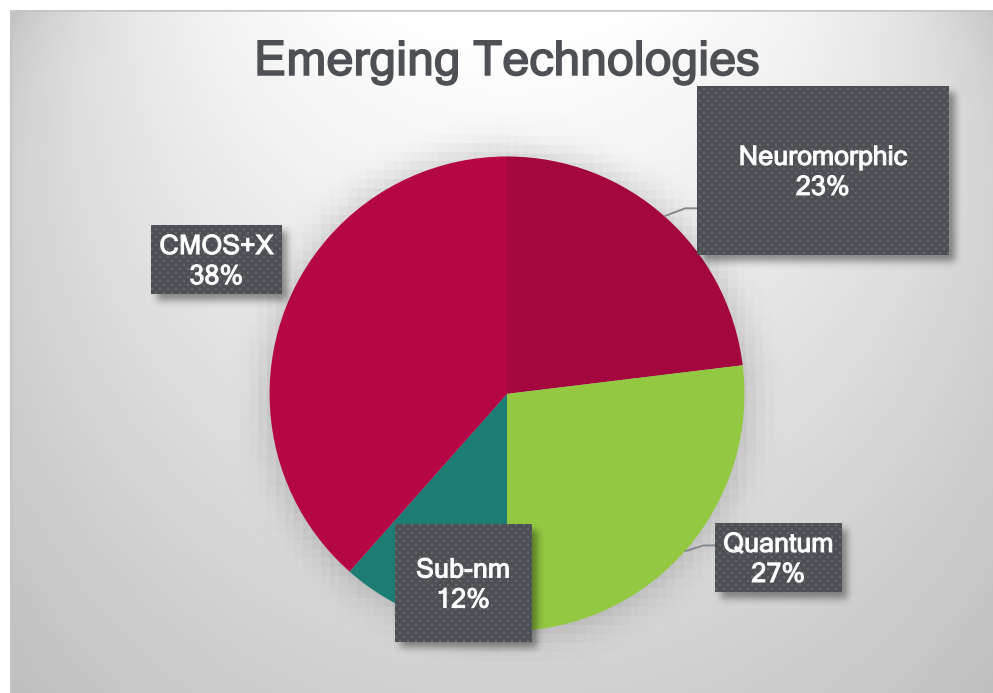


Figure 18: Further breakdown of new work enabled - Emerging Technologies

The following pie chart shows the breakdown, among the 26 responses dealing application areas with highest emphasis on research potential in hardware for Artificial Intelligence and

Machine Learning followed by applications in next generation communications technologies, e-Health and hardware for Internet of Things.




Figure 19: Further breakdown of new work enabled - applications

Slightly edited versions of some of the responses received for Question 4 on research enabled by requested infrastructure are as follows:

- A strong, open industrial-academic partnership will be a key asset for the nation to define and further a research agenda. Key areas include quantum computing, internet-of-everything, addressing climate change, technologies for space exploration, health and biomedical engineering, artificial intelligence ... All these application areas require advances in semiconductor technologies from cryogenic operation, rad-hard operation, to ultra-low power, ultra-reliable, novel computing paradigms for AI (probabilistic, approximate, analog computing,) ... as well as the associated circuit and architecture design research.
- Computational science (including data-intensive science) broadly construed would be poised for a major expansion and could much more effectively plan for the technology future.
- Investigate new methods across device/architecture/circuit/algorithm levels for mitigating non-ideal synaptic characteristics in brain-like computing systems. The results will bridge the gap between high-level accurate algorithms and low-level non-ideal devices, and therefore support memristor-based computer to foster breakthroughs and provide technology reserves for the large-scale applications and commercialization.
- There are entirely new opportunities possible in wearable and implantable medical devices enable by CMOS++ development and advanced packaging.
- More experiments in CMOS+X to address the research problems.
- CMOS+, ranging, sensing, LIDAR, optical sensors, spectral detection - a range of IoT and biomedical sensing would be unlocked.

- Provided expanded access to advanced nodes (such as 14nm and beyond) ... demonstration of advanced memory and device technologies - fully integrated with CMOS - for the emerging and expanding areas of neuromorphic computing and AI. These areas require custom hardware solutions.
- New RF and power electronics integrated with Silicon CMOS. Photonics in the short wavelength integrated with Silicon. CMOS integrated RF and MEMs and new forms of memories. Semiconductor/superconductor electronics.
- Given the limitations of fabricating PIM architectures, evaluation is limited to simulation or FPGA emulation. Access to fabrication would allow us to prototype and evaluate real hardware.
- We would have greater ability to work with circuit design collaborators using in-house fabrication. The ability to have multiple materials systems integrated into the fab would provide greater opportunity to explore circuits for different types of applications (e.g., logic, analog, power). Integrated testing facilities.
- In the mature area of test, academics can still contribute leapfrog ideas, but need industrial interest to validate them in volume production.
- Research programs that support wide disciplines in building full-stack research infrastructure for computing and security will very likely provide opportunities to build and advance those infrastructures.
- Projects that bring physics simulation, hardware systems, design automation, programming infrastructure, and HPC together will advance the infrastructures but also could improve the student enrollment.
- By providing semiconductor technology to a broader range of researchers and trained students, we can ... build a transparent, trustworthy computing platform where every layer of the abstraction is understood, analyzed, verified. I would assume that such a platform may crystallize around an existing infrastructure (such as RISC-V). A second effect of broad access to semiconductor technology is a renaissance in design automation. New and better tools, and entirely new ways of designing and implementing chips become possible.

Question 5 (Beyond facility access):

Facility access: How might existing facilities to which you may have access be augmented to address any further needs? Please note tools or processes that could be housed in an open-access facility.

With this question we sought to understand specific augmentations and enhancements that would improve the state of semiconductor research and education in the US. Perhaps because some of these were already mentioned in other responses, we received only 50 responses to this portion of the question. The most common responses were related to shared tools and shared facilities, enhancements to in-house fabrication facilities, and enhancements to available CMOS post-processing for CMOS+X activities.




Figure 20: Responses related to enhancing current facilities.

Responses related to the power of sharing facilities and tools were the most common response in this part of the question. Representative responses include:

- To this end, the NAE (National Academy of Engineering) and NSF developed iPodia (<https://ipodia.usc.edu/what-is-ipodia/>) tool for interactive collaborative student learning and training can be used to facilitate the ETRCs' intended purposes. The iPodia IT tool is currently used by <university> for such remote interactive learning and training of student across the globe. It enables the teaching of classes, as well as microelectronics-based projects (e.g., chips, boards, etc.) for students across many schools in the U.S. and globe. The use of iPodia-like tools will be essential in the

deployment of ETRCs, since they will enable wide access by many regional institutions, including HBCUs and other minority institutions, where students may find commute difficult at a time.

- Attention to university infrastructure should extend to facilities for metrology, CAD, system design and prototyping, testing, and packaging, and access to IC shuttle runs. Existing shared facilities should support these resources for the benefit of the entire community and CAD licensing arrangements and necessary cyberinfrastructure should be put in place to allow flexible access by the at-large student body.
- It is highly unproductive to have students spend years learning how to fabricate chips and then laying out chips for their research, especially if they do not envision future careers in VLSI. With a pool of properly trained engineers that could offer IC design as a service, we could have students (with the proper knowledge to interface with these designers) work with these engineers to design practical prototypes.
- It would be reasonable to stand up nodes with powerful FPGAs attached to them, similar to AWS F1. It would also be helpful to have a cluster with licenses & installs of well-known tools and simulators.
- Acquisition of the tools mentioned in the previous responses would be the most straightforward solution, and all of these could be placed in an open-access facility at our university.
- More updated, better community managed PDKs, access to fabrication technology would be a huge plus.
- CAD-cloud infrastructure: Facilities in the cloud (Amazon AWS, ...) could provide easy access to a standard CAD environment with pre-installed PDKs. The cloud scenario for CAD tools also makes it far easier to control (through login), as well as to revoke access to proprietary data. Furthermore, the CAD cloud infrastructure will be much lower cost as the installation, maintenance overhead is shared among many users. Finally, CAD cloud infrastructure may simplify legal hurdles involved in controlling access to PDK.
- The NSF Chameleon facility seems to start the edge to cloud services. Something going along that direction will help hardware researchers if companies would pitch in some open spec hardware for experiments.
- We need to continuously update our compute servers and data storage to host design kits.
- These should all be added to existing shared facilities - good facilities need great management, and shared facilities offer these.

Another theme was enhancements to in-house fabrication facilities. Five of these responses specifically mentioned adding 200mm wafer fabrication facilities to enhance in-house fabrication. For example:

- Eight (8) inch processing to enable commercial CMOS+X.
- The 200-mm wafer fabs are the preferred manufacturing environment for many commercially important families of products such as ICs based on legacy nodes (for automotive, industrial, and power markets, among others), specialty products such as MEMS, smart power, CMOS image sensors, analog, RF, LEDs, and strategic technologies such as 5G and millimeter-wave, GaAs, GaN power electronics, and SiC.
- Our existing nanofabrication facilities do not have the ability post-process CMOS chips reliably. This would involve standardizing on a large die size so that fixturing could be developed to allow photoresist spinning onto the die. Additional equipment needed

would be direct write optical lithography systems and eBeam direct write systems as neither of these would be bothered by the edge bead at the die boundary. In addition, ultra-fine-pitch capable die-to-die bonding systems to allow assembly of ultra-fine pitch bumps.

- Our case may be typical of smaller Electrical Engineering departments in the US, so a list of our needs is instructive: 1. Proper cleanroom facility with managed air, water, and gas handling. 2. In-lab access to all main process tools, such as metal evaporators, sputter coaters, dry etch machines, mask aligners, rapid thermal annealing systems etc. 3. In-lab access to imaging and diagnostic tools, such as scanning electron microscope (SEM), atomic force microscope (AFM), optical surface profilometer, X-ray diffraction system etc. 4. Trained lab support technicians who can commission, run and maintain the infrastructure and train new users on the various process tools.
- Over the past few decades, a worldwide trend in semiconductor foundries has focused on increasing the size of wafers from 100mm to 200mm to 300mm. The <nanoscale facility> has incredible capabilities at the four-inch wafer scale however, to make cutting edge advances in semiconductor research there is a need to adjust our facility to accommodate more 200mm wafer processing.
- Two primary tools that could prove beneficial to the <facility's> accommodation of 200mm wafers include an ICP silicon tool and an oxide etching tool. The increased capabilities afforded by these tools would allow the <facility> to perform most key backside processing on 200mm wafers further allowing for heterogeneous integration into sub-100nm CMOS nodes available from commercial vendors.
- <University> fab used to support 250nm CMOS but is now no longer in operation. MOSIS (www.mosis.org) or MUSE (www.musesemi.com) can be funded or strengthened to provide more research shuttles to US researchers.

Some responses again mentioned MOSIS by name. For example:

- MOSIS used to provide some of the services we need. Not sure though what is happening with MOSIS, they seem to be disengaging.
- MOSIS and other fabs need to provide some handholding as students are likely to be new and green. This is not the current model.
- We use MOSIS/MUSE for chip fabrication from leading foundries such as TSMC. But the available technology requires stringent NDA forms. Currently, we cannot access even 45nm technology due to TSMC's exclusive license requirement.
- We only have limited access to ISI MOSIS facility. Hence, augmenting MOSIS and other such university centered fabrication facilities for new microelectronics design will enable US to train a large pool of students.
- MOSIS is closed. I do not know any open-access facility.

Question 5.1 (IP issues):

If your answers involve specific intellectual property in this area, please let us know what that is, and be specific as possible [e.g., cell libraries, memories, input/output (I/O) systems, field-programmable-gate-array (FPGA) resources, accelerators, simulation data, data sets, execution traces, etc.].

This question was an attempt to drill down to very specific IP issues. We received 40 responses. we include some representative comments from the responses:

- TSMC 7nm CMOS standard cells library. AIM Photonics library and funding for fabrication at a special academic rate.
- Cell libraries, memories, FPGA resources, modeling tools, data sets.
- I need PDKs of advanced nodes and exotic technologies.
- Cell libraries (particularly for sub-22nm process technologies), memory templates (particularly for ReRAM and MRAM), FPGA-based testbeds that can interface with test chips
- Access to extensive IP libraries (digital cells, I/O cells, but also ADC, DAC, PLL, RF blocks, etc.) so we can engage in research at the system level.
- Cell libraries, memories, and IOs, especially those for 7nm and below, are critical.
- IP is really a very interesting issue because, unlike manufacturing, all costs beyond R&D are minimal - distributing a file to 1 or 1 million costs essentially zero, so the only reason they are not available is so that the developer can recoup their development costs. Maybe a model like with patents, where an IP developer must open source the IP after a period of several years would balance the developer need to recoup the costs with the society need to gain access to the IP.
- IP such as DDR, PCIe, PVT sensors, CPUs, etc. are very expensive. Architecture work would be substantially enhanced through simulation services such as access to the Cadence Palladium system that AFRL has.
- We use ARM physical IP blocks (memory compiler, standard cells, IOs) in our design and CAD tool testing. However, the choice of technology node is limited, and we lack analog design IPs.
- The FPGAs we use for emulation are extremely expensive, on the order of \$10-15K.
- We can benefit from access to better EDA resources that include better servers, software packages, commercial or open-source cell libraries and device qualification data.
- We are currently not hindered by access to any intellectual property.
- I would request access to 45nm cell libraries, specifically so students can train at feature tech sizes prior to industry. As of now, I'd request a Palladium Z1 SoC emulation system from Cadence, as they can simulate 140 million gates per hour.
- The CAD cloud infrastructure would be particularly helpful to scale up technology access. In many universities, the IT department is insufficiently experienced to build out a high-quality software environment.
- CMOS integration with emerging memory crossbar arrays (RRAM, PCRAM, MRAM, STRAM, etc.) would be great.
- Cell libraries (including design rules), memories, I/O systems, and FPGA resources
- State-of-the-art EDA tools from Cadence or Synopsys. Cell and memory libraries from fabs (TSMC, Samsung). Access to package design services and package houses.

- Pad frames. That's another major issue. Whenever we move to a new technology, pad frames are NOT provided, and it can take many months developing the pad frame before anything useful can be fabricated.
- RISC-V IP including accessible IP for I/O, standard cells memory, support of open-source projects developing this IP.
- As mentioned above, the big thing is compiled SRAMs, DRAM controllers, and other embedded memories.
- Transistor models - HSPICE 14nm to 2nm Transistor models - STTRAM Reliability data - STTRAM Architecture level power models McPAT for 7nm or lower with validated industry data
- We would love to have access to standard interfaces (e.g., DDR4 PHY and controllers), PCI express, etc.

Question 5.2 (Industry connections):

If your answers involve more general connections to semiconductor and microelectronic systems, please describe how these systems are connected to and critical to the research challenges that you described previously.

This question was designed to see if there were specific connections to industry that are missing. We received 21 responses to this question. Representative responses include:

- Close interactions between industry and academic are key to producing effective research in appropriate areas, well trained students, and improved products.
- Novel device arrays that require CMOS support circuitry can be demonstrated at modest scale either by connecting two separate dies together with ultra-fine pitch bumping technology or else the CMOS wafers must be available in planarized form to allow fabrication of the novel devices on top of completed CMOS wafers.
- Having access to expert knowledge and forecasting about future technologies and devices is critical to being able to use them for computer architectures.
- Having connections to foundries working on emerging memory arrays can allow us to identify partners who are willing to open memory macros and let us change some of the access circuitry to enable new architectural functionality.
- I work on devices and materials and need guidance from systems people as to what they need for a big win. They also need guidance from me and my colleagues on what is possible - the co-design scenario. Specifically, I need to know how much power is acceptable, what speeds are acceptable, what voltages are needed, what topologies are most useful and what applications really need help.
- joint call for proposals between US-Europe or US-Singapore US-Taiwan, US-Japan etc. on semiconductor research
- With the fabrication gets into deep nano scales, EDA especially physical design, are tightly related with fabs. It is necessary for researchers to have the access to and collaboration with fabs to conduct that research.
- Foundry is a big limitation to our research which focuses on "integrated" design. Currently, foundry access is very limited, costly and requires a lot of constraints. I hope we can have more choices for both research and education.
- Our desire to accommodate up to 200mm wafers directly relates to ongoing efforts on heterogeneous integration.

- The lack of foundry options is a larger concern for me and other relevant faculty. We currently only have executed contracts with two vendors for CMOS fabrication, one of which is significantly more expensive and has much longer turnaround times than the other. In the service we use, we are also limited to technologies above 45 nm. Although this isn't a concern for me now because our circuits don't require that level of performance, it could become a concern in the future.
- need a mechanism to have my student to stay in the industry for a summer yet paid by NSF like IUCIC program - we need a new simplified version
- Low substrate leakage and parasitics, as offered in fully depleted silicon-on-insulator (GF 22FDX) or silicon-on-sapphire (Peregrine) are critical to realizing ultra-low energy ML/AI compute-in-memory arrays.
- Need simplified interfaces with foundries.
- The semiconductor and microelectronics manufacturers should be more transparent in sharing their products reliability ratings to the research and user communities.

Question 6 (Other considerations):

Other considerations: Please let us know of other concerns or requirements that you think are important in research that involves semiconductor or microelectronic technologies. For example, describe organizational, process, learning and workforce development, access and sustainability, and diversity issues, along with any other issues more generally relevant that NSF/CISE should consider.

This question was designed as a catch-all to see if there were other issues or concerns from our respondents that were not captured in the preceding questions. We received 56 responses to this question. The general themes of the answers to this question are depicted in the chart below.

The bar to pie chart below shows the breakdown between various considerations. Interestingly, fab related issues are involved in at least 28% of the responses. Various flavors of these fab related responses are shown in the bar graph displayed at the right of the chart. Other considerations, e.g., education, workforce, need for scholarships/fellowships, student recruitment including underrepresented groups strongly feature in this representation. Yet other considerations include need for industrial partnership, tech transfer, startups, adverse effect of proprietary information on basic research, and need for an interdisciplinary approach.

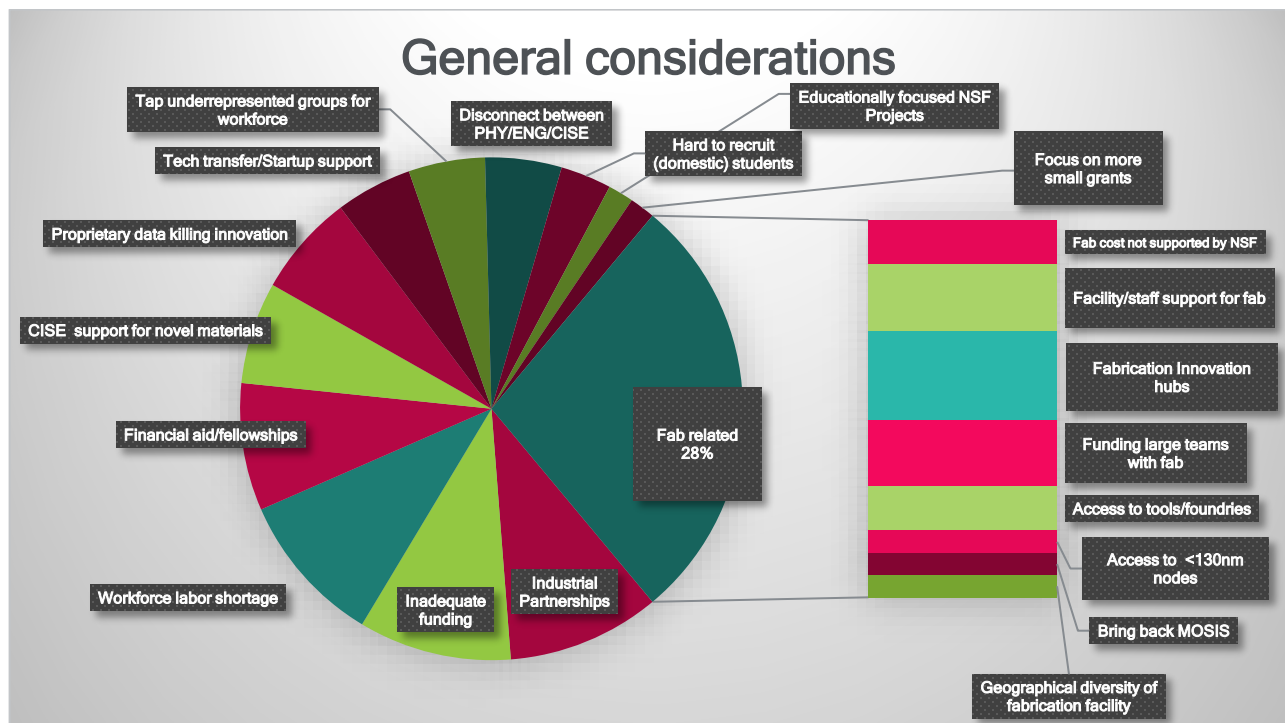


Figure 21: Some general considerations and recommendations for future

Some representative comments from the answers to this question include:

- We would love to know if there is an opportunity that NSF could help us with our new cleanroom.
- I think NSF/CISE is trying to engage and foster more industrial/academic interdisciplinary research. This will become even more critical as we move further into post-Moore, post-CMOS eras where classic separation of concerns of device electronics were achieved by industry under a standard and stable set of abstractions.
- I agree with the MIT paper that emphasized that a US MS microelectronics fellowship program is needed. Many US students graduate with debt which is a deterrence to going on.
- Multi-part response:
 - The human factor is as critical as buildings and instruments that are needed to support the academic contribution to the national microelectronics need. Highly qualified, well-motivated technical staff is an integral element of a successful operation.
 - A national program that aims to restore U. S. microelectronics leadership should also invest in the creation of new faculty slots at U.S. colleges and universities and provide flexible career-initiation grants for equipment and research support in the early years of a faculty career.
 - A renewed partnership in microelectronics between industry and academia should recruit seasoned and experienced researchers from industry to participate in university education and research activities.
 - Technology Translation programs are needed to facilitate the maturation of technology in appropriate university environments and the subsequent translation to external foundries and corporate R&D laboratories.
 - Programs to support the generation and nurturing of microelectronics startups by partially underwriting their user fees at shared university facilities.
- To facilitate both the discovery and application of new building blocks, a new form of innovation hubs would be effective. The innovation hub could be a multi-disciplinary, multi-university entity, that is established in the US. This could collaborate with the Department of Energy (DOE) national laboratories and industrial partners to take advantage of the investments in advanced fabrication and characterization capabilities.
- Removing barriers for multi-institution collaborations with easy access to PDKs and CAD tools and the possibility to share designs.
- Advances in semiconductor and microelectronics technologies and specifically on their fundamental building blocks ought to remain a focus of NSF research funding.
- My primary issue is access to PDKs and fabrication of advanced CMOS nodes and exotic technologies.
- One major hurdle that we have faced in our cross-stack research is the lack of funding programs that support teams pursuing work from software and architecture down to novel materials. At the NSF, a particular obstacle is the fact that the semiconductor and electronics research is primarily funded by ENG, while the architecture and software systems work is funded by CISE.
- One challenge we have run into is the teaming required to be competitive for funding.
- We need encourage young people to get into engineering, especially semiconductor and microelectronics. We have hard time to recruit graduate students who are American citizens.
- Realistically, the key strength for academia in engineering is workforce development. Major engineering research is happening in industry.

- Look at models that work and try to emulate - IMEC, CMC, CMP/Europractice, etc. We should at least be at the same level not behind.
- Very good semiconductor research dies because it is not at one of the very few well-funded universities or does not receive commercialization support. I believe this is one of the biggest issues faced today.
- A large section of semiconductor community that does not work on Silicon feel left out of large-scale initiatives at NSF despite being some of the largest growth areas in semiconductors.
- NSF allows for fundamental research but normally does not provide sufficient funds for the fabrication of chips. So, unfortunately, this has created a vacuum for VLSI/Chip/Analog research and in the US. Other countries recognize the importance of keeping these industries and provide support for academic universities. This must change if we want our students to continue to innovate within this critical industry.
- The full spectrum of the workforce development i.e., rooting semiconductor knowledge from junior high school towards doctoral degree is suggested,
- I want to highlight one important missing element is the industry collaboration. Semiconductor is largely driven by industry and commercial usage. But the current collaboration between company and universities are weak and not systematic.
- In addition to capital equipment acquisition, other essential support is needed for a working semiconductor device fabrication facility. This includes a set of qualified support personnel, graduate studentships, consumable allocations and facilities and equipment insurance and maintenance contracts.
- Today's challenges when demonstrating technology innovations at scale would be more easily addressed if the equipment housed at the NNCI cleanrooms provided start to finish fabrication needs and included at-scale (200mm - 300mm) state-of-the-art hetero-integration and packaging capabilities. Staff would be able to hone expertise in FEOL and BEOL design and processing.
- I consider the lack of access for underrepresented groups to EDA/VLSI/CMOS+X to be our nation's single largest emerging challenge.
- Frequently, I found access to CAD know-how one of the main showstoppers for progress. That is, the answer to a design issue may be found in a manual somewhere - but often it's unknown where to start. Another experience is that the use of older technology nodes in academia (e.g., 130nm) often makes it harder to find good answers to questions, because the older technology is no longer considered a main support concern.
- The CISE is currently disconnected from fundamental physics research while placing too much emphasize on engineering implementation.
- I hope we can make more clusters for device fabrication for each state at least one for the state level.
- MOSIS is an important to support Universities for semiconductor or microelectronic technologies, but it is closed now. I really hope it can re-open.
- NSF (and the funding agencies in general) have decimated microelectronics research over the past 20 years. A simplistic argument made is that hardware is not exciting and expensive. But the consequence of this inaction is what we are reaping now.
- I believe that minority schools are a relatively untapped source for workforce development in microelectronics. However, support is needed to maintain the fabrication facilities and give access to a larger number of students.
- I think encouraging university, industry and government partnership is still important.

- Funding is inadequate; consider tapping into collaborations with philanthropy and industry.
- The most limiting factor is PhD student support. grants or fellowships to cover those students would significantly impact outcomes.
- We are at an inflection point where university IC design research in the US is about to shrink dramatically. The US benefits significantly from the ~\$100B fabless design industry (Apple, Broadcom, Qualcomm, NVIDIA, etc.). But unfortunately, this industry does not support university research or education. It is critical that NSF supports cutting-edge research and education in IC design.
- Proprietary data and models are really killing innovation. Industry will not release the models, and academics will not accept the phrase "validated with industry partners" in publication submissions.
- Please stay away from large center grants and focus on single (or few) investigator grants which usually produce a much larger return for the money.
- Organizational, process, learning and workforce development, workforce development, sustainability, diversity is all hard in semiconductor education/research because the lack of funding to the right instructor/faculty. There should be funding for education. No education, no good students for future research and workforce. Stop funding private equipment projects that are only accessible by very limited people and low usage rate.

Endnote:

The RFI was primarily targeted to the NSF/CISE community, which to a large extent also overlaps with the community supported by the NSF Engineering Directorate with focus of device engineering, and to a lesser extent by the NSF Mathematical and Physical Sciences (MPS) Directorate with focus on materials research.

As for the term, “semiconductor technology”, we have taken it to broadly refer to all forms of micro- and nano-electronics, photonics, sensors, and actuators, as well as the circuit and system architecture design, manufacturing, and packaging technologies. For simplicity, we use the terms “semiconductor technology” and “micro-/nano-electronics” synonymously.

Recently, there have been several reports from various entities assessing the state of semiconductor research in the US. These include the SRC decadal plan (https://www.semiconductors.org/wp-content/uploads/2020/10/Decadal-Plan_Interim-Report.pdf), the MIT report on Reasserting US leadership in microelectronics - the Role of Universities (<https://usmicroelectronics.mit.edu/>) and several NSF workshops (e.g., https://nsfedaworkshop.nd.edu/assets/429148/nsf20_foundry_meeting_report.pdf and https://nsfedaworkshop.nd.edu/assets/432289/nsf20_edaworkshop_report.pdf). By contrast, the present report is exclusively based on the responses retrieved from the community survey and is thus purely data driven. Many of the views expressed in these latter mentioned reports largely resonate with the present survey.

Acknowledgement:

We would like to thank members of the CISE Semiconductor Working group for helping us with data collection and analysis for this RFI summary, especially Robert Beverly, CISE/OAC, and Alex Jones, CISE/CNS. We would also like to thank Goli Yamini, CISE/CCF for her help in extracting the original data and generating demographic information on the respondents to this RFI.

Glossary of Acronyms:

5G: Fifth generation wireless standard
6G: Sixth generation wireless standard
ADC: Analog-Digital Converter
AI: Artificial Intelligence
ASIC: Application Specific Integrated Circuit
AFM: Atomic Force Microscope
AFRL: Air Force Research Lab
AWS: Amazon Web Services
BEOL: Back-End of Line
CAD: Computer Aided Design
CEA: Commissariat à l'Énergie Atomique et aux Énergies Alternatives
CERN: Conseil Européen pour la Recherche Nucléaire
CCF: Computing and Communications Foundations
CISE: Computer and Information Sciences and Engineering
CMC: Canadian Microelectronic Corporation
CMOS: Complementary Metal Oxide Semiconductor
CNS: Computer and Network Systems
CPU: Central Processing Unit
CS: Computer Science
DAC: Digital-Analog Converter
DCL: Dear Colleague Letter
DRAM: Dynamic Random Access Memory
ECE: Electrical and Computer Engineering
EDA: Electronic Design Automation
FEOL: Front End of Line
FET: Field Effect Transistor
ETRC: Education and Teaching Resource Center
EU: European Union
FeRAM: Ferro-electric Random Access Memory
FinFET: Fin-based Field Effect Transistor
FPGA: Field Programmable Gate Array
HBCU: Historically Black College and University
HPC: High Performance Computing
IMEC: Inter-university MicroElectronic Centre (Belgium)
IC: Integrated Circuit
ICP: Inductively Coupled Plasma
I/O: Input / Output
IoT: Internet of Things
IP: Intellectual Property
ISI: Information Systems Institute
ISSCC: International Solid State Circuits Conference
IT: Information Technology
ITAR: International Traffic in Arms Regulations
LETI: Laboratoire d'électronique des technologies de l'information
LIDAR: Light Detection and Ranging
MEMS: Micro-Electrical-Mechanical System
MIMO: Multiple Input, Multiple Output

ML: Machine Learning
MOS: Metal Oxide Semiconductor
MOSIS: MOS Implementation Service
MPW: Multi-Project Wafer
MRAM: Magnetic Random Access Memory
NDA: Non-Disclosure Agreement
NEMS: Nano-Electrical-Mechanical System
nm: nanometer
NNCI: National Nanotechnology Coordinated Infrastructure
NVM: Non-Volatile Memory
OAC: Office of Advanced Cyberinfrastructure
PCIe: Peripheral Component Interconnect express
PCM: Phase-Change Memory
PCRAM: Phase-Change Random Access Memory
PDK: Process Design Kit
PHY: Physical Layer device
PIM: Processor in Memory
PLL: Phase-Locked Loop
PVT: Process, Voltage, Temperature
RAM: Random Access Memory
RRAM: Resistive Random Access Memory
ReRAM: Resistive Random Access Memory
RF: Radio Frequency
RFI: Request for Information
RISC: Reduced Instruction Set Computer
SEM: Scanning Electron Microscope
SOI: Silicon on Insulator
SOT-MRAM: Spin-Orbit Torque Magnetic Random Access Memory
SRC: Semiconductor Research Corporation
SSD: Solid State Disk
STTRAM: Spin Torque Transfer Random Access Memory
TSMC: Taiwan Semiconductor Manufacturing Company
TSV: Through-Silicon Via
VLSI: Very Large Scale Integration