

NSF CISE Committee of Visitors Report

For the Divisions of

Computing and Communication Foundations
Computer and Network Systems
Information and Intelligent Systems

Azer Bestavros

Boston University
(Chair)

Mary Hall

University of Utah
(CCF Co-Chair)

Sonia Fahmy

Purdue University
(CNS Co-Chair)

Julia Hirschberg

Columbia University
(IIS Co-Chair)

November 12, 2019

Executive Summary

Charged with assessment of program operations and management pertaining to proposal decisions in the Directorate for Computer and Information Science and Engineering (CISE), the 2019 Committee of Visitors (CoV) reviewed the quality and effectiveness of the merit review process, the selection of reviewers, the management of the programs under review, and the resulting portfolio of awards for the three divisions comprising CISE, namely Computing and Communications Foundations (CCF), Computer and Network Systems (CNS), and Information and Intelligent Systems (IIS), for the 5-year period from FY'14 through FY'18. In addition, the CoV considered thematic issues that cut across all three divisions and opportunities to enhance CISE's ability to deliver on its mission.

Overall, the CoV found CISE activities to be of exceptional quality: (a) management is grounded in effective processes and supported by an exceptional team; (b) the review process is robust and proactively refined to improve quality and efficiency; (c) reviewers and panelists are well-qualified and balanced; and, (d) the program portfolio addresses national priorities and relevant research, reflecting the agency's mission. This report elaborates on these findings and provides specific recommendations, which we summarize below.

This CoV review comes at a unique time of significant expansion of the reach and impact of computing on all aspects of our society and economy and a corresponding explosive growth in the size of the academic community served by CISE. This is a time of unprecedented opportunities to tackle critical computing research challenges with significant implications on the competitiveness of the US and its ability to maintain its historic leadership of the field. In light of the increasing dependence of the computing community on NSF funding, and notwithstanding CISE's exemplary operational and management processes and the admirable efforts of its dedicated staff, it is clear to the CoV that the current level of funding for CISE is compromising its ability to sustain the computing research community, and jeopardizing the future competitiveness of the US research enterprise. It is this context that leads the CoV to recommend that (1) to maintain US competitiveness in a changing global context, NSF must increase the overall funding for CISE to be at least commensurate with the growth in competitive proposal pressure and the increasing costs of conducting research in Computer Science.

The CoV review has also identified a number of other findings and recommendations, which are detailed in the remainder of this report, most notably: (2) to improve portfolio balance and to allow enough agility to target emerging areas of research, CISE should realign or sunset programs more often; (3) to stay true to its original intent at a time of significant growth of junior faculty, CISE should consider introducing adjustments to the CAREER program; (4) to grow understanding and awareness of broader impacts throughout the research community, CISE should pursue approaches that intentionally clarify and effectively evaluate the broader impacts of proposed work; (5) to increase diversity, CISE should continue to develop and evaluate initiatives that promote the participation of under-represented groups in research, proposal review, and in the computing pipeline; (6) to reduce workload, introduce flexibility, and increase access to larger pools of qualified and diverse reviewers, CISE should consider the use of rolling proposal submissions; (7) to enable analysis of pre/post-panel decision-making processes, CISE should introduce mechanisms to document reviewer selection practices and factors influencing fund/no-fund decisions; (8) to improve program evaluation and reduce unnecessary burden, CISE should reconsider the need for divisional annual reports and focus instead on regular data collection and longer-term trend analysis; (9) to enhance productivity in award selection and oversight, CISE should pursue efforts to develop needs-based tools that are integrated with existing business processes; and, (10) to attract highly qualified and diverse members of the CS research community as IPA rotators, NSF should address disincentives hindering their recruitment and consider mechanisms that would allow them to resume an active research agenda upon returning to their home institutions.

1. Context of the CoV

The CISE 2019 CoV review comes at a unique time in the relatively young history of computer science (CS). It comes at a time of significant opportunities and also significant challenges.

1.1 Opportunities

The CoV review comes at a time when technological convergence has weaved computing into all aspects of our society, economy, and public discourse, prompting unprecedented demand for a CS-educated workforce¹ and a corresponding skyrocketing enrollment in CS degree programs at all levels.² It comes at a time when academic institutions are making significant investments by expanding their capacity for CS research and education through hiring of new faculty,³ introduction of novel programs, and development of physical and cyber infrastructures. It comes at a time when social and humanistic dimensions are deeply intertwined into CS research and education on issues ranging from security and privacy to algorithmic fairness and bias in AI. It comes at a time when basic CS research is poised to make significant leaps in a number of critical areas where the US has been a traditional superpower.

CISE's Division of **Computing and Communications Foundations** (CCF) supports foundational research in computing and communication devices and their usage, and is making major contributions to society in developing the technology underlying smart cities, smart grids, environmental monitoring, and more generally, the Internet of Things. As silicon technology is reaching scaling limits, CCF researchers are exploring non-silicon device technology, including bio-molecules, carbon nanotubes, and optics; and, fundamentally new approaches to computational models and their associated realization in architecture, including quantum and neuromorphic computing. CCF research on approximate computing technologies will make it possible to deploy future systems in heavily-constrained and fragile environments. CCF researchers develop the algorithms that maintain privacy, the systems that protect the security of our infrastructure, and the programming technologies that ensure efficiency, quality and sustainability of the software that undergirds our economy and culture.

CISE's Division of **Computer and Network Systems** (CNS) supports innovative research in designing and evaluating networked computer systems that support emerging applications such as large scale data analytics, next generation datacenter networks, virtual, augmented and mixed reality, the Internet of Things, and smart infrastructure, as well as the continually evolving Internet. The systems and protocols needed to support emerging applications must be secure, scalable, survivable, sustainable, usable, and manageable. At the same time, CNS researchers

¹ Employment of computer and information technology occupations is projected to grow 12 percent from 2018 to 2028, much faster than the average for all occupations tracked by the Bureau of Labor and Statistics. These occupations are projected to add about 546,200 new jobs with wages that are 250% higher than the median annual wage for all occupations. <https://www.bls.gov/ooh/computer-and-information-technology/>.

² See CRA Taulbee Survey for 2018 at https://cra.org/wp-content/uploads/2019/05/2018_Taulbee_Survey.pdf.

³ Responses to a survey from 176 institutions indicate that on average there were 2 open tenure-track faculty positions per department at these institutions, and that on average 75% of these positions were successfully filled. <https://cra.org/crn/2018/08/2018-computer-science-tenure-track-faculty-hiring-outcomes/>.

must continuously harness technology trends in the underlying infrastructure of computing, storage, and communication networks, in order to develop systems that can be widely deployed to benefit large segments of society. As examples, CNS researchers are designing systems and networks to increase the energy efficiency and resilience of data centers, embedded systems, and smartphones; CNS researchers are developing abstractions and algorithms to enable fog, edge, and cloud computing systems to support applications such as real-time machine learning and virtual reality with stringent latency and privacy requirements; and CNS research on 5G networking is designing resource management algorithms and network protocols for high mobility, high node density, high heterogeneity, and high capacity (with mmWave or even THz bandwidth) systems to enable applications such as autonomous vehicles and e-health.

CISE's Division of **Information and Intelligent Systems (IIS)** supports fundamental research that is critical to society, including various facets of artificial intelligence (AI) such as computer vision, speech and natural language processing, machine learning and robotics; human-computer interaction; data mining and management; social media and social network analysis; computational biology, healthcare and neuroscience; and many more areas that leverage machine intelligence to investigate and model human behavior. IIS research currently underway on societally important topics include creating new voices for people who have lost theirs or never had a voice that sounded like themselves; detection of deception, hate speech, and disinformation; robots to help children with autism develop social skills and help seniors remain in their homes; disaster relief modeling; smart agriculture; and scoring sporting competitions such as gymnastics to remove human bias. Advances enabled by IIS research have given rise to new challenges which themselves are now the subject of extensive research: as smart systems become smarter, they fall prey to reliance on large amounts of data making them operate as black boxes, thus creating the need for research in Explainable AI, and often perpetuating historical demographic biases, thus creating the need for research to recognize bias in systems used for machine translation, face recognition, job selection, loan worthiness, emergency response, court decisions on likely recidivism, and even healthcare diagnoses.

Among all science and engineering disciplines, computing research is notable in the extent to which advances in one area of the field create opportunities for research in another. As a result, while many traditional areas of computing research are associated with a single CISE division, increasingly most of the emerging challenges for computing research defy that association. Machine Learning has its foundational roots in CCF, is made feasible by scalable cloud computing research in CNS, and is fueling IIS research advances in the many applications of AI. Mitigating security, privacy, and ethical concerns of data-driven IIS applications that make use of CNS systems and network research innovations relies on CCF research in areas as diverse as cryptography, differential privacy, and algorithmic fairness. For robotics and augmented reality IIS research informs advances in CNS research for embedded systems and wireless sensor networks, which in turn leverage CCF innovations in hardware design.

The cross-cutting nature of emerging and exceedingly impactful computing research extends well beyond CISE divisions. Advances in computing research have proven to be transformative to many other science and engineering disciplines, from material science to synthetic biology, from advanced manufacturing to economics, and from computational physics to mechanized mathematics. In addition to the catalytic nature of computing research, the tools and platforms

of computing in general, and of data science in particular, have emerged as the *lingua franca* not only for research across the landscape of disciplines in academe, but also for bridging systems and solutions that span multiple sectors of the economy.⁴

1.2 Challenges

The CoV review also comes at a time of mounting challenges that stand to hinder the ability of the computing research community to capitalize on the above opportunities. It comes at a time when the CS research community is increasingly sought out for interdisciplinary collaborations, notwithstanding the challenges of recruiting and retaining academic talent in light of competition from industry. It comes at a time when federal funding for CS research is relatively flat and increasingly dependent on NSF.⁵ It comes at a time when the US leadership of CS research is challenged by the rise in global competitiveness due to the significant infusion of funding for basic research by major global players⁶ with negative implications on the ability of the US to compete for the best talents. It comes at a time of severe workforce shortages despite the booming enrollments in CS. It also comes at a time when our research enterprises as well as the innovation economy workforce is insufficiently diverse with negative implications on the ability of our field to address socio-technical challenges.

2. Findings and Recommendations

We present the findings and associated recommendations of the CoV review completed on Friday, November 8, 2019. Our findings are presented in an order that reflects what the CoV believes to be the most critical issues related to its charge to review program operations and program-level matters pertaining to proposal decisions and the resulting portfolio of awards.

2.1 Impact of the Untenable Proposal Pressure on CISE's Ability to Fulfill its Mission

By far, the most concerning of the CoV findings relates to the escalating strain on CISE's resources due to proposal pressure (total number of proposal actions), which will significantly harm CISE's ability to fulfill its mission.

Across all three divisions, there was clear evidence of an unmistakable trend of declining success rates for competitive proposals across all programs, including core, cross-cutting, and CAREER

⁴ In the US, it is predicted that data-driven technologies in healthcare alone will result in as much as \$300 billion in added value per year (<https://www.mckinsey.com/business-functions/digital-mckinsey/our-insights/big-data-the-next-frontier-for-innovation>) and we may need as many as 1.5 million data scientists to oversee the growing place of data analytics in corporate management and planning (<https://www.nsf.gov/cise/ac-data-science-report/CISEACDataScienceReport1.19.17.pdf>)

⁵ Over the last decade, the fraction of CS research supported by NSF has hovered between 83% and 87% as reported by NSF in the President's Budget Requests (see <https://www.nsf.gov/about/budget/>).

⁶ As identified in the S&E Indicators report published by the National Science Board (NSB) and the NSF NCSES, Global R&D funding increased by 250% from 2000 to 2015. The largest contributor by far was China, which accounted for 31% of that growth, followed by the US and the European Union, which accounted for 19% and 17%, respectively. <https://www.nsf.gov/statistics/2018/nsb20181/report/sections/research-and-development-u-s-trends-and-international-comparisons/highlights/>.

programs. Using figures for the total CISE budgets and the number of unique PIs and co-PIs submitting proposals and those receiving awards over the period of the CoV review (FY'14 – FY'18), the CoV concluded that the average annual funding per PI per awarded project has been consistently declining from an average of approximately \$68K in FY'14 to an average of \$58K in FY'18. That is a 15% decline. An even more alarming trend for the CoV was the extent to which competitive and even highly competitive proposals were declined. In FY'14 the total amount of funding requested by PIs of competitive (and even highly competitive) proposals topped \$269M, representing well over 40% of the overall CISE budget for that year. By FY'17, that percentage exceeded 50% with over \$418M in highly-rated, yet declined proposals from the CS community.

This state of affairs is untenable.

At a time when the US should be capitalizing on the unique opportunities highlighted at the outset of this report – by significantly boosting its investment in CS research – these trends are going in exactly the opposite directions.

Recommendation #1: *To maintain US competitiveness in a changing global context, NSF should increase overall funding for CISE to be at least commensurate with the growth in competitive proposal pressure and the increasing costs of conducting research in Computer Science.*

The relatively flat funding of CISE, combined with the increasing volume of proposal submissions (and the associated increase in proposal pressure) is also having significant negative implications on the workload of CISE program directors and staff, which we will discuss in later sections of this report. This is a major, escalating concern given the accelerating growth of computer scientists nationwide, which will result in a significant increase in the number of submitted proposals.

2.2 Program Proliferation and Portfolio Balance

As we alluded at the outset, the computing research field has a rapidly evolving landscape that is driven by formidable technological and societal challenges. Since the last CoV review in 2014, numerous research areas and communities have been born with a corresponding expansion of the research lexicon. Among many others, terms such as quantum supremacy, algorithmic fairness, explainable AI, blockchain, software guard extensions, and deep-fake technologies would not be recognizable only five years ago.

This rapidly changing research landscape calls for an agile approach to the management of CISE programs that aim to energize and direct the research community towards needed timely research advances. The computing community has been very effective in this regard, by self-organizing around new research areas and challenges, and CISE has been equally effective in launching targeted programs in support of these initiatives, which tend to span multiple CISE divisions (e.g., FMITF, SPX) and NSF directorates (e.g., NRI, CPS). Some of these programs are meant to be broad and expected to be long-running (e.g., SaTC), while others are meant to bootstrap new multi-disciplinary communities (e.g., S&CC). About 50% of CISE's budget for FY'18 is supporting these targeted programs as opposed to CISE's core programs in CCF, CNS, and IIS.

The CoV recognizes the significant value from targeted programs and commends CISE on its community-driven approach to the development of these programs. The CoV also recognizes that these programs signal to the research community research directions that should be perceived as important, and thus these programs have an impact (however indirect) on steering the national research agenda for computing.

Portfolio balance aside, the CoV also noted the increasing number and wide diversity of programs in the CISE portfolio. The list of solicitations given to the CoV included approximately 10 solicitations led by the CNS division alone, in addition to CNS participation in a large number of cross-cutting programs led by other NSF divisions and its participation in multiple foundation-wide programs, which together total 15 or so additional solicitations per year. The number of cross-cutting programs handled by the CCF division has increased from three CCF-led and 15 CCF-participating programs in 2014, to four CCF-led and 20 CCF-participating programs in 2018. Due to the additional overhead associated with managing proposals in cross-cutting programs (*e.g.*, coordination with other divisions, directorates, and organizations), any increase in the number of such programs results in a significant increase in workload on the CISE staff.

The number of partnerships with industry (*e.g.*, Intel and VMware) and other agencies (*e.g.*, NIH) has also increased over the past five years. These partnerships offer several benefits. For example, they can accelerate CISE's entry into emerging areas, such as smart health. They offer exposure to real-world problems and accelerate technology transfer; papers with industry coauthors have 6-8 times higher citation counts. And, they provide researchers with additional funding and access to unique infrastructures and data. However, industry partnerships introduce challenges related to intellectual property and licensing of technology. The required customized partnership agreements that address these challenges add to CISE program director workloads.

International collaboration programs also add to CISE's workload. CISE has recently partnered with foreign science funding agencies such as those in Israel, the European Union, Finland, Japan, and Germany. Some of these programs are highly successful due to a common view of standards, policies, and regulations.

The CoV also noted that there was a significant variation in the size, depth, and potential impact across current CISE programs. Careful consideration is needed when starting, terminating, sunseting, consolidating, and splitting programs. Setting up a new program, especially a cross-cutting program or partnership, requires considerable effort. Hence it is most effective when that start-up cost can be amortized over a long period of time, and the program can fund a significant number of projects. A large number of small, fragmented programs creates undue burden.

In its consideration of CISE's portfolio, and in light of the tight CISE budget, it is evident to the CoV that regular reviewing of the mix of targeted programs in CISE is necessary to determine if any such programs should be sunset or realigned. This would free up or retarget resources that would contribute to CISE's agility in responding to new opportunities for investments advanced by the research community as well as CISE's need to properly support its core programs. While the CoV was satisfied with the mix of targeted and core programs in the CISE portfolio, some members expressed concerns about the relatively stagnant funding of core programs and the corresponding decline in success rates.

To adapt to emerging areas while managing the workload on program directors, the CoV recommends that CISE continuously explore the potential for focusing on a smaller number of broader programs. Programs can be smoothly sunset by issuing dear colleague letters that inform PIs about alternative programs. Other programs can be consolidated, *e.g.*, CNS recently combined the NeTS and CSR programs and held a joint PI meeting. CISE should regularly review its program portfolio and solicitations to determine which programs to consolidate and which programs to sunset, in order to focus its efforts on a manageable number of programs.

Recommendation #2: To improve portfolio balance and to allow enough agility to target emerging areas of research, CISE should realign or sunset programs more often.

The CoV recognizes that making these decisions may be difficult given their impact on the smaller communities that advocate for (and are the primary beneficiaries from) these targeted programs.

2.3 Effectiveness of the CAREER Program

The CAREER and CRII programs are extremely important to cultivate scholars at early stages of their careers. The CoV expressed major concerns about the impact of the rapidly increasing number of junior faculty members in CS on the review process and the success rates of these programs.⁷ For example, many universities have hired five to nine faculty members per year over the past five years, as opposed to one or two per year as was the case over the previous decade. This significant growth will be especially impactful on the CRII and CAREER programs in the near term, and on all programs in the longer term. It is expected that this growth will lead to a doubling or more of CAREER and CRII proposal submissions. This will tremendously increase the pressure on CISE program directors, and adversely impact the CAREER and CRII success rates.

To remain effective, it is important that the CAREER and CRII programs not become excessively competitive with exceedingly low success rates. For instance, CCF aims to fund at least 30% of CAREER proposals to ensure a healthy research pipeline. Unfortunately, the CCF CAREER success rate has consistently declined from FY'14 to FY'18. The situation is similar for the decline in the success rate for CAREER proposals submitted to the CNS division. CAREER proposals submitted to the IIS division core programs did not fare much better.

This is an alarming trend, which is largely due to the unabated increase in number of CAREER proposals submitted to all three divisions, reflecting the trends we noted above regarding the recent and continued expansion of CS academic departments nationally.

The CoV noted three negative impacts of this increased selectiveness on the vitality of the pipeline of computing academics, and consequently on the competitiveness of US research.

First, this increased selectiveness makes faculty members wait to submit their CAREER proposals until fairly late in their junior careers, after they have become more experienced (and mentored) in writing proposals. This can be detrimental to their ability to recruit graduate students and grow their research groups during their first five years, ultimately harming their research progress.

⁷ See supporting data for 2018 at <http://web.cs.wpi.edu/~cew/papers/outcomes18.pdf>

More to the point, this delay does not align with the original purpose and rationale for the CAREER program – namely to have faculty with limited experience compete with each other for resources that will help them bootstrap their research agendas.

Second, CAREER awards are becoming a “gating function” for tenure, *i.e.*, the receipt of a CAREER award has become an important consideration in tenure cases at many universities. This has increased the number of submissions and resubmissions of proposals to the program from junior faculty – even from those who are close to tenure and those who have been successful in getting other forms of NSF funding – adding to proposal pressure and to the burden on CISE staff and on the reviewing community. Even worse, this “gating function” coupled with increased selectivity is exacerbating the problem of junior faculty leaving academia, which adversely impacts the training of the future computing research workforce.

Finally, the introduction of the CRII program and the excellent course-correction that CISE has implemented for that program⁸ provide a good opportunity to disentangle these two programs by making it clear that both are targeting early-career academics in the first few years of their appointment, with CRII meant to make up for insufficient startup resources, whereas CAREER is meant to support a well-thought-out plan that would lead to a successful integration and tenure in an academic unit.

Recommendation #3: To stay true to its original intent at a time of significant growth of junior faculty, CISE should consider introducing adjustments to the CAREER program.

The CoV explored possible approaches to increase the effectiveness of the CAREER and CRII programs, and to align them to their original objective of cultivating scholars at early career stages. First, increasing funding to these programs will be extremely important to keep up with the explosive growth of the number of junior faculty members. Second, it will be important to explore the pros and cons of reducing the dollar amount of CAREER awards, in order to increase the success rates, *e.g.*, by reducing their duration from five years to three years, or to a duration that does not extend beyond expected tenure and promotion review. Third, it may be useful to consider the advantages and disadvantages of requiring faculty to apply to the CAREER award in the first three years of their academic appointment, as is the current practice with the CRII program. While such adjustments may have implications on the CAREER award’s ability to support a junior faculty for a long-enough period (to allow for development of a thriving research agenda), they will serve as incentives for faculty members to seek that support earlier and will level the playing field.

⁸ The purpose of the CRII program is to provide essential resources to PIs to achieve “research independence” after completing their PhDs and obtaining their first academic appointments. In recent years, CISE observed that CRII success rates for PIs from highly-ranked, well-resourced departments were much higher than success rates for PIs from others. To realign the program with its original intent, CISE changed the latest CRII solicitation to emphasize that the program is geared towards PIs who do not otherwise have adequate resources to launch their research agendas.

2.4 Articulation and Evaluation of Broader Impacts

There have been substantial efforts to clarify NSF’s Merit Review Criteria — Intellectual Merit and Broader Impacts — and evaluate their suitability and success in advancing societal goals of scientific research. Most notably, in 2011, the National Science Board published an extensive review that determined the criteria were appropriate, but “revisions were needed; both to draw a clearer connection of the Criteria to core principles and to better articulate the essential elements of each criterion.” It was particularly noted that there were “persistent anecdotal reports about confusion related to the Broader Impacts criterion, and inconsistency in how the criterion was being applied.” Consequently, the review provided new language to describe Broader Impacts. Starting in 2012, this language was integrated into the NSF Proposal & Award Policies & Procedures Guide (PAPPG) and all NSF solicitations.

Since these changes, CISE has had two CoV reviews (2014 and 2019), and both have concluded that there remains confusion among both the PIs and the reviewers regarding: (1) what Broader Impacts means; (2) what comprises suitable Broader Impacts for a proposal; and, (3) what role it is supposed to play in the evaluation process. The CoV identified that the CISE core program solicitations may be contributing to this confusion; it mentions Broader Impacts numerous times, but does not define it until the last mention. Further, this definition is spread across two paragraphs, with an intervening paragraph listing five review elements for both Intellectual Merit and Broader Impacts. In addition to the definition, if the expectation is that Broader Impacts should have a concrete plan with deliverables commensurate with the level of funding, then this expectation should be more clearly articulated by NSF and evaluated during the review process.

Additional lengthier messaging would assist the community in understanding the role of Broader Impacts in the broader context of advancing societal goals. It would also help to have more examples of good broader impact plans,⁹ and further, clear articulation of commonly-used elements that should not be in a broader impacts plan. An additional consideration is the relationship between the Broadening Participation in Computing (BPC) plans and Broader Impact plans; BPC is just one possible Broader Impact of a proposal. For proposals with robust BPC plans, the expectations for additional Broader Impacts should be well articulated.

Recommendation #4: To improve the understanding and awareness of what constitutes broader impacts throughout the research community, CISE should pursue approaches that intentionally clarify and effectively evaluate the broader impacts of proposed work.

To improve the CISE community’s understanding of Broader Impacts, we therefore recommend that CISE: (1) review the relevant material in its solicitations to improve clarity; (2) advertise to PIs and reviewers supplementary material that clearly articulates the role of broader impacts and provides good examples and counterexamples of what constitutes a broader impact plan; and, (3) in light of the expanding role of BPC plans within CISE, consider whether disentangling BPC from Broader Impacts would facilitate further clarification.

⁹ Examples of broader impact plans are available from the website at <http://cisebroaderimpacts.org/>.

2.5 Broadening Participation Challenges and Impact on the Merit Review Process

There is a severe underrepresentation in computing of certain population groups as compared to the US population. Per the NSF CISE 2012 Strategic Plan for Broadening Participation, these groups include women, African Americans, Hispanics, Native Americans and indigenous peoples, and persons with disabilities.¹⁰ As previously noted, this underrepresentation negatively impacts the quality of the technology that is produced by the field as it is not informed by a community representing the users of that technology. It also limits US competitiveness. While the field struggles to fill computing jobs with qualified candidates, underrepresentation cuts off large segments of the population who could be contributing to the talent pool.

Beyond the negative consequences of underrepresentation of demographic groups on the field in general, it also presents challenges specifically in the CISE merit review process. In spite of significant efforts by program directors to engage diverse reviewers, the CoV observed instances where there was insufficient diversity of panels, with low participation of women and other underrepresented groups among the panelists. One reason that was cited by program directors is that researchers from underrepresented groups are overburdened with such requests, and more likely to decline a request to serve on a panel. Overall, the CoV observed comparable funding rates for underrepresented groups as for the general population of principal investigators. At the same time, there was reduced participation of underrepresented groups in new programs, and it was unclear how Minority Serving Institutions were considered as part of the funding decision process. Further, the number of proposals submitted by these groups was small, and hence the number of awards in the portfolios was similarly small, especially for minorities and those with disabilities. We also note that the information on membership in underrepresented groups is typically underreported, and hence it is extremely difficult to track. We learned from CISE that this underreporting in the past was partially due to not having a single PI identifier, which has been corrected recently and should be less of a problem for the next CoV.

To address the systemic challenge of acute underrepresentation in CS, CISE is attempting to bridge the gap between its research activities and its efforts to broaden participation in computing (BPC). Most notably, in 2018 CISE piloted a new initiative that requires all awards in selected CISE programs to include a BPC plan that describes how the project will specifically and measurably conduct activities to better include underrepresented groups. In the current system, the BPC plans are evaluated by BPC experts and refined with program directors after the award decision has been made. This approach is an appropriate transitional step as the community gains the required expertise. To further bridge the expertise gap, CISE has also funded efforts to raise awareness of BPC in the CS community, such as the resource portal bpcnet.org, which includes resources on how to develop a BPC plan, and access to organizations with which individuals and CS departments can partner. As a compelling strategy for easing the burden of BPC plans on individual PIs, a Departmental BPC Plan will foster capacity-building in BPC and consolidate individual efforts into larger-scale programs that achieve more effective institutional change. CISE is funding workshops to guide CS department leadership in developing such plans, which

¹⁰ Women, in particular, earned about 50% of Bachelor's degrees in science and engineering, but only 19% of Bachelor's degrees in CS in 2016, down from 27% in 1997, according to the 2019 NCSSES report entitled Women, Minorities, and Persons with Disabilities in Science and Engineering. See <https://nces.nsf.gov/pubs/nsf19304/>.

can then be made available and used as exemplars for other departments. Indeed, these collective efforts towards motivating the creation of meaningful BPC plans represent a major step towards the following goal in the CISE 2012 Strategic Plan for Broadening Participation: "the inclusion of BP efforts as an accepted and expected part of its research and education award portfolios."

Towards broadening participation in computing at all levels, starting with K-12 activities and throughout the computing education and research pipeline, the CoV makes the following recommendations. We encourage CISE to establish best practices for panel composition and tracking of reviewers to identify and recruit more diverse researchers. For this purpose, we recommend that CISE continue its efforts to engage underrepresented groups and organizations to widen the reviewer pool to enable more equitable and effective merit review. We encourage NSF to leverage technology to assist in employing best practices in reviewer selection and reviewer tracking. It is important to systematically capture information about the experience, geography, type of institution, gender, ethnicity, disability, and institution of reviewers and panelists to ensure broad representation and perspectives are taken into account in all aspect of reviewing and decision making. Collecting this information facilitates tracking how well individual panels, programs and divisions are doing in meeting their best practices. CISE should also consider the use of software to provide recommendations for reviewers and evaluations of proposed panels to ensure best practices are being followed. This technology could also be employed when inviting participants to workshops related to new programs, to increase participation of underrepresented groups in development of new approaches. The result of these efforts will likely contribute to greater proposal submissions from groups currently underrepresented in computing, and improve the quality of scientific advances across the computing research community as a whole.

The CoV enthusiastically supports the approach CISE has taken with regards to the BPC pilot. We recommend that CISE continue to monitor and measure progress, both in terms of how BPC plans are created and evaluated, as well as the effectiveness of the BPC activities themselves. Furthermore, sustained messaging and training is needed to advance the community's understanding of how to conduct impactful BPC activities.

***Recommendation #5:** To increase diversity, CISE should continue to develop and evaluate initiatives that promote the participation of under-represented groups in research, proposal review, and in the computing pipeline.*

2.6 Quality and Integrity of the CISE Merit Review Processes

The CoV found that the quality and integrity of the CISE merit review processes were high and the review methods were appropriate. The triage of low-rated proposals also seems to be working well, as it allows more time to discuss the more competitive proposals. While in-person panels remain a strong mechanism for proposal review, providing an opportunity for reviewers to share wisdom and perspective in order to arrive at more robust decisions, the increased use of virtual panels that leverage state-of-the-art video conferencing software allowed for more

flexibility in recruiting more diverse panels by including those with travel constraints. Where necessary, the use of ad-hoc reviews added important expertise and some proposals were passed to other programs where appropriate. The CoV suggests that program directors continue to educate reviewers on writing detailed and constructive reviews by distributing examples of anonymized, high quality reviews that are evaluative in terms of both intellectual merit and broader impacts. Panel summaries were generally of high quality, and the Review Analyses provided detailed information on the overall rationale for award decisions.

In evaluating the CISE merit review processes, it was evident to the CoV that the selection of a qualified set of diverse reviewers was one of the major concerns of program directors. In part, this is due to the over-constrained nature of the problem, stemming from proposal pressure, added reviewing load due to resubmission of previously-declined proposals, unavailability of reviewers due to teaching duties during academic year months, conflicts of interest constraints that reduce the pool of potential reviewers, and the need to organize a large number of panels for a given program in a relatively short period of time. One way to provide added flexibility to program directors is the use of rolling proposal submissions for core programs.

Recommendation #6: *To reduce load, introduce flexibility, and increase access to larger pools of qualified and diverse reviewers, CISE should consider the use of rolling proposal submissions for core programs.*

In addition to spreading out reviewing burden and providing flexibility that could help program directors deal with conflict of interest, using rolling proposal submissions could potentially reduce proposal pressure due to the removal of the “deadline trigger” for resubmission of previously-declined proposals, especially if there is a limit on the total number of proposals under consideration. The SaTC program piloted (and continues to use) a rolling proposal submission mechanism, which seems to have achieved these desirable outcomes. While there may be tradeoffs associated with adoption of rolling submissions and that “mileage may vary” for various programs, the CoV believes that considering this approach or variants thereof is warranted.

Another suggestion is that the Review Analysis include information on how panelists for the panel were selected and recruited. It might be useful for CISE to develop some “best practices” guidelines to pass along to new program directors in reviewer selection and how reviewers are asked to rate their expertise in different areas. The CoV encourages CISE also to expand its use of virtual panels where this will help to increase the diversity of the panels.

Recommendation #7: *To enable analysis of pre/post-panel decision-making processes, CISE should introduce mechanisms to document reviewer selection practices and factors influencing fund/no-fund decisions.*

For proposals not discussed in panels (NDPs) for which PIs only get (at least) three strong reviews, it was suggested that the context statement sent to those PIs include information on how to contact a program director to obtain more information than just the reviews.

2.7 Efficiency and Agility of CISE Management

Management of CISE programs is grounded in solid, well-functioning processes, exceptional staff, and effective management. CoV members were impressed with the management of CISE and its various programs (both core and cross-cutting) across all three divisions, which has contributed to notable successes. The most important contributor to these successes is the quality and hard work of program directors. The CoV also finds that the current practice of complementing permanent staff with rotators (IPAs) works very well, providing a healthy mix of institutional memory, fresh ideas, and a robust connection with members of the research community. CISE has various mechanisms in place to enable IPAs to work effectively from when they arrive and throughout their time at NSF.

The CoV was also impressed by the creativity, resourcefulness, and agility of CISE program directors. In many instances, including the breakout sessions during the in-person meeting at NSF, when members of the CoV made suggestions for tweaking of processes or procedures, they were pleasantly surprised to find out that CISE program directors have been piloting the same or very similar ideas. Also notable is CISE's collaborative approach to program management across divisions, including the effective use of clusters of program directors to decide on assignments of proposal to review panels, the recruitment of reviewers and panelists, and the assignment of proposals to panelists.

We commend CISE on the collaborative approach of its program management philosophy as well as its agility and willingness to experiment with processes for effective operations.¹¹

NSF is keen on limiting the administrative overhead to a relatively low percentage of its budget. This is much appreciated by the research community since it leaves more resources for NSF's main mission of supporting research activities. But, in a period of relatively flat funding levels, it also means that resources for staff are relatively flat as well, and the size of the scientific staff is not keeping up with the increasing workload.

Over the period of 10 years from FY'08 to FY'18 the number of CISE proposals has increased by over 50% (from 6,067 to 9,151). Over that same period, CISE's staff FTE allocation has remained relatively flat. Coupled with the unabated growth in the number of proposal submissions (observed since CoV'14 and expected in the next few years), this flat staffing level is concerning as it has serious implications not only on the quality of life of CISE program directors and staff, but also on CISE's ability to fulfill its mission.

CISE staff workload depends not only on the number of proposals and awards, but also on the number and complexity of the programs under management. Over the period of the review, the number of programs involving CISE increased by over 25% from a total of 31 programs in FY'14 to a total of 39 programs in FY'18. Cross-cutting programs, which are growing in number, tend to

¹¹ Examples include spreading out proposal submission windows in the CPS and S&CC programs with solicitations in alternating years for FY'19 and FY'20; introducing single solicitation for all CISE core programs and limitations on the number of proposals submitted by PIs to certain solicitations; using triage in review panels, where a panel may elect to forego discussion of a proposal with uniformly low reviews; piloting asynchronous review panels in FY'14 and a rebuttal phase for PIs to provide input ahead of a panel meeting in FY'15 and FY'16; and piloting a process for requesting certain proposal sections at time of award rather than time of initial submission.

have unique requirements to be coordinated with other divisions, directorates and organizations. Over the period of the review, the number of cross-cutting programs involving CISE as a non-leading directorate increased by over 36% from a total of 11 in FY'14 to a total of 15 in FY'18.

As we noted earlier in the report, the expected increase in proposal pressure and the cross-cutting nature of computing research suggest that both of these trends will continue to accelerate. Until these systemic issues are resolved, the CoV suggests that CISE consider reducing its workload by re-evaluating and possibly adjusting some of its business processes in a way that assists CISE with more effective planning and oversight.

One example along these lines is the burden imposed on administrative staff to compile CISE division annual reports. As part of its review, the CoV examined these reports from FY'14 to FY'18. While extensive, these reports provided snapshots of the state of CISE which were not designed for the task at hand – namely the identification and analysis of longitudinal trends, which members of the CoV requested from CISE staff in some instances, or figured out themselves. Such trends are far more valuable, not only for CoV purposes, but indeed for planning and oversight purposes as well as for effectively communicating the state of CISE to the research community.

***Recommendation #8:** To improve program evaluation and reduce unnecessary burden, CISE should reconsider the need for CISE divisional annual reports and focus instead on regular data collection and longer-term trend analysis.*

In general, and as a matter of course, CISE should continue to look for ways to streamline its administrative burden.

As noted in the CoV'14 report, another contributor to CISE staff workload is the quality of the software tools and of the data infrastructures that support its work. In reviewing progress since 2014, the CoV was very pleased to note upgrades in tools available for proposal processing, including the Program Director Proposal Panel Portfolio Organizer (PD-3PO) and Robo-RA, , which streamline panel formulation and proposal processing; the MyNSF panel management system and access to the Enterprise Reporting data analytics portal, which allow better tracking of budgets, annual project reports, and program portfolios, saving time and effort; and the new ways to identify potential reviewers, including the use of SurveyMonkey and the SuggestReviewer tool. CoV members commend CISE for the resourcefulness of its program directors, many of whom developed their own tools some of which have been used by other program directors across NSF. Despite these improvements, the current slate of software tools and the underlying data analytics infrastructure available to CISE staff remain problematic.¹²

***Recommendation #9:** To enhance productivity in award selection and oversight, CISE should pursue efforts to develop needs-based tools that are integrated with existing business processes.*

¹² For example, CoV members were informed that program directors have to work on information that is split between two NSF systems (FastLane and EJacket) which are not integrated.

2.8 Recruitment of IPA Rotators

CISE employs a mix of rotators (called IPAs) and permanent staff to serve as program directors. IPAs also fill leadership roles such as Division Director and CISE Assistant Director. This combination of personnel is essential to maintaining research excellence in the leadership of CISE programs. While the permanent staff form the backbone of the organization, maintain consistency and preserve institutional knowledge, the IPAs complement this experience with recent academic research experience. Therefore, IPAs are representative of the academic community they serve. Because they serve from one to four years, including the IPAs in the CISE organization, NSF is able to hire new people with new ideas on a regular basis. These short-term employees incur expenses to CISE such as onboarding, and financial support for temporary housing and maintaining their research activities at their home institutions. In spite of these extra costs, the value proposition of including IPAs is their profound contribution to the vibrancy and the intellectual excitement of the CISE community.

It has always been a challenge to recruit IPAs, who by design must disrupt their current work, set up temporary households, and learn a new job. CISE must continuously identify a diverse set of candidates who are able and willing to do this and whose home institutions are supportive. In recent years, recruiting IPAs has become more difficult due to cost-cutting changes in NSF policy. Specifically, NSF supports only one trip home per month for IPAs to remain connected to their research groups, and limits the salary commitment from NSF for IPAs to be more consistent with the salaries of permanent staff. Home institutions are then expected to fill in any gaps in salary, and home institutions or IPAs themselves cover the costs of any additional trips home to remain connected to their research groups. The value proposition for the home institution to support its employees in serving at NSF then changes. The home institution loses the employee during their service, and now must contribute more of the cost. Indeed, the very people for whom the cost-sharing burden is greatest — the leadership positions that recruit very senior people — are also highly productive individuals for whom the loss of their leadership at home is felt most acutely.

The CoV is deeply concerned about the interrelationship between cost-cutting measures at NSF related to use of IPAs and its impact on CISE's ability to recruit highly-qualified and diverse applicants for these positions. We recommend that all stakeholders be engaged in these decisions in the future to evaluate their impact.

Recommendation #10: *To attract highly qualified and diverse members of the CS research community as IPA rotators, NSF should address disincentives hindering their recruitment and consider mechanisms that would allow them to resume an active research agenda upon returning to their home institutions.*

2.9 CISE Responsiveness to CoV'14 Recommendations

CoV members were impressed by the extent to which CISE responded to recommendations of the previous CoV review in 2014. Simply put, CISE's response to the comments of the last CoV has been exemplary. The 2015 CISE response to the CoV'14 report shows that all three divisions took the CoV recommendations quite seriously. The responses give justification for the actions

taken and not taken, though in some cases we believe that the recommendations need to be reconsidered in light of the current CoV review.

CoV members took special note of (1) the use of modern video conferencing software (BlueJeans) to conduct virtual panels, which enabled CISE to expand and diversify the pool of qualified reviewers through the use of hybrid and virtual panels; (2) the use of triage to avoid lengthy discussion of proposals that receive poor scores by reviewers, which is a welcome change, saving precious panel time that allows more time to focus on the best proposals; (3) the diligence of program directors in diversifying the composition of their panels; (4) the overall balance of funding by gender and minority status, despite differences in the number of submissions; and, (5) the proactive outreach to the research community through funding of workshops in emerging areas to encourage transformational proposals as well as to identify potential reviewers.

3. Concluding Remarks

3.1 Summary Statement

Notwithstanding mounting proposal pressures, insufficient overall funding, significant staffing limitations, and antiquated support infrastructures that will significantly harm its ability to fulfill its mission in the longer term, CISE continues to deliver impressive results when it comes to the quality and integrity of the merit review process, the selection of reviewers, and the management of the programs under its purview. CISE program-level technical and management processes are resulting in proper proposal funding decisions that contribute to advancement of scientific endeavors, and also to the maintenance of an excellent portfolio balance.

3.2 CoV Process and Basis for Findings

The timeline for the CoV process is summarized in Appendix A. During the period from when the CoV was charged on September 18, 2019 until the in-person meeting was held at NSF starting on November 6, 2019, co-Chairs independently managed the work by CoV members in their subcommittees, including assignment of randomly-sampled proposal jackets and of template questions. For the four weeks preceding the in-person meeting, CoV co-chairs held regular one-hour video conferences (facilitated by CISE) with members of their subcommittees, which were joined by the CoV Chair to coordinate and share best practices across the three subcommittees.

In preparing for its review, CoV members relied on the following information:

- CoV 2014 report and the CISE response;
- NSF strategic plans for FY'14-FY'18 and FY'18-FY'22;
- Strategic plans and various documents for various recent national initiatives;
- Division annual reports for CCF, CNS, and IIS over the period of the review FY'14-FY'18;
- Presentations to CCF, CNS, and IIS subcommittees about division programs/outcomes;
- Proposal jackets initially provided by CISE to the CCF, CNS, and IIS subcommittees;
- Additional proposal jackets provided in response to requests by CoV members; and
- Data from CISE in response to CoV Chair requests before/during the in-person meeting.

In total, CoV members reviewed 589 jackets, with 182 reviewed by the CCF subcommittee, 215 by the CNS subcommittee, and 192 by the IIS subcommittee.

3.3 Effectiveness of the CoV Process and Suggested Improvements

Overall, members of the CoV found the CoV process to be very well designed with no substantive recommendations to make in terms of adjustments. The CoV recognizes that this process was adjusted in response to recommendations from CoV'14 and confirms that the changes to the process (by completing a solid draft of the subcommittee CoV reports ahead of the in-person meeting) allowed for a very productive in-person meeting at NSF. The agenda for that meeting is provided in Appendix B.

The CoV also found that the information provided to the CoV (annual reports and jackets) was extensive and adequate for the CoV to accomplish its task. That said, a few tweaks were suggested to make the process of digesting the information easier. These mostly concern the presentation of the data in the divisional annual reports (and/or data compiled for CoV members in advance), including a preference for summaries that reveal longitudinal trends as opposed to snapshots of current statistics (see Recommendation #8).

Members of the CoV were very appreciative of CISE's efforts to minimize the COI period during which they could not submit proposals. They also suggested that these schedules and constraints be communicated as early as possible, preferably at the time of initial invitation. Some CoV members noted that they may have under-estimated the time commitments needed ahead of the in-person meeting at the NSF, suggesting that it would be advisable for CISE to communicate clearly that prospective CoV members are advised not to accept the invitation to serve if they have other significant commitments during the 4-6 weeks preceding the CoV in-person meeting.

Members of the CoV had nothing but praise for the agenda of the in-person meeting at NSF. In particular, CoV members found that breakout sessions on the first day of the meeting were invaluable. Topics for the breakout sessions were organized around cross-cutting issues identified by the CoV Chair ahead of the meeting (based on initial findings by the three subcommittees and based on discussions with CISE in advance).

They felt that the amount of time dedicated to each part of the meeting was adequate and that the topics selected for the parallel breakout sessions on the first day were on point. They also felt that their interactions with fellow CoV members from other divisions and with a large number of CISE program directors and staff during the parallel breakout sessions were very informative and by far the most rewarding part of the meeting.

CoV members underscored the importance of having weekly meetings over the month preceding the CoV in-person meeting, and the importance of setting an early deadline for completing the division templates, which forced them to complete the initial analysis early and have time to read the responses of others. CoV members also noted that completing an initial draft of the divisional templates ahead of the meeting was valuable in the discussions with the division leadership and in the breakout groups because it meant that CoV members could ask thoughtful questions and get deeper responses. The handful of CoV members who served on CoV reviews in 2009 and/or in 2014 noted that the organization of the 2019 CoV was much better than that of prior ones.

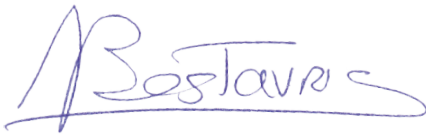
3.4 Acknowledgment of NSF Staff and CoV Members

In closing, as co-authors of this report, we wish to acknowledge the efforts of all 20 CoV members for their significant and meticulous work over the course of the last two months. We know that we speak on behalf of the entire computing research community in saying “thank you” for a job well done.

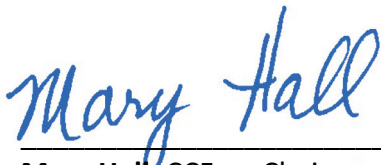
We are also grateful for the tireless work by CISE leadership, its program directors, and staff in preparation for and throughout the entire CoV review process. Your dedication and service to our research community is inspiring. We extend our special thanks to all CISE program directors and staff members who participated in the in-person CoV meeting at NSF. In particular, we would like to single out CISE Division Directors, Deputy Division Directors, and Operations Managers Rance Cleaveland, Thyagarajan (Thyaga) Nandagopal, and Velma Swales in CCF, Ken Calvert, Jeremy Epstein, and Tracey Zeigler in CNS, and Henry Kautz, Joydip (JD) Kundu, and Chantini Reid in IIS; as well as CISE staff Gwen Owens, Directorate Operations Officer, Carmen Whitson, Staff Associate for Budget and Science Strategy, and Kristen Oberright, Senior Program Analyst.

Finally, we are truly indebted to Acting Assistant Director for CISE, Erwin Gianchandani, and to Deputy Division Director for IIS, JD Kundu, for going well beyond the call of duty in supporting the work of the CoV. Your seemingly infinite energy and enthusiastic dedication to our research community are evident and contagious. Thank you!

Respectfully submitted to NSF on Tuesday, November 12, 2019 by



Azer Bestavros, CoV Chair
Boston University



Mary Hall, CCF co-Chair
University of Utah



Sonia Fahmy, CNS co-Chair
Purdue University



Julia Hirschberg, IIS co-Chair
Columbia University

Appendix A: Timeline of the CoV Process

- Late Fall 2018: Identification and commitment of Azer Bestavros to serve as the overall CoV Chair
- Mid Spring 2019: Identification and commitment of Mary Hall, Sonia Fahmy, and Julia Hirschberg to serve as CoV co-Chairs for the CCF, CNS, and IIS Divisions, respectively
- June 11, 2019: Introductory video conference between CoV Chair, CoV Co-Chairs, and members of the CISE leadership responsible for management of the CoV process
- August 1, 2019: Identification and commitment of 20 additional members of the CoV to serve on the subcommittees for CCF (6), CNS (8), and IIS (6), respectively
- August 7, 2019: Introductory video conference with all CoV members during which CISE leadership walked through the CoV process, timeline, and the prior CoV report
- September 18, 2019: Virtual meeting with entire CoV for CISE to formally kickoff the CoV, discuss the CoV charge, conflict of interest guidelines, and provide an overview of CISE
- September 19, 2019: Formal registration and Conflict of Interest forms for all CoV members were requested
- September 30, 2019: Virtual meeting of the IIS sub-committee during which CISE IIS leadership presented overview of the division, followed by Q&A about process
- October 1, 2019: Virtual meeting of the CNS sub-committee during which CISE CNS leadership presented overview of the division, followed by Q&A about process
- October 4, 2019: Virtual meeting of the CCF sub-committee during which CISE CCF leadership presented overview of the division, followed by Q&A about process
- October 4, 2019: SharePoint site and EJacket system for CoV set up and made available to CoV members
- October 4-30, 2019: Review of Jackets and prior annual reports by CoV subcommittees organized by CoV co-chair and coordinated through periodic calls joined by the CoV Chair
- October 30, 2019: Initial drafts of the CoV subcommittee reports shared with CoV Chair to identify cross-cutting themes to be considered for the in-person meeting
- November 1, 2019: Initial set of questions and requests for more information compiled by the CoV Chair and communicated with CISE
- November 4, 2019: Initial response from CISE including answers to preliminary CoV questions and additional information shared with CoV
- November 5, 2019: Agenda for the in-person meeting at the NSF is finalized and shared with CoV members by the CoV Chair in consultation with CISE
- November 6-7, 2019: In-person meeting of the full CoV membership at the NSF, during which subcommittee reports were completed and cross-cutting issues discussed
- November 8, 2019: In-person meeting of the CoV Chair and co-Chair to finalize set of recommendations and initial outline of the CoV report
- November 12, 2019: CoV report finalized and submitted to CISE by the CoV Chair

Appendix B: Agenda for the In-person CoV Meeting at the NSF

Wednesday, November 6 – Friday, November 8, 2019

Wednesday (11/6)

- | | |
|------------------|---|
| 08:15am-08:30am: | Refreshments |
| 08:30am-09:00am: | CISE Acting AD welcome and reminder of CoV charge |
| 09:00am-09:30am: | CoV co-chairs Welcome and CoV member introduction |
| 09:30am-10:00am: | Subcommittees convene in separate rooms to continue drafting division templates <ul style="list-style-type: none">• CCF subcommittee in Room W 3210• CNS subcommittee in Room E 3450• IIS subcommittee in Room W 3160 |
| 10:00am-11:00am: | Each subcommittee meets with corresponding division leadership (division director, deputy division director, and operations manager) in their separate subcommittee rooms |
| 11:00am-12:00pm: | Subcommittees continue drafting division templates and prepare report-out to full CoV in their separate subcommittee rooms |
| 12:00pm-01:30pm: | Working lunch: Report out/discuss draft division templates to full CoV in Room E 2030 |
| 01:30pm-02:45pm: | Breakouts on select issues, part 1 <ul style="list-style-type: none">• Room E 3450: Merit review process and portfolio balance• Room W 3210: PI support over career pipeline• Room W 3160: Broadening participation |
| 02:45pm-03:15pm: | Coffee Break: Develop report-outs to full CoV |
| 03:15pm-04:30pm: | Breakouts on select issues, part 2 <ul style="list-style-type: none">• Room E 3450: Broader impacts• Room W 3210: Partnerships• Room W 3160: CISE Budget, staffing, recruiting, and workload |
| 4:30pm-05:00pm: | Develop report-outs to full CoV |
| 5:30pm | CoV members only Dinner
Galae Thai, 215 Swamp Fox Rd, Alexandria, VA 22314 |

Thursday (11/7)

08:15am-08:30am:	Refreshments
08:30am-10:00am:	Report out of discussion in breakout sessions to full CoV in Room E 2030
10:00am-10:30am:	Coffee Break
10:30am-12:00pm:	Subcommittees reconvene separately to finalize division templates in light of report/discussion of breakouts <ul style="list-style-type: none">• CCF subcommittee in Room W 3210• CNS subcommittee in Room E 3450• IIS subcommittee in Room W 3160
12:00pm-01:00pm:	Working lunch -- Report out/discuss final division templates to full CoV in Room E 2030
01:00pm-01:30pm:	Break / CoV co-Chairs prepare observations of cross-cutting issues
01:30pm-02:00pm:	Co-chairs present observations of cross-cutting issues to full CoV and provide outline of report on cross-cutting themes
02:00pm-03:30pm:	Full CoV discussion of cross-cutting themes
03:30pm-04:00pm:	Coffee Break: CoV co-Chairs prepare cross-cutting findings and recommendations
04:00pm-05:00pm:	CoV Chair presents cross-cutting findings and recommendations to CISE management
05:00pm	Most CoV members depart

Friday (11/8)

08:30am-04:30pm:	CoV co-Chairs finalize report and write executive summary in Room E 3450
------------------	--

CORE QUESTIONS and REPORT TEMPLATE for FY 2019 NSF COMMITTEE OF VISITOR (COV) REVIEWS

Guidance to NSF Staff: This document includes the FY 2019 set of Core Questions and the COV Report Template for use by NSF staff when preparing and conducting COVs during FY 2019. Specific guidance for NSF staff describing the COV review process is described in the “COV Reviews” section of NSF’s Administrative Policies and Procedures which can be obtained at <https://inside.nsf.gov/tools/toolsdocuments/Inside%20NSF%20Documents/Policy,%20Procedures,%20Roles%20and%20Responsibilities%20for%20COV%20Reviews%20and%20Program%20Portfolio%20Reviews.pdf>¹.

NSF relies on the judgment of external experts to maintain high standards of program management, to provide advice for continuous improvement of NSF performance, and to ensure openness to the research and education community served by the Foundation. COV reviews provide NSF with external expert judgments in two areas: (1) assessments of the quality and integrity of program operations; and (2) program-level technical and managerial matters pertaining to proposal decisions.

The program(s) under review may include several sub-activities as well as NSF-wide activities. The directorate or division may instruct the COV to provide answers addressing a cluster or group of programs – a portfolio of activities integrated as a whole – or to provide answers specific to the sub-activities of the program, with the latter requiring more time but providing more detailed information.

The Division or Directorate may add questions relevant to the activities under review. Copies of the report template and the charge to the COV should be provided to OIA prior to forwarding to the COV. In order to provide COV members adequate time to read and consider the COV materials, including proposal jackets, COV members should be given access to the materials in the eJacket COV module approximately four weeks before the scheduled face-to-face meeting of the COV members. Before providing access to jackets, the Conflict of Interest and Confidentiality briefing for COV members should be conducted by webinar, during which, NSF staff should also summarize the scope of the program(s) under review and answer COV questions about the template.

Suggested sources of information for COVs to consider are provided for each item. As indicated, a resource for NSF staff preparing data for COVs is the Enterprise Information System (EIS) –Web COV module, which can be accessed by NSF staff only at <http://budg-eis-01/eisportal/default.aspx>. In addition, NSF staff preparing for the COV should consider other sources of information, as appropriate for the programs under review.

For programs using section IV (addressing portfolio balance), the program should provide the COV with a statement of the program’s portfolio goals and ask specific questions about the program under review. Some suggestions regarding portfolio dimensions are given on the template. These suggestions will not be appropriate for all programs.

Guidance to the COV: The COV report should provide a balanced assessment of NSF’s performance in the integrity and efficiency of the **processes** related to proposal review. Discussions leading to answers of the Core Questions will require study of confidential material such as declined proposals and reviewer comments. **COV reports should not contain confidential material or specific information about declined proposals.** The reports generated by COVs are made available to the public.

¹ This document has three parts: (1) Policy, (2) Procedures, and (3) Roles & Responsibilities.

We encourage COV members to provide comments to NSF on how to improve in all areas, as well as suggestions for the COV process, format, and questions. For past COV reports, please see <http://www.nsf.gov/od/oia/activities/cov/>.

FY 2019 REPORT TEMPLATE FOR NSF COMMITTEES OF VISITORS (COVs)

The table below should be completed by program staff.

Date of COV: November 6-8, 2019			
Program/Cluster/Section: Core Programs: <ul style="list-style-type: none"> • Algorithmic Foundations (AF) • Communications and Information Foundations (CIF) • Foundations of Emerging Technologies (FET) (started in FY2019) • Software and Hardware Foundations (SHF) Cross-Cutting Programs: <ul style="list-style-type: none"> • Expeditions in Computing (EiC) • Formal Methods in the Field (FMitF) • NSF/Intel Partnership on Foundational Microarchitecture Research (FoMR) • Scalable Parallelism in the Extreme (SPX) • Transdisciplinary Research in Principles of Data Science (TRIPODS) • Algorithms in the Field (AitF) • Cyber-Enabled Sustainability Science and Engineering (CyberSEES) • Exploiting Parallelism and Scalability (XPS) • NSF/Intel Partnership on Computer Assisted Programming for Heterogeneous Architectures (CAPA) 			
Division: Computing and Communication Foundations (CCF)			
Directorate: Computer and Information Science and Engineering (CISE)			
Number of actions reviewed: 182			
Awards: 100			
Declinations: 79			
Other: 3			
Total number of actions within Program/Cluster/Division during period under review:			
	# of Awards	# of Declines	Total
FY 2014	417	1184	1601
FY 2015	455	1352	1807
FY 2016	459	1220	1679
FY 2017	416	1299	1715
FY 2018	442	1210	1652

Manner in which reviewed actions were selected:

- (1) Random selection from bins for programs over the covered period, evenly divided between awards and declinations;
- (2) Random selection of EAGER, RAPID, travel grant, and workshop awards; and
- (3) Additional random selection from bins requested by the COV.

COV Membership

	Name	Affiliation
COV Chair or Co-Chairs:	Overall COV Chair: Azer Bestavros Subcommittee Chair for CCF: Mary Hall	Boston University University of Utah
COV Members:	Joanne Atlee Randy Bryant Martine Ceberio Dexter Kozen Teresa Przytycka Baruch Schieber	University of Waterloo Carnegie Mellon University University of Texas, El Paso Cornell University National Institute of Health New Jersey Institute of Technology

MERIT REVIEW CRITERIA

An understanding of NSF's merit review criteria is important in order to answer some of the questions on the template. Reproduced below is the information provided to proposers in the Grant Proposal Guide about the merit review criteria and the principles associated with them. Also included is a description of some examples of broader impacts, provided by the National Science Board

1. Merit Review Principles

These principles are to be given due diligence by PIs and organizations when preparing proposals and managing projects, by reviewers when reading and evaluating proposals, and by NSF program staff when determining whether or not to recommend proposals for funding and while overseeing awards. Given that NSF is the primary federal agency charged with nurturing and supporting excellence in basic research and education, the following three principles apply:

- All NSF projects should be of the highest quality and have the potential to advance, if not transform, the frontiers of knowledge.
- NSF projects, in the aggregate, should contribute more broadly to achieving societal goals. These broader impacts may be accomplished through the research itself, through activities that are directly related to specific research projects, or through activities that are supported by, but are complementary to, the project. The project activities may be based on previously established and/or innovative methods and approaches, but in either case must be well justified.
- Meaningful assessment and evaluation of NSF funded projects should be based on appropriate metrics, keeping in mind the likely correlation between the effect of broader impacts and the resources provided to implement projects. If the size of the activity is limited, evaluation of that activity in isolation is not likely to be meaningful. Thus, assessing the effectiveness of these activities may best be done at a higher, more aggregated, level than the individual project.

With respect to the third principle, even if assessment of Broader Impacts outcomes for particular projects is done at an aggregated level, PIs are expected to be accountable for carrying out the activities described in the funded project. Thus, individual projects should include clearly stated goals, specific descriptions of the activities that the PI intends to do, and a plan in place to document the outputs of those activities. These three merit review principles provide the basis for the merit review criteria, as well as a context within which the users of the criteria can better understand their intent.

2. Merit Review Criteria

All NSF proposals are evaluated through use of two National Science Board approved merit review criteria. In some instances, however, NSF will employ additional criteria as required to highlight the specific objectives of certain programs and activities.

The two merit review criteria are listed below. Both criteria are to be given full consideration during the review and decision-making processes; each criterion is necessary but neither, by itself, is sufficient. Therefore, proposers must fully address both criteria. ([PAPPG Chapter II.C.2.d.\(i\)](#) contains additional information for use by proposers in development of the Project Description section of the proposal.) Reviewers are strongly encouraged to review the criteria, including [PAPPG Chapter II.C.2.d.\(i\)](#), prior to the review of a proposal.

When evaluating NSF proposals, reviewers will be asked to consider what the proposers want to do, why they want to do it, how they plan to do it, how they will know if they succeed, and what benefits could accrue if the project is successful. These issues apply both to the technical aspects of the proposal and the way in which the project may make broader contributions. To that end, reviewers will be asked to evaluate all proposals against two criteria:

- **Intellectual Merit:** The Intellectual Merit criterion encompasses the potential to advance knowledge; and
- **Broader Impacts:** The Broader Impacts criterion encompasses the potential to benefit society and contribute to the achievement of specific, desired societal outcomes.

The following elements should be considered in the review for both criteria:

1. What is the potential for the proposed activity to:
 - a. Advance knowledge and understanding within its own field or across different fields (Intellectual Merit); and
 - b. Benefit society or advance desired societal outcomes (Broader Impacts)?
2. To what extent do the proposed activities suggest and explore creative, original, or potentially transformative concepts?
3. Is the plan for carrying out the proposed activities well-reasoned, well-organized, and based on a sound rationale? Does the plan incorporate a mechanism to assess success?
4. How well qualified is the individual, team, or organization to conduct the proposed activities?
5. Are there adequate resources available to the PI (either at the home organization or through collaborations) to carry out the proposed activities?

3. Examples of Broader Impacts

The National Science Board described some examples of broader impacts of research, beyond the intrinsic importance of advancing knowledge.² “These outcomes include (but are not limited to) increased participation of women, persons with disabilities, and underrepresented minorities in science, technology, engineering, and mathematics (STEM); improved STEM education at all levels; increased public scientific literacy and public engagement with science and technology; improved well-being of individuals in society; development of a globally competitive STEM workforce; increased partnerships between academia, industry, and others; increased national security; increased economic competitiveness of the United States; and enhanced infrastructure for research and education. These examples of societally relevant outcomes should not be considered either comprehensive or prescriptive. Investigators may include appropriate outcomes not covered by these examples.”

² [NSB-MR-11-22](#)

INTEGRITY AND EFFICIENCY OF THE PROGRAM'S PROCESSES AND MANAGEMENT

Briefly discuss and provide comments for *each* relevant aspect of the program's review process and management. Comments should be based on a review of proposal actions (awards, declinations, returns without review, and withdrawals) that were *completed within the past four fiscal years*. Provide comments for *each* program being reviewed and for those questions that are relevant to the program(s) under review. Quantitative information may be required for some questions. Constructive comments noting areas in need of improvement are encouraged.

I. Questions about the quality and effectiveness of the program's use of merit review process. Please answer the following questions about the effectiveness of the merit review process and provide comments or concerns in the space below the question.

QUALITY AND EFFECTIVENESS OF MERIT REVIEW PROCESS	YES, NO, DATA NOT AVAILABLE, or NOT APPLICABLE
<p>1. Are the review methods (for example, panel, ad hoc, site visits) appropriate?</p> <p>Comments:</p> <p>For programs managed by CCF, the primary approach to reviewing proposals is to identify a panel of external reviewers who are experts in the areas of the collection of proposals to be reviewed and have those panelists participate in an on-site review at NSF for roughly 2 days. Sometimes, these panels may include one or more virtual panelists, but the majority of panelists are physically present at the meeting. Depending on the year, 81-91% of panels have been conducted this way, with remaining panels entirely virtual. It appears the use of virtual panels is substantially higher than in the previous report. During the COV meeting with CCF division leadership, it was noted that in 2013, a goal of having 15% virtual panels was set, and this goal was surpassed in 2014 with 18% virtual panels. Since this pilot, the frequency of virtual panels has been used as appropriate, but has lessened awaiting better videoconferencing technology.</p> <p>To broaden the expertise of the panels, ad hoc mail-in reviews are used in addition to those of panelists for a modest 3-5% of the reviews. A small percentage of proposals (6-8.5%) are internally evaluated and awarded at the discretion of the program director. CISE varies strategies for different programs, such as using pre-proposals for larger projects and broadening panels beyond computer scientists for interdisciplinary programs.</p> <p>The analysis of proposal jackets revealed a number of strategies to reduce reviewer and program director load that seemed to strike a good balance with respect to feedback to most principal investigators. The triaging of proposals, and choosing not to discuss proposals with low review marks, seems like an</p>	<p>YES</p>

<p>effective way for panels to spend more time discussing the merits of proposals that stand a chance of being funded. More detail about the latter set of proposals help the program directors make better funding decisions. Moreover, the questions in the review forms guide panelists to create reviews that convey essential information in a concise way and are straightforward to integrate into panel summaries.</p> <p>The committee agreed that these time-tested methods for reviewing proposals are appropriate and effective. A suggestion to CISE is an echo of the 2014 COV report. In-person panels provide high-quality interactions. Nevertheless, the cost of travel in terms of time and significant environmental impact of travel emissions have prompted several conferences in the CISE CCF community to move to virtual program committee meetings. In this context, we encourage NSF CISE to expand their use of virtual panels whenever feasible.</p>	
<p>2. Are both merit review criteria addressed</p> <ul style="list-style-type: none"> a) In individual reviews? b) In panel summaries? c) In Program Officer review analyses? <p>Comments:</p> <ul style="list-style-type: none"> a) Most reviews make concrete comments about both intellectual merit and broader impact, though comments about the former are more comprehensive and richer than comments about the latter. The committee found some individual reviews did not mention broader impacts or comment on these aspects of the proposals. The committee noted that what constitutes broader impact is not well understood by the community. In many proposals the claims of broader impact are rather generic (e.g., will solve research problem, will create a course, will recruit diverse students); and most reviewers accept that as being reasonable. A few researchers are more thoughtful in considering the possible impacts of their research – on their research community, on society, and on their students. The committee also noted that program-specific criteria (like collaboration plan, mentoring plan, etc.) are often not addressed in individual reviews. b) Panel summaries are generally comprehensive, commenting on both intellectual merit and broader impact. The lowest-rated (NDP) proposals are not discussed by panels and therefore do not receive panel summaries. Instead, the summary reiterates the review criteria and states that panel agrees that these were not met and the program director concurs. The committee agreed this approach was an appropriate way to manage reviewer and program director workload. c) Review analyses of proposals where funding is approved are comprehensive in assessing both intellectual merit and broader impact. Review analyses of proposals where funding is declined tend to rely 	<p>YES</p>

more on boilerplate text. Recommendations for improving review analyses appear in Section I.5 below.

In all three of the above, the committee observed that there was often little correlation between the strength of the broader impacts statement and the final funding decision. Moreover, there seems to be confusion among both the PIs and the reviewers what role it is supposed to play in the evaluation process.

<p>3. Do the individual reviewers giving written reviews provide substantive comments to explain their assessment of the proposals?</p> <p>Comments:</p> <p>Most reviews are very detailed, providing good explanations of what the reviewer thought were strengths and weaknesses of the proposal, and also offering constructive feedback. Several reviews are exceptionally insightful. Other reviews are sufficient for making a funding decision, but they are fairly terse, in terms of useful feedback to the proposer in the event the proposal is not funded. It is particularly important to provide extensive feedback to new PIs and to PIs of CAREER and CRII proposals to clarify what is expected of a high-quality proposal. We recommend that panelists be mindful of this during discussion. On occasion, a proposal will receive reviews that are thin and unsubstantiated, particularly for lower-ranked proposals. These lightweight reviews might be related to the quality of proposals, but such potential correlation might have other sources – e.g., it is possible that such reviewers have insufficient expertise. The committee recommends CISE add mechanisms to aid in identifying such reviews as soon as possible so that ad-hoc reviews can be requested; for example, an expertise rating for the reviewer could be added to the review template.</p>	<p>YES</p>
<p>4. Do the panel summaries provide the rationale for the panel consensus (or reasons consensus was not reached)?</p> <p>Comments:</p> <p>For the proposals discussed by a panel, the panel summaries provide a compilation of the comments of the reviews and account for divergent opinions. In some cases, the panel summaries provide more detail and insight than the sum of the reviews, presumably because of issues that come to light during the panel discussion. However, since panel summaries consist largely of remarks extracted from the individual reviews, they sometimes suffer from the same weaknesses of individual reviews noted in the response to the previous question.</p> <p>The committee was initially concerned that proposals with an NDP rating did not receive a panel summary. We learned from CCF division leadership that all programs other than CCF Small Core, including CRII and CAREER, are discussed in panels and receive a proposal summary. In addition, even for NDP-rated proposals, the program directors ensure that individual reviews provide sufficient feedback. The committee agreed that this approach achieves the appropriate balance between quality feedback to the proposers and management of workload for panelists and program directors.</p>	<p>YES</p>

5. Does the documentation in the jacket provide the rationale for the award/decline decision?

YES

Comments:

Generally, the Review Analysis is supposed to provide rationale for the award or decline decision. In some cases, the CCF committee observed decisions that required the program director to argue against a panel's recommendation (e.g., funding a Low Competitive proposal or not funding a Highly-Competitive proposal), and such decisions were thoroughly justified in the Review Analysis. The committee also observed decisions where Competitive proposals were funded over other Competitive proposals in the same pool. While some of these were well justified, the decision rationale were not consistently documented well, even in cases where proposals were funded. The CISE response to the 2014 COV report mentioned that explanations of other factors impacting decisions would be provided in the Context Statement, but in the jackets reviewed by the committee, these consist of boilerplate text listing the factors that may be considered (balancing subfields, available funds, and NSF policies). The committee did not observe Review Analyses that mentioned specific factors being part of a decision on a specific proposal.

As a result, there is a large gap between the panel summaries and proposal rankings (which are just recommendations to the program director) and the eventual funding recommendations (after aggregating proposal evaluations from multiple panels, after discussion with other PDs in the program cluster, after considering a number of other factors such as geographic diversity, the PI comes from an underrepresented population, diversity of research problems and directions funded, etc.). Because the rationale for how a proposal progressed from the review-panel recommendation to a final-funding recommendation/decision are not consistently documented, the COV panel (and NSF) does not have the data to really answer this question. There is research that suggests that without a rubric for the criteria for evaluating submissions, it is easy to introduce bias in the decision process.

In order to (1) increase the documentation of funding decisions, (2) mitigate against increasing the workload of PDs, and (3) reduce the potential for bias in the decision process*, the COV recommends adding a checkbox-like template to the Review Analysis that would allow PDs or division directors to note factors that played a **strong role** vs. **weak role** vs. **no role**, and that **raised** competitiveness of a proposal or **lowered** it. Example candidate checkboxes to consider include

- The ranking of a proposal in the overall discussion of all proposals being considered together by program clusters
- Diversity of research problems or solution approaches
- Degree of innovation or potential for transformation
- Strength of proposal's Broader Impact statement
- Institutional or geographic diversity of funded proposals
- Consideration that the PI is a member of an underrepresented population
- EPSCoR
- EAGER

- | | |
|--|--|
| <ul style="list-style-type: none">• New/junior PI• Historic or current level of funding that the PI is holding• Other factors that have historically gone into funding recommendations | |
|--|--|

The checkboxes can be text that the RA copies into the form in order to experiment with different factors to be recorded.

<p>6. Does the documentation to the PI provide the rationale for the award/decline decision?</p> <p>Comments:</p> <p>Documentation to the PI seems to consist primarily of the panel reviews, panel summaries and context statement. Where reviews and panel summaries are comprehensive, this information appropriately conveys the rationale for the award/decline decisions. In cases where the best description of the rationale was in the Review Analysis, it appears that this information was not shared with the PI. The PO comment field in the jacket could be used for this purpose. The feedback to the PIs, in cases where the award is declined, often lacks sufficient feedback to help the proposers reformulate their research strategies and write more compelling proposals. The PIs on declined proposals should be encouraged to communicate with their program directors if there is any question about the reasons for the declination.</p>	<p>YES</p>
<p>7. Additional comments on the quality and effectiveness of the program's use of merit review process:</p> <p>There is confusion among PIs and reviewers about what Broader Impacts means and what comprises suitable Broader Impacts for a proposal. The committee recommends efforts to better communicate through improving language about Broader Impacts in the solicitations, and providing supplementary information on Broader Impacts; e.g., referring PIs to cisebroaderimpacts.org. The committee noted that the current CCF core solicitation appears to contribute to the confusion; it mentions BI eight times, and does not start to define it until the last mention. Further, this definition of BI is spread across two paragraphs, with an intervening paragraph listing five review elements for both Intellectual Merit and Broader Impacts. It was noted during the BI breakout session that this defining language was added in 2012 and is mandated, but perhaps moving the definition above its uses, and placing the definition paragraphs together followed by the review elements would improve the clarity.</p> <p>The committee thought the more focused discussion of BI in Section II of the PAPPG was clearer, which preceded the discussion of Merit Review Criteria in Section III. As a suggestion, in addition to examples of what could be BI, a list of what is not BI might help PIs steer away from common mistakes, e.g., "I will teach my standard classes".</p> <p>Communication with PIs, reviewers, and program directors is crucial, and we encourage NSF to keep up its current communication efforts (communication to potential PIs, reviewer training, program director reboot camps), including the above recommendations about language and focus.</p> <p>Additional criteria, such as education components of proposals, postdoctoral mentoring plans, and data management plans, are not systematically addressed in individual reviews, panel summaries, and review analyses. If</p>	

these are important to the decision or to PI feedback, then they should be reviewed more carefully.	
---	--

II. Questions concerning the selection of reviewers. Please answer the following questions about the selection of reviewers and provide comments or concerns in the space below the question.

SELECTION OF REVIEWERS	YES, NO, DATA NOT AVAILABLE, or NOT APPLICABLE
<p>1. Did the program make use of reviewers having appropriate expertise and/or qualifications?</p> <p>Comments:</p> <p>The reviewer mix was well balanced, including some experts in the specific area and some somewhat remote to give an external perspective. Review panels included a diversity of viewpoints. In technology-related proposals, reviewers from industry were utilized. Some reviewers were from international institutions and government labs. In interdisciplinary proposals, reviewers came from a mixture of the relevant disciplines. Generally, the reviews that most panelists provide are detailed and are evidence of their appropriate expertise. In the uncommon cases where reviews are lightweight, it is impossible to determine the cause, but may be because the panelist did not have the depth of expertise required. In such cases, ad-hoc reviewers were sometimes invited to evaluate the projects, and their evaluations were taken into account in the review analysis and final decision.</p>	YES
<p>2. Did the program recognize and resolve conflicts of interest when appropriate?</p> <p>Comments:</p> <p>In reviewing the proposal jackets, conflicts-of-interest were regularly identified. In unusual cases, conflicts were detected after a panelist was assigned, and these reviews were marked “unreleasable”. In those cases, the program director was careful to take all appropriate steps to prevent the conflict from influencing the ultimate funding decision.</p>	YES
<p>3. Additional comments on reviewer selection:</p> <p>In reviewing jackets, the committee sometimes observed insufficient diversity of panels, with low participation of women and other underrepresented groups among the panelists. CISE must continue to strive for the composition of its panels to match the composition of the underrepresented communities it serves.</p>	

A program director's network is currently considered to be important for efficient selection of qualified reviewers. However, building a panel relying on such a network has the potential to introduce subconscious bias.

To address both these concerns, the committee makes several recommendations: (1) increase the pool of reviewers through building a database that is used by the program directors in identifying panelists and reviewers; (2) establish best practices on how to use the database to derive the desirable balance of panelists as well as reviewers for a specific proposal; and, (3) employ additional technology such as an automatic recommendation system to suggest reviewers that may not already be part of the program director's network, and assist in creating panel diversity.

To analyze how well the collection of reviewers for a proposal meet the best practices, it is important to collect data about who the reviewers are: demographics, home institution types, geographic location, etc.

III. Questions concerning the management of the program under review. Please comment on the following:

MANAGEMENT OF THE PROGRAM UNDER REVIEW
<p>1. Management of the program.</p> <p>Comments:</p> <p>Overall the program is very well managed. It receives about 1700 (+/- 100) competitive proposals yearly. More than 70% of the proposals have dwell time of at most 6 months. This meets the GPRA performance goal of making award/declination decisions on 70% of submissions within six months of either receipt of the proposal or applicable solicitation due date. Yet, in 2014-2015 this number was nearly 90%, so it appears that in recent years there is an increase in dwell time.</p> <p>Almost all proposals that receive external reviews are discussed in a merit review panel.</p> <p>The committee is aware that a promising proposal is sometimes held longer in case funding becomes available a little later in the fiscal year, which may account for some of this increase as funding pressure on CISE has increased. Even though CCF is outperforming its dwell time goal, we recommend NSF continue to monitor to prevent dwell time from slipping below the goal.</p>
<p>2. Responsiveness of the program to emerging research and education opportunities.</p> <p>Comments:</p> <p>CCF, in conjunction with other divisions in CISE, has run workshops both on emerging technologies (quantum computing, synthetic biology) and ways to apply computing technology in the real world (formal methods for security, advanced manufacturing). In addition, the Computing Community Consortium (CCC) serves as a mechanism to gather insights into important research directions from the computing community.</p> <p>CCF runs targeted programs in technology areas (parallel computing, heterogeneous architectures, microarchitectures, applied algorithms, formal methods), as well as applications of computer technology (sustainability, data science). Some of these are in conjunction with industry partners. CCF also participates with the rest of CISE in targeted programs (SaTC, Expeditions)</p> <p>Overall, there seems to be an appropriate balance of fostering new research priorities, while being responsive to research directions set by the PIs themselves, with proposals handled as part of the core program.</p>
<p>3. Program planning and prioritization process (internal and external) that guided the development of the portfolio.</p>

Comments:

CCF, like other divisions of CISE, is primarily supported by its program directors (PDs). However, PDs also collaborate across divisions and directorates to support cross-cutting programs within CISE such as Expeditions in Computing (EiC), Algorithms in the Field (AitF), Exploiting Parallelism and Scalability (XPS), and across directorates, such as Biological and Computing Shared Principles (BCSP), Cyber-Physical Systems (CPS), and Secure and Trustworthy Cyberspace (SaTC).

CCF seems to take action and react promptly in adjusting its portfolio of programs depending on the scientific context and needs. Focused scientific workshops are held regularly that help inform about trends and needs, like the Workshop on Formal Methods and Security (in 2016), the Workshop on Exploiting Parallelism and Scalability (in 2015), and the CCC document on the next steps of Quantum Computing (by Margaret Martonosi and Martin Roetteler). As a result, the XPS program evolved into Scalable Parallelism in the eXtreme (SPX), and Formal Methods in the Field (FMitF) was added in 2017.

Quantum Computing and Computational Biology were added as separate topical areas to the core programs in 2017. It was explained to the committee that these two areas are not separate programs, but are identified internally by separate funding codes for administrative purposes. In FY17 and FY18, proposals in these two areas are submitted to one of the other core programs AF, CIF or SHF. Starting in FY19, they have been part of the FET program, both internally and externally.

4. Responsiveness of program to previous COV comments and recommendations.

Comments:

The 2015 Response to the 2014 CISE/CCF COV report shows that CISE takes the COV recommendations quite seriously. The responses give justification for the actions taken and not taken (though in some cases we believe that the recommendations need to be considered again).

Examples of 2014 COV recommendations that have been addressed:

1) The 2014 report recommended that future COV committees have advanced access to information needed to complete their reports, including advance access to eJacket, and an early detailed description about the NSF review process. All recommendations were implemented as part of the 2019 COV review process.

2) The 2014 COV panel raised a concern that the funding rate of large proposals seemed to be declining. CISE responded that they would monitor the program, taking into account the value of large projects with respect to portfolio balance, and would engage the PI community in best practices for proposing large projects. In 2018, CCF discontinued accepting large proposals while they reconsider their value. Given the large awards granted to those proposals, this seems prudent.

3) There was a concern that CCF had difficulty recruiting high-quality panelists in some programs, in particular interdisciplinary programs. 2014 COV recommendations included (1) tool support to manage a database of potential reviewers, track their performance, enable automatic searches of panelists based on expertise; (2) greater use of ad-hoc reviews for proposals that don't fit well into any panel topic; and, (3) greater use of virtual or hybrid panels to increase the involvement of panelists who are unable or unwilling to travel. NSF has experimented with, and makes use of, all of the above. In-person panels result in the highest-quality review decisions, but PDs do use

virtual/hybrid panels or ad-hoc panelists rather than lose out on good reviewers. Some tools have been rolled out to assist with panel formation: SurveyMonkey to recruit reviewers for the reviewer database and SuggestReviewer to identify reviewers within the database. Tools tend to solve siloed problems and their data is not integrated, so PDs have to consult with multiple and sometimes overlapping tools, like both FastLane and Research.gov to manage their program and proposal portfolios.

4) There was a concern about the onboarding of new Program Directors since they do not always start at the same time and about keeping the training of permanent staff fresh and up-to-date. CISE responded that a Program Manager Survival Guide was being revised, and that the NSF Academy was running training on merit review and NSF policies in the winter-spring timeframe as well as in the Fall; and that it has revived a long-term professional development program for permanent staff. CISE also pointed out that the Program Directors work in clusters, and that new PDs shadow experienced ones in part to facilitate mentoring and the transfer of knowledge.

Examples of 2014 COV recommendations that were not addressed, but where CISE justifications for not taking action make sense:

5) The 2014 COV noted that the recruiting of high-quality rotators could be hampered by the constraint that PDs cannot submit NSF proposals while serving on staff. They suggested some workarounds, such as providing automatic extensions to those who have expiring awards or creating a restart grant for rotators once they return home. The CISE decided against acting on either recommendation, identifying alternative workarounds (e.g., securing non-NSF funding, being a co-PI on an NSF grant where someone else serves as the submitting and negotiating PI). Given NSF's strict policies on conflict of interest and against any change that might introduce bias into the review process, this decision makes sense.

6) As a means to improve review quality, it was suggested that reviewers could provide more informed feedback on proposals if they had access to a PI's funding history, including panel feedback on past related proposals. Again, the CISE demurred saying that past reviews could bias reviewers' opinions, which should be based only on the contents of the current proposal. Worse, it would be impossible to guard against revealing the identities of past reviewers, which could violate conflict-of-interest rules.

7) The COV noted that some Review Analyses addressed only the review process and outcome, and not intellectual merit. The CISE mentioned that it would be standard practice to address intellectual merit, except for low competitive or not recommended proposals, where (for workload reasons) standard practice is to rely on reviews and panel summary. This process makes sense, as long as there is an informative panel summary; it may need to be revisited for cases where a proposal is not discussed in a panel.

Examples of long-standing issues:

8) The high workload of Program Directors and CCF staff seems to be a perpetual problem. CCF did increase the staff complement after the 2014 COV. But the number of proposals to CISE has increased by 30% over these same 6 years (7000 to 9000). In addition, the number of cross-cutting programs has increased (from 3 CCF-led and 15 CCF-participating programs in 2014 to 4 CCF-led and 20 CCF-participating programs in 2018). Because of the extra overhead associated with managing proposals in cross-cutting programs (e.g., coordination with other divisions, directorates, organizations), this increase in programs is a greater increase in workload than it might look to be on paper.

To mitigate against the increased workload, CISE has experimented with a number of innovative approaches and efficiencies to reduce workload, such as

- limiting each PI to a small number submissions per year in some programs
- increasing use of review templates and panel summary templates
- triaging proposals and not discussing in panels (NDP) those proposals ranked very low by reviewers
- not providing substantive review analyses or panel summaries for NDP proposals
- using virtual/hybrid panels and ad-hoc reviewers to make it easier to recruit reviewers

The PDs make use of these where they make sense (ease the task of forming a panel, improve the quality of reviews).

One remaining recommendation to study is to eliminate the submission deadlines for core programs, which has the potential to reduce workload for program directors, assist principal investigators in submitting on their own timeline, and consequently, improve proposal quality. This rolling submissions of proposals represents a substantive change, and CISE will need to develop new processes before jumping to rolling submissions. A detailed analysis of rolling submissions can be found in Question 1 under Other Topics at the end of this document.

9) Achieving consistently high-quality reviews for all proposals is always a challenge for any organization seeking reviews. CISE does amazingly well in recruiting panelists who are committed to a high-quality process and who overall produce comprehensive, insightful, and critical reviews of proposals. But it does occasionally happen that a proposal is not well reviewed. The 2009 and 2014 COV panels both made recommendations that aimed to mitigate against weak reviews due to lack of reviewer expertise, in particular (1) asking reviewers to rate their expertise and confidence in their ability to review a proposal **after** completing their review rather than asking this question pre-review, based only on a reviewer's reading of a proposal's title or abstract; (2) insisting that reviews be submitted well in advance of the panel date, so that (3) last-minute ad-hoc reviewers can be solicited in cases where the set of reviews for a proposal is not sufficiently informative. CISE has consistently responded that the PDs use their own knowledge of reviewer expertise and the reviewers' pre-review self-assessment of expertise. During the 2019 COV visit, CISE PDs mentioned that they can easily determine reviewers' expertise during the panel and, in cases where a proposal is lacking any qualified review, can go out for an ad-hoc review after the panel to solicit an expert review; and can reconvene the panel virtually to discuss the expert review.

10) Both the 2009 and 2014 COV panels made multiple recommendations of IT solutions to assist Program Directors with panel formation, reviewer performance, confidence ratings, data analytics of submissions, portfolios, PIs, and funded institutions. CISE mentioned a number of tools that have been rolled out to help PDs: *SurveyMonkey* to populate the reviewer database and *SuggestReviewer* to help automate some aspects of the review-panel process. In general, it seems NSF uses mostly homegrown tools to ensure that their data is secure, which limits how quickly they can develop and deploy new tools. Other problems are (1) the homegrown tools tend to solve one problem in isolation, and do not integrate with the rest of the PDs tool suite; and, (2) the tools are developed and PDs are consulted after the fact with focus groups. Instead, the CCF panel recommends more lightweight prototyping of tools to get PD feedback before full implementation.

IV. Questions about Portfolio. Please answer the following about the portfolio of awards made by the program under review.

<p style="text-align: center;">RESULTING PORTFOLIO OF AWARDS</p>	<p style="text-align: center;">APPROPRIATE, NOT APPROPRIATE, OR DATA NOT AVAILABLE</p>
<p>1. Does the program portfolio have an appropriate balance of awards across disciplines and sub-disciplines of the activity?</p> <p>Comments:</p> <p>Over the years FY14 to FY17, between 57% and 64% of CCF funds went to core programs, with the remainder going to cross-cutting programs shared with other divisions.</p> <p>There has been almost no growth in the total CCF budget over the past five years, which is a cause for concern.</p> <p>Funding for core programs has been flat for the last five years, which is a cause for some concern. Of the three main core programs AF, CIF, and SHF, the last tends to receive about half the core funding, with the remainder split evenly between AF and CIF. This split has been fairly consistent over the last five years. The reason for this, as stated in the annual reports, is due to the large and comprehensive nature of the SHF cluster, whereas AF and CIF tend to be more focused. In FY17, the Quantum Computing and Computational Biology clusters were introduced.</p> <p>The cross-cutting programs are highly diverse, with a noticeable increase in the number and diversity of programs over the past five years.</p> <p>Success rate of core proposals, which was previously stable, has declined in FY17 and FY18. These figures are essentially uniform across all CCF clusters. The success rate for cross-cutting programs is considerably less and varies widely, with some declines due to increased proposal pressure, notably in the CyberSEES program.</p> <p>Most troubling is the decline in success rate of CAREER proposals. NSF promotes CAREER awards as “prestigious awards in support of early-career faculty.” Indeed, this prestige is well recognized within universities, and the receipt of a CAREER award has become an important consideration in tenure cases. As noted in the reports, CCF aims to fund at least 30% of CAREER proposals to ensure a healthy research pipeline. Unfortunately, the success rate has declined steadily from a high in FY14 to lower rates in FY 17-18. This is an alarming trend. It is largely due to an increase in number of proposals in the last two years. CCF claims to have increased the funding for CAREER awards in response, but more needs to be done.</p>	<p style="text-align: center;">MOSTLY APPROPRIATE</p>

<p>Assuming a fixed budget for CAREER awards, one possibility would be to reduce the amount of each award, e.g., by reducing their duration from 5 to 3 years. Although this would reduce their effectiveness in helping young faculty establish their research careers, the fact is that many applicants wait until they are close to the end of the eligibility window before proposing in order to maximize their prospects. Thus, such awards are not really serving this purpose.</p>	
<p>2. Are awards appropriate in size and duration for the scope of the projects?</p> <p>Comments:</p> <p>Over the past five years, the number of competitive projects funded by CCF and average duration of the projects were remarkably consistent. The number of projects funded varied between 342 and 378; the average duration varied between 2.7 and 2.84 years. In each year, certain special programs tended to have longer four-year durations: CyberSEES in FY14, XPS and AiTF in FY15-16, AiTF in FY17, and FMitF in FY18. Most projects fall in the Small range with award sizes of \$100,000-\$150,000 and \$150,000-\$250,000 (annualized). The CoV regards this consistency as a sign of stability.</p> <p>It was observed in the 2014 COV report that awarded budgets average 91% of the requested budgets, and over 90% of awards receive at least 80% of the requested funds. It was also noted with some concern that there was a significant decline in the funding rate of Large proposals over the past four years. Large proposals were suspended in 2018 while CCF reevaluates their value; they consume a lot of funding and were rarely being rated as HC.</p>	<p>APPROPRIATE</p>
<p>3. Does the program portfolio include awards for projects that are innovative or potentially transformative?</p> <p>Comments:</p> <p>CISE employs processes that actively seek to elicit innovative or transformative proposals. In the review templates, one of the five elements to be considered is whether the “activities suggest and explore creative, original, or potentially transformative concepts.” This information is conveyed to principal investigators via the PAPPG.</p> <p>Beyond the process, the committee identified the extent to which jackets we reviewed explored new research priorities mentioned in question 10 below (exascale computing, quantum, security, privacy, AI, applications areas). Most of these new areas were included in the jackets we reviewed.</p>	<p>APPROPRIATE</p>
<p>4. Does the program portfolio include inter- and multi-disciplinary projects?</p> <p>Comments:</p>	<p>APPROPRIATE</p>

<p>The portfolio includes several multi-disciplinary programs, exemplified by the computational biology initiative that brings together two distinct disciplines. In addition, some programs, like TRIPODS and CyberSEES are actually designed to foster Transdisciplinary Research. Sampled TRIPODS jackets suggested these are typically in rather closely-related fields like interfaces of mathematics, statistics, and theoretical computer science. CyberSEES includes several projects that bridge more distant fields, for example computer science and geoscience, computer science and ecology, or computer science and synthetic biology.</p>	
<p>5. Does the program portfolio have an appropriate geographical distribution of Principal Investigators?</p> <p>Comments:</p> <p>The analysis is based on the Annual Reports from 2015 to 2018 since the Annual Report for 2014 seems to have much higher numbers than the rest of the reports, making comparisons difficult.</p> <p>There is a significant difference in the award ratio between EPSCoR and Non-EPSCoR states (in favor of Non-EPSCoR states), based on NUMBER of awards.</p> <p>It might well be the case that the award funding allocation per state is correlated with the number of students per state, but insufficient data was given to the committee to verify this. In general, it seems that more accessible data is necessary to make a thorough analysis whether the geographical distribution is appropriate.</p> <p>Some program directors mentioned that they use the funding dedicated to EPSCoR states to increase the total number of awards and not specifically increase the number of awards to EPSCoR states, while other program directors mentioned that they do use the EPSCoR funding to award borderline proposals from EPSCoR states. We do not have data to analyze what fraction of the EPSCoR funding was used for each of these goals. We recommend that the role of EPSCoR in funding decisions be documented so that such an analysis can be done.</p>	<p>INSUFFICIENT DATA TO MAKE A DETERMINATION</p>
<p>6. Does the program portfolio have an appropriate balance of awards to different types of institutions?</p> <p>Comments :</p> <p>It is appropriate for NSF to consider the balance of awards or amount of funding across different types of institutions.</p> <p>The top 100 PhD-granting research-intensive universities receive the bulk of the funding (out of 250 PhD-granting programs in CS (195), CE (34), and IS (24)). 2014 : 85.6% (top 100), 9.6% (other PhD institutions) 2015 : 82.0% (top 100), 12.3% (other PhD institutions)</p>	<p>APPROPRIATE</p>

<p>2016 : 84.3% (top 100), 10.2% (other PhD institutions) 2017 : 82.1% (top 100), 12.5% (other PhD institutions) 2018 : 85.1% (top 100), 9.3% (other PhD institutions)</p> <p>These percentages do not seem out of line. The top 100 PhD-granting research-intensive institutions are training the bulk of the PhD students who are pursuing their doctoral studies in the US, thus it makes sense for those institutions to be awarded the bulk of the funding. The COV panel did not have information on what percentage of all PhD students are studying at the top 100 research institutions, but 85% (or more) seems to be very believable.</p>	
<p>7. Does the program portfolio have an appropriate balance of awards to new and early-career investigators?</p> <p>Comments:</p> <p>New investigators within CCF have been funded at a lower but not significantly lower rate than returning investigators. Over the period of interest, new investigators have been mostly funded at a rate comparable to the general population of all investigators</p> <p>Regarding the number of proposals submitted by new investigators, it is fairly stable, in comparison to the number of proposals by returning investigators.</p> <p>More specifically, looking into the various programs of CCF, although the largest numbers of submissions by new investigators are usually observed in the core program, the imbalance of funding is more stringent in some of these core programs. For AF in particular, there are some significant variations across the multiple years reviewed by the COV. On the other hand, the core SHF program more consistently achieves even (or close to even) funding rates (2014, 2015, 2018).</p> <p>Looking at cross-cutting CCF programs, we observe that CyberSEES and Exploiting Parallelism & Scalability were fairly effective at attracting new investigators in 2014-2016, with the largest numbers of submissions by new investigators observed in these programs specifically these years.</p> <p>CRII effectively attracts a healthy number of new investigators' submissions. Concerns were raised during the COV discussions that CRII proposals were disproportionately going to faculty at top institutions where faculty already have available resources from their university, and changes have been implemented to use CRII as a tool to better assist new faculty where startup funding is less available.</p> <p>CAREER attracts a good and increasing number of proposals. However, we are concerned with the sharp decline in CISE-wide CAREER success rates over the analysis period. This is discussed in more depth in IV.1.</p>	<p>APPROPRIATE</p>
<p>8. Does the program portfolio include projects that integrate research and education?</p>	<p>APPROPRIATE</p>

<p>Comments:</p> <p>For the most part, the only proposals that require an education component are the CAREER awards, which are designed to support early faculty in initiating their research and building new curriculum around it. While the number of CAREER proposals awarded has increased during the period of review, funding rates have not kept pace with the dramatic growth of faculty resulting from increased demand for computing-related degrees and computing professionals. This reduction in CAREER funding rates poses a risk for initiating education programs in new areas, and negatively impacts the success of new faculty in establishing their careers.</p> <p>Beyond CAREER, most recent proposals in their Broader Impacts plans describe educational plans that migrate the research developed by the project into classes at their institutions or tutorials in their research communities. This component of Broader Impacts offers a path to develop and disseminate new curricular materials in new research areas. Such an approach is a cost-effective way for CISE to support education in cutting-edge research areas of interest and incentivizes curricular upkeep.</p>	
<p>9. Does the program portfolio have appropriate participation of underrepresented groups³?</p> <p>Comments:</p> <p>Determining what level of participation by underrepresented groups is “appropriate” requires considerations far outside the scope that can be addressed by the COV. In addition, the data available is at a fairly high level, and many of the categories of interest have such low numbers that a statistical analysis would not be meaningful.</p> <p>Combining data from the CCF reports for the years 2014–2018 provides some insights into the funding rates overall versus for specific groups. The rate for women seems to indicate that they are getting an appropriate consideration. The rate for minorities is well below the overall average and bears investigation. One problem is that the numbers are very small. The success rate for new proposers is also below the average, but one might expect inexperienced proposers to have a lower success rate. The success rate for those with disabilities is well below average, but the numbers are so low that the statistics are not very meaningful.</p>	<p>INSUFFICIENT DATA TO MAKE A DETERMINATION</p>
<p>10. Is the program relevant to national priorities, agency mission, relevant fields and other constituent needs? Include citations of relevant external reports.</p>	<p>APPROPRIATE</p>

³ NSF does not have the legal authority to require principal investigators or reviewers to provide demographic data. Since provision of such data is voluntary, the demographic data available are incomplete. This may make it difficult to answer this question for small programs. However, experience suggests that even with the limited data available, COVs are able to provide a meaningful response to this question for most programs.

<p>Comments:</p> <p>The CCF portfolio includes research areas relevant to national priorities and relevant fields, based on relevant external reports.</p> <p>The National Strategic Computing Initiative (July 2015) called for (1) algorithms, networking, hardware and software ecosystems; (2) exploiting similarities between scientific simulation and data analytics; and (3) industry-government-academia cooperation, to provide a viable path forward Post-Moore's Law Era. These goals are supported by the following programs:</p> <ul style="list-style-type: none"> • Scalable Parallelism in the Extreme (SPX) solicitation makes direct reference to the NSCI Executive Order, and addresses item (1). • NSF/Intel Partnership on Computer Assisted Programming for Heterogeneous Architectures (CAPA) addresses items (1) and (3). • Transdisciplinary Research in Principles of Data Science addresses item (2). <p>The committee identified proposal jackets from these programs that were responsive to NSCI.</p> <p>The National Quantum Initiative Act (2018) calls for multidisciplinary curriculum and research to fill the knowledge gap for quantum computing. The goals of NQIA are addressed by the Quantum Information Science program code under Foundations of Emerging Technology, and the committee identified proposal jackets that were responsive.</p> <p>The Cancer Moonshot Blue Ribbon Panel Report (2016) identified 10 required research areas, which could potentially be supported by the Computational Biology program code under Foundations of Emerging Technology.</p>	
<p>11. Additional comments on the quality of the projects or the balance of the portfolio:</p> <p>At the partnership breakout group, there was a mention about joint NIH/NSF panel. It is unclear how this is working right now, but it appears to be a good way to increase synergy between the quickly changing landscape of practical questions enabled by new types of data and foundational research to support them.</p>	

OTHER TOPICS

1. Please comment on any program areas in need of improvement or gaps (if any) within program areas.

CCF is seriously considering moving to rolling submissions as a way to address the anticipated increase in proposal pressure in the near term. This would be a major step, as it would involve considerable retooling of the review mechanism. Other NSF units have tried this recently and there is some information available regarding its advantages and disadvantages.

This move would bring immediate benefits and costs, both significant. The chief benefit that has been observed by those units that have gone this route is an immediate drop in the number of submissions, apparently by as much as 50%, followed gradually by a rebound back to 75%. If these numbers are accurate, this would be a significant reduction in workload for NSF staff. It is conjectured that the reason for this effect is simply that the forcing function is removed. Another benefit seems to be an improvement in the quality of the submitted proposals, although the evidence for this effect is so far anecdotal.

The major cost would be the revamping of the review mechanism. The panel system would have to move to a rolling schedule. This would involve waiting for enough submissions to convene a panel, which could negatively effect dwell time. There may also have to be a mechanism for dealing with orphan proposals.

The transition may also require creativity in managing the annual budget, which currently presumably depends on hard deadlines for accounting purposes or alignment with the fiscal year. With rolling submissions, this dependence would somehow have to be relaxed. Perhaps a workable compromise would be a more frequent deadline schedule, say quarterly, which would lessen the forcing function and spread the work out over the year while still allowing budgeting projections.

If this move is to be seriously considered, the decision should be made on the basis of hard evidence from the units that have gone to rolling deadlines. There should be data available about drop in submissions, average quality of submissions before and after, and possible negative effect on dwell time. The units that have implemented rolling submissions should be able to provide advice on how best to implement it.

Apart from the benefit to NSF, there would be a major benefit to university grant administrators as well, as it would alleviate the crunch around deadline time. PIs will also benefit as deadline flexibility allows them to select the most appropriate time in their schedule to develop their proposal.

2. Please provide comments as appropriate on the program's performance in meeting program-specific goals and objectives that are not covered by the above questions.
3. Please identify agency-wide issues that should be addressed by NSF to help improve the program's performance.
4. Please provide comments on any other issues the COV feels are relevant.
5. NSF would appreciate your comments on how to improve the COV review process, format and report template.

The COV felt that in general the process is very well designed and there is not much of substance to recommend. The information provided to the COV (annual reports, jackets) is complete and fully adequate for the COV to accomplish its task. However, we do have a few minor suggestions for improvement that would make the process a little easier. These mostly concern the presentation of the data.

There are many acronyms and it is sometimes difficult to keep track of them all. There is a glossary provided, but it was rather difficult to find, and some items were missing completely (e.g. EPSCoR), and there were no pointers to where the acronyms were defined. Also, there is some inconsistency in the labeling of the tables: sometimes acronyms are used, sometimes the names are spelled out. It would be good for the sake of consistency to use only acronyms. In a perfect world, hovering over

an acronym would bring up a text box with the definition, but this is probably too much to ask for.

In the annual reports, many of the tables are summarized in the early part of the report, and the tables presented later on. It might be better to have the summaries and the tables together in the same location to avoid having to page back and forth.

Some of the tables are presented as raw numbers and unsorted, making comparisons difficult. It might be helpful to have them presented in a more visual form, such as bar charts, or at least sorted so that comparisons can be made.

Some tables are given as images (.png, etc). This makes it impossible to copy and paste for the purpose of doing any data analysis. To do any such analysis, the reviewer must type it in. It would be good to have the tables given in a form that allows the text to be selected and copied/pasted.

The data and reports from several years are in separate files. For the purpose of analyzing trends, it would be good to have bar charts with several years' data in the same chart to compare. This would be especially important for analyzing trends in funding levels and success rates for the various programs, CAREER awards, etc. over a period of several years.

Although the CoV was provided copies of previous annual reports, it would be helpful if that data could be aggregated in the form of general trends. As it was, members of the CoV did this individually, and some were provided to us by NSF staff in response to explicit requests. Having the NSF anticipate this need in advance would make the CoV review more efficient and effective.

The Committee of Visitors is part of a Federal advisory committee. The function of Federal advisory committees is advisory only. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the Advisory Committee, and do not necessarily reflect the views of the National Science Foundation.

SIGNATURE BLOCK:



For the CCF COV Sub-committee
Mary Hall
Chair

CORE QUESTIONS and REPORT TEMPLATE for FY 2019 NSF COMMITTEE OF VISITOR (COV) REVIEWS

Guidance to NSF Staff: This document includes the FY 2019 set of Core Questions and the COV Report Template for use by NSF staff when preparing and conducting COVs during FY 2019. Specific guidance for NSF staff describing the COV review process is described in the “COV Reviews” section of NSF’s Administrative Policies and Procedures which can be obtained at <https://inside.nsf.gov/tools/toolsdocuments/Inside%20NSF%20Documents/Policy,%20Procedures,%20Roles%20and%20Responsibilities%20for%20COV%20Reviews%20and%20Program%20Portfolio%20Reviews.pdf>¹.

NSF relies on the judgment of external experts to maintain high standards of program management, to provide advice for continuous improvement of NSF performance, and to ensure openness to the research and education community served by the Foundation. COV reviews provide NSF with external expert judgments in two areas: (1) assessments of the quality and integrity of program operations; and (2) program-level technical and managerial matters pertaining to proposal decisions.

The program(s) under review may include several sub-activities as well as NSF-wide activities. The directorate or division may instruct the COV to provide answers addressing a cluster or group of programs – a portfolio of activities integrated as a whole – or to provide answers specific to the sub-activities of the program, with the latter requiring more time but providing more detailed information.

The Division or Directorate may add questions relevant to the activities under review. Copies of the report template and the charge to the COV should be provided to OIA prior to forwarding to the COV. In order to provide COV members adequate time to read and consider the COV materials, including proposal jackets, COV members should be given access to the materials in the eJacket COV module approximately four weeks before the scheduled face-to-face meeting of the COV members. Before providing access to jackets, the Conflict of Interest and Confidentiality briefing for COV members should be conducted by webinar, during which, NSF staff should also summarize the scope of the program(s) under review and answer COV questions about the template.

Suggested sources of information for COVs to consider are provided for each item. As indicated, a resource for NSF staff preparing data for COVs is the Enterprise Information System (EIS) –Web COV module, which can be accessed by NSF staff only at <http://budg-eis-01/eisportal/default.aspx>. In addition, NSF staff preparing for the COV should consider other sources of information, as appropriate for the programs under review.

For programs using section IV (addressing portfolio balance), the program should provide the COV with a statement of the program’s portfolio goals and ask specific questions about the program under review. Some suggestions regarding portfolio dimensions are given on the template. These suggestions will not be appropriate for all programs.

Guidance to the COV: The COV report should provide a balanced assessment of NSF’s performance in the integrity and efficiency of the **processes** related to proposal review. Discussions leading to answers of the Core Questions will require study of confidential material such as declined proposals and reviewer comments. **COV reports should not contain confidential material or specific information about declined proposals.** The reports generated by COVs are made available to the public.

¹ This document has three parts: (1) Policy, (2) Procedures, and (3) Roles & Responsibilities.

We encourage COV members to provide comments to NSF on how to improve in all areas, as well as suggestions for the COV process, format, and questions. For past COV reports, please see <http://www.nsf.gov/od/oia/activities/cov/>.

FY 2019 REPORT TEMPLATE FOR NSF COMMITTEES OF VISITORS (COVs)

The table below should be completed by program staff.

<p>Date of COV: November 6-8, 2019</p>
<p>Program/Cluster/Section:</p> <p>Core programs</p> <ul style="list-style-type: none">• Computer Systems Research (CSR)• Networking Technology and Systems (NeTS) <p>Cross-cutting programs</p> <ul style="list-style-type: none">• CISE Community Research Infrastructure (CCRI)• CISE Research Infrastructure: Mid-Scale Infrastructure - NSFCloud (CRI: NSFCloud)• Computer Science for All (CSforAll:RPP)• Critical Resilient Interdependent Infrastructure Systems and Processes (CRISP)• Cyber-Physical Systems (CPS)• Enabling Access to Cloud Computing Resources for CISE Research and Education (Cloud Access)• Enhancing Access to the Radio Spectrum (EARS)• Future Internet Architectures -- Next Phase (FIA-NP)• Inclusion Across the Nation of Communities of Learners of Underrepresented Discoverers in Engineering and Science (NSF INCLUDES)• Industry-University Cooperative Research Centers Program (IUCRC)• Japan-US Network Opportunity 2 (JUNO)• NSF/Intel Partnership on Cyber-Physical Systems Security and Privacy (CPS-Security)• NSF/Intel Partnership on Information-Centric Networking in Wireless Edge Networks (ICN-WEN)• NSF/Intel Partnership on Machine Learning for Wireless Networking Systems (MLWiNS)• NSF/VMware Partnership on Edge Computing Data Infrastructure (ECDI)• NSF/VMware Partnership on Software Defined Infrastructure as a Foundation for Clean-Slate Computing Security (SDI-CSCS)• Platforms for Advanced Wireless Research (PAWR): Establishing the PAWR Project Office (PPO) (PAWR/PPO)• Research Experiences for Undergraduates (REU)• Secure and Trustworthy Cyberspace (SaTC)• Secure and Trustworthy Cyberspace: Secure, Trustworthy, Assured and Resilient Semiconductors and Systems (SaTC: STARSS)• Smart and Connected Communities (S&CC)• Spectrum Efficiency, Energy Efficiency, and Security (SpecEES): Enabling Spectrum for All• US Ignite: Networking and Research Application Prototypes Leading to Smart and Connected Communities• US-EU Internet Core & Edge Technologies (ICE-T)• Wireless Innovation between Finland and US (WiFiUS)

Division: Computer and Network Systems (CNS)			
Directorate: Computer and Information Science and Engineering (CISE)			
Number of actions reviewed: 215			
Awards: 97			
Declinations: 112			
Other: 6			
Total number of actions within Program/Cluster/Division during period under review:			
	# of Awards	# of Declines	Total
FY 2014	620	1843	2463
FY 2015	671	1796	2467
FY 2016	658	2166	2824
FY 2017	714	2380	3094
FY 2018	762	2697	3459
Manner in which reviewed actions were selected:			
<ul style="list-style-type: none"> (1) Random selection from bins for programs over the covered period, evenly divided between awards and declinations; (2) Random selection of EAGER, RAPID, travel grant, and workshop awards; and (3) Additional random selection from bins requested by the COV. 			

COV Membership

	Name	Affiliation
COV Chair or Co-Chairs:	Overall CoV Chair: Azer Bestavros CNS Subcommittee Chair: Sonia Fahmy	Boston University Purdue University
COV Members:	Terry Benzel Jill Denner Katherine Guo Anthony Joseph Philip Levis Baochun Li Cristina Nita-Rotaru Ambareen Siraj	University of Southern California ETR Nokia-Bell Labs University of California, Berkeley Stanford University University of Toronto Northeastern University Tennessee Tech University

MERIT REVIEW CRITERIA

An understanding of NSF's merit review criteria is important in order to answer some of the questions on the template. Reproduced below is the information provided to proposers in the Grant Proposal Guide about the merit review criteria and the principles associated with them. Also included is a description of some examples of broader impacts, provided by the National Science Board

1. Merit Review Principles

These principles are to be given due diligence by PIs and organizations when preparing proposals and managing projects, by reviewers when reading and evaluating proposals, and by NSF program staff when determining whether or not to recommend proposals for funding and while overseeing awards. Given that NSF is the primary federal agency charged with nurturing and supporting excellence in basic research and education, the following three principles apply:

- All NSF projects should be of the highest quality and have the potential to advance, if not transform, the frontiers of knowledge.
- NSF projects, in the aggregate, should contribute more broadly to achieving societal goals. These broader impacts may be accomplished through the research itself, through activities that are directly related to specific research projects, or through activities that are supported by, but are complementary to, the project. The project activities may be based on previously established and/or innovative methods and approaches, but in either case must be well justified.
- Meaningful assessment and evaluation of NSF funded projects should be based on appropriate metrics, keeping in mind the likely correlation between the effect of broader impacts and the resources provided to implement projects. If the size of the activity is limited, evaluation of that activity in isolation is not likely to be meaningful. Thus, assessing the effectiveness of these activities may best be done at a higher, more aggregated, level than the individual project.

With respect to the third principle, even if assessment of Broader Impacts outcomes for particular projects is done at an aggregated level, PIs are expected to be accountable for carrying out the activities described in the funded project. Thus, individual projects should include clearly stated goals, specific descriptions of the activities that the PI intends to do, and a plan in place to document the outputs of those activities. These three merit review principles provide the basis for the merit review criteria, as well as a context within which the users of the criteria can better understand their intent.

2. Merit Review Criteria

All NSF proposals are evaluated through use of two National Science Board approved merit review criteria. In some instances, however, NSF will employ additional criteria as required to highlight the specific objectives of certain programs and activities.

The two merit review criteria are listed below. Both criteria are to be given full consideration during the review and decision-making processes; each criterion is necessary but neither, by itself, is sufficient. Therefore, proposers must fully address both criteria. ([PAPPG Chapter II.C.2.d.\(i\)](#) contains additional information for use by proposers in development of the Project Description section of the proposal.) Reviewers are strongly encouraged to review the criteria, including [PAPPG Chapter II.C.2.d.\(i\)](#), prior to the review of a proposal.

When evaluating NSF proposals, reviewers will be asked to consider what the proposers want to do, why they want to do it, how they plan to do it, how they will know if they succeed, and what benefits could accrue if the project is successful. These issues apply both to the technical aspects of the proposal and the way in which the project may make broader contributions. To that end, reviewers will be asked to evaluate all proposals against two criteria:

- **Intellectual Merit:** The Intellectual Merit criterion encompasses the potential to advance knowledge; and
- **Broader Impacts:** The Broader Impacts criterion encompasses the potential to benefit society and contribute to the achievement of specific, desired societal outcomes.

The following elements should be considered in the review for both criteria:

1. What is the potential for the proposed activity to:
 - a. Advance knowledge and understanding within its own field or across different fields (Intellectual Merit); and
 - b. Benefit society or advance desired societal outcomes (Broader Impacts)?
2. To what extent do the proposed activities suggest and explore creative, original, or potentially transformative concepts?
3. Is the plan for carrying out the proposed activities well-reasoned, well-organized, and based on a sound rationale? Does the plan incorporate a mechanism to assess success?
4. How well qualified is the individual, team, or organization to conduct the proposed activities?
5. Are there adequate resources available to the PI (either at the home organization or through collaborations) to carry out the proposed activities?

3. Examples of Broader Impacts

The National Science Board described some examples of broader impacts of research, beyond the intrinsic importance of advancing knowledge.² “These outcomes include (but are not limited to) increased participation of women, persons with disabilities, and underrepresented minorities in science, technology, engineering, and mathematics (STEM); improved STEM education at all levels; increased public scientific literacy and public engagement with science and technology; improved well-being of individuals in society; development of a globally competitive STEM workforce; increased partnerships between academia, industry, and others; increased national security; increased economic competitiveness of the United States; and enhanced infrastructure for research and education. These examples of societally relevant outcomes should not be considered either comprehensive or prescriptive. Investigators may include appropriate outcomes not covered by these examples.”

² [NSB-MR-11-22](#)

INTEGRITY AND EFFICIENCY OF THE PROGRAM'S PROCESSES AND MANAGEMENT

Briefly discuss and provide comments for *each* relevant aspect of the program's review process and management. Comments should be based on a review of proposal actions (awards, declinations, returns without review, and withdrawals) that were *completed within the past four fiscal years*. Provide comments for *each* program being reviewed and for those questions that are relevant to the program(s) under review. Quantitative information may be required for some questions. Constructive comments noting areas in need of improvement are encouraged.

I. Questions about the quality and effectiveness of the program's use of merit review process. Please answer the following questions about the effectiveness of the merit review process and provide comments or concerns in the space below the question.

QUALITY AND EFFECTIVENESS OF MERIT REVIEW PROCESS	YES, NO, DATA NOT AVAILABLE, or NOT APPLICABLE
<p>1. Are the review methods (for example, panel, ad hoc, site visits) appropriate?</p> <p>The CNS CoV found that the review methods are appropriate. Panel reviews were supplemented by external (ad hoc) reviews, typically written by an expert who was unable to participate in the panel.</p> <p>There is an increasing utilization of virtual and hybrid panels, which is convenient for panelists and, more importantly, expands the panelist pool and its diversity.</p> <p>The CNS CoV noted that the correspondence with the PIs recorded by program directors in the jackets shows that some proposals were transferred among programs (e.g., CSR to NeTS) with the consent of the PI. The CNS CoV commends the program directors on seeking to place each proposal in the most appropriate program panel.</p> <p>Evaluation of many larger proposals involved site visits and reverse site visits. The CNS CoV felt that these visits were very valuable in the decision process. There was evidence in the material reviewed that one visit simultaneously addressed some reservations panelists had about the proposal text, while also allowing the program director to determine that a reduced but still impressive scope of work (and funding) was appropriate.</p>	YES
<p>2. Are both merit review criteria addressed</p> <p>a) In individual reviews?</p> <p>b) In panel summaries?</p> <p>c) In Program Officer review analyses?</p>	<p>YES</p> <p>YES</p> <p>YES</p>

<p>Most reviews read by the CNS CoV were found to address the strengths and weaknesses of both merit criteria, and to give helpful suggestions for improvement to the PI. A few reviews, however, were rather terse without helpful suggestions for improvement.</p> <p>There was at least one review that included only a single sentence about the broader impact of a proposal: the CNS CoV suggests continuing to educate reviewers on the depth required in their reviews, e.g., at least include one strength and one weakness for broader impacts. Some reviewers only commented on the broadening participation aspect in their review of broader impacts. There was one proposal where a reviewer said that the rating was based only on intellectual merit, which conflicts with NSF guidelines to consider both intellectual and broader impact criteria.</p> <p>Additional information was provided in some reviews, which the CNS CoV appreciated. For example, certain solicitations required separate collaboration plans, management plans, or validation plans. In these cases, additional criteria were included in the reviews, panel summaries, and program director analysis. Some reviews, summaries and analyses also included specific sections that discuss the relevance to the particular program/solicitation.</p>	
--	--

<p>3. Do the individual reviewers giving written reviews provide substantive comments to explain their assessment of the proposals?</p> <p>Most of the reviews read by the CNS CoV contained detailed comments with clear descriptions of the strengths and weaknesses of different aspects of the proposal. Some reviews included specific suggestions for focusing on certain promising tasks of the proposed research, and omitting other less novel or less viable tasks.</p> <p>There is some variance in review quality, however, and a few of the reviews that the CNS CoV read could have been more detailed and more constructive. Perhaps program directors can educate reviewers by distributing examples of anonymized, high quality, reviews that are evaluative in terms of both intellectual merit and broader impact, and can ensure that reviewers revise their reviews when necessary to better reflect NSF reviewing guidelines and standards.</p>	<p>YES</p>
---	------------

<p>4. Do the panel summaries provide the rationale for the panel consensus (or reasons consensus was not reached)?</p> <p>The CNS CoV found that panel summaries typically indicate that the panel felt that the proposal was strong and they rated it as highly competitive or competitive, or that the panel felt that the proposal was lacking in several respects and they rated it as low competitive or non-competitive. The summaries typically explained the considerations that led to the panel conclusion/ranking, and, in so doing, showed which of the comments from the individual reviews “bubbled up” to being the major points of consideration.</p>	<p>YES</p>
---	------------

<p>The CNS CoV observed very few cases where a summary does not clearly explain why a rating was “low competitive” despite relatively high ratings by most of the individual reviewers.</p>	
<p>5. Does the documentation in the jacket provide the rationale for the award/decline decision?</p> <p>The review analyses by the program directors are very detailed for funded proposals. Several CNS CoV members noted that the analyses are extremely well-written.</p> <p>For declined proposals, most of jackets only included a review analysis with a generic statement listing typical reasons proposals are declined, not a specific reason for the particular proposal. However, at least one exception was observed for a proposal that received an excellent by a reviewer and a poor by another, and the review analysis explained the panel dynamics and the reason for the final decision.</p> <p>The diary notes or correspondence sections sometimes included detailed responses from the PI to concerns raised by the panel. In most of these cases, the program director analysis agreed with the responses.</p> <p>In a few cases, although the reviewer ratings were mixed, the proposal was awarded due to the innovative intellectual merit or broader impacts, or the under-served community that will be impacted by the award. The CNS CoV generally believed that these cases worked well and that the program directors ensured that the process was fair.</p> <p>For some international collaboration programs, a separate review is conducted by the corresponding foreign funding agency, so the funding decision is influenced by that agency in many cases. The CNS CoV recommends periodically reviewing the effectiveness of these collaborative programs.</p>	<p>YES</p>

<p>6. Does the documentation to the PI provide the rationale for the award/decline decision?</p> <p>For the most part, the reviews and panel summary contained sufficiently detailed information to explain the decision, but in a few cases a relatively highly rated proposal was not funded. The CNS CoV suggests that more explanation is given to the PI about how competitive funding is in these cases.</p> <p>For education programs, failing to address unrepresented population needs or lack of Research Practitioner Partnership (RPP) was communicated clearly to the PI.</p>	<p>YES</p>
<p>7. Additional comments on the quality and effectiveness of the program's use of merit review process:</p> <ul style="list-style-type: none"> • The reviewer ranking is objective, while the ranking in the panel summary is relative to other proposals being considered. This can lead to inconsistencies, e.g., if a proposal with all "V" ratings is in competition with proposals that are all "E"s, the proposal rated "V" may be have a panel summary ranking of "Low Competitive." The CNS CoV recommends that program directors ensure that the panel summary accurately describes the panel discussion in these cases; e.g., feedback such as "the panel agreed that drawback A overwhelmed strength B, and this is why the proposal was judged to be Low Competitive" would be extremely valuable to the PI. • As stated in the previous CoV report, the broader impact and broadening participation requirements need to continue to be <i>uniformly</i> and <i>clearly</i> communicated to both the PIs and the reviewers, e.g., by linking to relevant white papers and web sites. The program directors are encouraged to work more closely with panelists to provide meaningful feedback on broader impact, and to explain the extent to which broader impact influences decision making. One suggestion from the CNS CoV is to include panelists or ad hoc reviewers with the necessary expertise and ability to evaluate broader impact in the review process. Another suggestion is to modify the review template to be consistent with other communications and program director views on broader impact. 	

II. Questions concerning the selection of reviewers. Please answer the following questions about the selection of reviewers and provide comments or concerns in the space below the question.

SELECTION OF REVIEWERS	YES, NO, DATA NOT AVAILABLE, or NOT APPLICABLE
<p>1. Did the program make use of reviewers having appropriate expertise and/or qualifications?</p> <p>The CNS CoV was impressed by the qualifications of the reviewers. The reviewer expertise was appropriate based on the program. For example, in the case of the Cyber-Physical Systems program, the panel expertise was very broad, whereas for the CSR and NeTS programs, the panels were appropriately focused. Panelists included industry participants, and academic participants from a broad range of schools.</p> <p>Most panels had a diverse pool of panelists, e.g., in terms of gender representation, but perhaps even more attention can be paid to the diversity aspect in the future.</p> <p>As expected, the education program panels had only few panelists from CS departments; the remainder being in education but well-known in the field of CS Education. The CNS CoV commended the CNS program directors on including panelists from public school districts on these panels.</p>	YES
<p>2. Did the program recognize and resolve conflicts of interest when appropriate?</p>	YES
<p>3. Additional comments on reviewer selection:</p> <p>Overall, the review selection process appears to be working well, but continued efforts on diversity (including diversity of institutions) and inclusion of junior faculty members will remain important. One suggestion from the CNS CoV is to ensure that a more uniform process is followed across panels for reaching out to junior researchers, and to researchers at different types of institutions.</p>	

III. Questions concerning the management of the program under review. Please comment on the following:

MANAGEMENT OF THE PROGRAM UNDER REVIEW

1. Management of the program.

- Overall, the CNS CoV was extremely satisfied with the management of the program, the broad range of topics in the solicitations, and the integrity of the review process.
- There appear to be challenges to maintain the six month timeframe for providing a decision on a submitted proposal, but most proposals are handled in under nine months. The number of proposals with a decision provided within six months has gone down from 76% in 2015 to 35% in 2018. This is not surprising, since this may be partly due to the relocation to a new office -- an extremely disruptive event in the life of the organization.
- One suggestion that the CNS CoV has is increasing the interaction and cooperation with OAC, since CNS also handles a number of infrastructure programs.
- Another suggestion is educating the community about the do's and don'ts with respect to including industry participants and collaborators on proposals to regular programs.
- The CNS CoV found that the quality of project annual reports submitted by the awarded PIs varies significantly, with some reports not being sufficiently detailed. The community needs to be better educated about the level of depth required in project annual reports. A few CNS CoV members also noted that the broader impact outcomes listed in the project annual reports mention research mentorship and successful career progressions for students, but few indicate the outcomes for underrepresented groups that were promised in the proposals.
- The CNS CoV suggests further engagement with the research community, e.g., through CRA, the CCC, ACM SIGs, IEEE TCs, and department head mailing lists. Conducting Q&A webinars, emailing presentation slides, arranging regional visits by CNS staff, and partnering with programs such as Women in CyberSecurity (WiCyS) and the Tapia conference can also be effective methods for reaching underrepresented researchers and institutions.

2. Responsiveness of the program to emerging research and education opportunities.

The CNS CoV found that new and updated solicitations respond extremely well to emerging needs, which are sometimes based on reports from community-driven workshops or industry collaboration. One example is the new AI initiatives. Another example is the partnership programs with VMware and Intel. A third example is the EU-US collaboration program where the solicitation was largely based on a community workshop that identified specific sub-areas where international collaboration is required to achieve a deeper understanding of the inter-domain problems in the Internet, and leverage combined research resources available in the North American and European continents.

Additionally, the increasing number of cross-cutting programs reflects an appropriate response to the broadening importance of computing. However, this broadening scope poses significant challenges in handling a large number of solicitations and programs, with broad communities of PIs with different backgrounds.

3. Program planning and prioritization process (internal and external) that guided the development of the portfolio.

The CNS CoV commends the CNS staff on balancing the demands on resources according to national priorities. CNS continues to improve the review processes and public access to NSF-funded peer-reviewed publications. However, CNS, and CISE in general, are under tremendous pressure due to the increasing number of proposals, emerging research areas and programs, and a very tight budget.

4. Responsiveness of program to previous COV comments and recommendations.

- The CNS CoV believes that CNS has addressed the most critical issues within its purview. For example, the program directors have utilized more ad hoc reviews, virtual panels, and hybrid virtual and in-person panels. They have emphasized the broader impact requirements in the proposal review process.
- An important issue discussed in the last CoV report that remains a concern is the lack of powerful tools for (i) analyzing long-term trends and making complex queries, and (ii) automating tasks to streamline operations and reduce CISE staff workload. The CNS CoV suggests exploring the idea of employing undergraduate and graduate students as summer interns to develop such software tools. Programs such as CyberCorps emphasize service, and partnering with these programs can be an effective method for recruiting interns. As discussed above, the CRA, the CCC, ACM SIGs, IEEE TCs, and department head mailing lists can aid in gathering data about the research community, aspiring PIs, and potential proposal reviewers.
- Another issue that remains a major concern is the load on the program directors, especially due to the increase in cross-cutting programs which require much more coordination effort, sometimes across programs, divisions, directorates, or agencies. The program directors need increased resources, more support, and better productivity tools to manage this load.

IV. Questions about Portfolio. Please answer the following about the portfolio of awards made by the program under review.

<p align="center">RESULTING PORTFOLIO OF AWARDS</p>	<p align="center">APPROPRIATE, NOT APPROPRIATE, OR DATA NOT AVAILABLE</p>
<p>1. Does the program portfolio have an appropriate balance of awards across disciplines and sub-disciplines of the activity?</p> <p>The CNS CoV found that the proposals awarded had a good balance of foundational and applied research. The portfolio includes a wide range of research, education, infrastructure, and interdisciplinary projects. Some programs are open-ended such as the CSR, NeTS, and core SaTC programs, while others have a specific focus such as radio spectrum programs EARS and SpecEES.</p> <p>Clearly, the portfolio is limited by the numbers of submitted proposals in response to a particular solicitation.</p> <p>Keeping the core programs as the largest funding elements was viewed positively by several members of the CNS CoV: although there are many cross-cutting programs, by definition the core programs are broader and more general. Some CNS CoV members noted that SaTC is relatively large, but many of today's societal problems that need immediate solutions are in this space, so it should continue to be of high priority.</p>	<p>YES</p>
<p>2. Are awards appropriate in size and duration for the scope of the projects?</p> <p>The CNS CoV found that there is a good balance of small and large projects, with appropriate sizes and durations. One suggestion is to link the award sizes to the increasing cost of conducting research/inflation. This would ensure that a small award can continue to fund at least one graduate student.</p> <p>In addition to the traditional three to five year awards, the CNS CoV noted that there is judicious use of EAGER and RAPID awards and other seed funding. The CNS CoV suggests exploring whether a larger number of small seed awards can be given to new and early career PIs.</p>	<p>YES</p>
<p>3. Does the program portfolio include awards for projects that are innovative or potentially transformative?</p> <p>The CNS CoV noted several cases of innovative and potentially transformative project ideas that were high risk, high reward. The CNS CoV members observed several funded projects that propose highly innovative approaches that are known to have spurred lots of discussion and follow-up work in the community. The CNS CoV was pleased to observe at least one</p>	<p>YES</p>

case when an innovative research proposal was awarded despite mixed reviewer ratings; the text of the highest and lowest score reviews showed that this was a high risk project with significant disagreement on its chances for success.

The education awards address important questions about how to prepare teachers to teach CS, and how to broaden participation in computing by investigating how to support children from underrepresented groups (e.g., English Language Learners).

<p>4. Does the program portfolio include inter- and multi-disciplinary projects?</p> <p>The CNS CoV noted a number of cases, especially in the Cyber-Physical Systems, Secure and Trustworthy Cyberspace, and Smart & Connected Communities programs where the funded projects included experts from multiple disciplines. The CNS division received more funds from other divisions than it distributed, which is expected given the wide variety of programs and solicitations led by the division.</p>	<p>YES</p>
<p>5. Does the program portfolio have an appropriate geographical distribution of Principal Investigators?</p> <p>The CNS CoV found that, as expected, the distribution generally follows the presence of top-25 research institutions, with high numbers of awards in California, Texas, Massachusetts, and New York. Those four states mostly had the highest success ratios as well.</p>	<p>YES</p>
<p>6. Does the program portfolio have an appropriate balance of awards to different types of institutions?</p> <p>The CNS CoV found that, as expected, there are large discrepancies among different types of institutions in the number of proposals, number of awards, and success rates, which are low for non-PhD-granting institutions. Very few proposals are submitted by two-year institutions. To increase the number of awards to institutions that do not have research-intensive programs, an increase in the number of programs or awards tailored to such institutions, such as REU sites, can be employed. The CNS CoV also suggests reaching out to investigators at different types of institutions to participate in workshops on writing effective proposals, and to serve on proposal review panels.</p>	<p>YES</p>
<p>7. Does the program portfolio have an appropriate balance of awards to new and early-career investigators?</p> <p>The success rate for established PIs is higher than that for new PIs, which is expected. There had been some confusion in previous years among junior faculty regarding whether to submit to CRII or CAREER or both. The CNS CoV commends CISE staff on re-examining the focus of the CRII program, and revising its solicitation to clearly communicate eligibility requirements to aspiring PIs.</p> <p>The CNS CoV was extremely concerned about the impact of the rapidly increasing number of junior faculty members in Computer Science on the proposal review process. This growth will be especially impactful on CRII and CAREER proposals in the near term, and subsequently on proposals to all programs. The significant growth in junior faculty (e.g., recent hiring of 3-7 new faculty members per year in many CS departments, instead of 1-2 as was historically typical) will lead to a doubling or more of CAREER and CRII</p>	<p>YES</p>

<p>proposals. To remain effective, it is important that these programs not become excessively competitive. The success rate for CAREER has dropped in the past 5 years, and several members of the CNS CoV noted the negative impacts of this increased selectiveness, such as faculty submitting their CAREER proposals later in their junior career, and universities emphasizing CAREER awards when considering tenure cases.</p>	
<p>8. Does the program portfolio include projects that integrate research and education?</p> <p>The CNS CoV noted several examples of awarded projects that include strong integration of research and education. For example, CAREER awards included curriculum module development, research involvement by students at all levels including undergraduate students, and development of tools and demos that can be used in the classroom. The REU sites and supplements also enable undergraduate students to learn while conducting research, benefitting both research and education.</p>	<p>YES</p>
<p>9. Does the program portfolio have appropriate participation of underrepresented groups³?</p> <p>The CNS CoV found that success rates are roughly similar for proposals from females, minorities, and those with disabilities to proposals from other PIs. However, the number of proposals submitted by these underrepresented groups is small for the most part, and hence the number of awards in the portfolio is similarly small, especially for minorities and those with disabilities. Female PIs are much better represented in education programs.</p> <p>The information on membership in underrepresented groups is typically underreported, and hence it is extremely difficult to track and make meaningful conclusions about their participation. The CNS CoV felt that NSF needs to explore how to better address the data gathering aspect for gender, minority, and disability status. For gender, one would need to include additional categories for gender neutral and non-binary.</p>	<p>YES</p>
<p>10. Is the program relevant to national priorities, agency mission, relevant fields and other constituent needs? Include citations of relevant external reports.</p> <p>The CNS CoV noted several examples of projects that address national priorities, such as the SaTC program addressing critical cybersecurity needs, and the CPS program for cyberphysical systems, and programs responding to the AI Executive Order and the National Quantum Initiative Act.</p>	<p>YES</p>

³ NSF does not have the legal authority to require principal investigators or reviewers to provide demographic data. Since provision of such data is voluntary, the demographic data available are incomplete. This may make it difficult to answer this question for small programs. However, experience suggests that even with the limited data available, COVs are able to provide a meaningful response to this question for most programs.

11. Additional comments on the quality of the projects or the balance of the portfolio:	
---	--

OTHER TOPICS

1. Please comment on any program areas in need of improvement or gaps (if any) within program areas.

The CNS CoV did not observe any significant gaps, but noted that there was a significant variation in the size, depth and potential impact across areas. This is to be expected, since there is also variation in the number of submitted proposals. Given the concerns of the workload on program directors, the CNS CoV suggests that program directors consider focusing on a smaller number of broader programs.

2. Please provide comments as appropriate on the program's performance in meeting program-specific goals and objectives that are not covered by the above questions.

Overall, the CNS CoV committee was impressed with the merit review process and research portfolio of CNS over the period of review. However, this positive conclusion was shadowed by alarm at the trends and what the state of CISE will be in four to five years, when the next CoV convenes. Many academic members of the CoV have seen or encountered similar trends within their departments with the increase in undergraduate enrollments. The growth in Computer Science enrollments has had tremendous and difficult impact on departments. The CoV fears that there will be a similar sea change within CISE in the next four to five years, as these increased undergraduate enrollments play out as growing departments with many junior faculty. Given that program directors are already stretched thin, proactively handling this expected surge in proposals will be critical for the long-term health of CISE.

3. Please identify agency-wide issues that should be addressed by NSF to help improve the program's performance.

The workload on the program directors and staff remains a major concern, especially with the increasing number of computer scientists nationwide, which will result in a significant increase in the number of submitted proposals. The CNS CoV recognizes that, with the increase in the number of solicitations and proposals, the workload on the program directors has significantly increased, and we commend the program directors on adapting to the increasing workload thus far, without harming the quality of the review process. Moving forward, the CNS CoV encourages needs-motivated design of productivity and data analysis tools that reduce the workload on NSF staff, including exploring the idea of employing undergraduate and graduate students as summer interns to develop such tools.

4. Please provide comments on any other issues the COV feels are relevant.
NSF would appreciate your comments on how to improve the COV review process, format and report template.

- The CNS CoV members greatly benefitted from interactions with the CCF and IIS CoV members, and from discussions with the NSF program directors and staff.

- More time for the review of the materials would have benefitted the CNS CoV.
- Most CNS CoV members found that the eJacket system requires a lot of manual effort to navigate across pages and download many separate small Word and PDF documents instead of having all the information consolidated on a few web pages or documents. The search facilities in the system were inadequate to query the system based on various criteria, or to link jackets of collaborative proposals in an automatic fashion. It would be more convenient to download all the needed information with a single or a few mouse clicks.
- Some of the tables in the division annual reports were difficult to parse as they were split across multiple pages or not labeled very clearly, and only PDF (as opposed to the source data) was provided. It took some effort to spot trends as one had to open five reports simultaneously and scroll to comparable pages. If it would be possible to plot trends across the review period, the process can be much more efficient. Providing the data as Excel spreadsheets would help the CoV, and potentially reduce the effort by NSF staff to create the reports.
- A few of the questions in the template of this report are difficult to address with the provided information, or are hard to interpret, since what is considered “appropriate” is subjective.
- The CNS CoV suggests distributing to the next CoV, at the start of the process, a carefully written document describing the precise goals of the CoV, and the information that needs to be reviewed and where to find it, in order to make the process more efficient and effective.

The Committee of Visitors is part of a Federal advisory committee. The function of Federal advisory committees is advisory only. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the Advisory Committee, and do not necessarily reflect the views of the National Science Foundation.

SIGNATURE BLOCK:



For the CNS CoV Sub-committee
Sonia Fahmy
Chair

CORE QUESTIONS and REPORT TEMPLATE for FY 2019 NSF COMMITTEE OF VISITOR (COV) REVIEWS

Guidance to NSF Staff: This document includes the FY 2019 set of Core Questions and the COV Report Template for use by NSF staff when preparing and conducting COVs during FY 2019. Specific guidance for NSF staff describing the COV review process is described in the “COV Reviews” section of NSF’s Administrative Policies and Procedures which can be obtained at <https://inside.nsf.gov/tools/toolsdocuments/Inside%20NSF%20Documents/Policy,%20Procedures,%20Roles%20and%20Responsibilities%20for%20COV%20Reviews%20and%20Program%20Portfolio%20Reviews.pdf>¹.

NSF relies on the judgment of external experts to maintain high standards of program management, to provide advice for continuous improvement of NSF performance, and to ensure openness to the research and education community served by the Foundation. COV reviews provide NSF with external expert judgments in two areas: (1) assessments of the quality and integrity of program operations; and (2) program-level technical and managerial matters pertaining to proposal decisions.

The program(s) under review may include several sub-activities as well as NSF-wide activities. The directorate or division may instruct the COV to provide answers addressing a cluster or group of programs – a portfolio of activities integrated as a whole – or to provide answers specific to the sub-activities of the program, with the latter requiring more time but providing more detailed information.

The Division or Directorate may add questions relevant to the activities under review. Copies of the report template and the charge to the COV should be provided to OIA prior to forwarding to the COV. In order to provide COV members adequate time to read and consider the COV materials, including proposal jackets, COV members should be given access to the materials in the eJacket COV module approximately four weeks before the scheduled face-to-face meeting of the COV members. Before providing access to jackets, the Conflict of Interest and Confidentiality briefing for COV members should be conducted by webinar, during which, NSF staff should also summarize the scope of the program(s) under review and answer COV questions about the template.

Suggested sources of information for COVs to consider are provided for each item. As indicated, a resource for NSF staff preparing data for COVs is the Enterprise Information System (EIS) –Web COV module, which can be accessed by NSF staff only at <http://budg-eis-01/eisportal/default.aspx>. In addition, NSF staff preparing for the COV should consider other sources of information, as appropriate for the programs under review.

For programs using section IV (addressing portfolio balance), the program should provide the COV with a statement of the program’s portfolio goals and ask specific questions about the program under review. Some suggestions regarding portfolio dimensions are given on the template. These suggestions will not be appropriate for all programs.

Guidance to the COV: The COV report should provide a balanced assessment of NSF’s performance in the integrity and efficiency of the **processes** related to proposal review. Discussions leading to answers of the Core Questions will require study of confidential material such as declined proposals and reviewer comments. **COV reports should not contain confidential material or specific information about declined proposals.** The reports generated by COVs are made available to the public.

¹ This document has three parts: (1) Policy, (2) Procedures, and (3) Roles & Responsibilities.

We encourage COV members to provide comments to NSF on how to improve in all areas, as well as suggestions for the COV process, format, and questions. For past COV reports, please see <http://www.nsf.gov/od/oia/activities/cov/>.

FY 2019 REPORT TEMPLATE FOR NSF COMMITTEES OF VISITORS (COVs)

The table below should be completed by program staff.

Date of COV: November 6-8, 2019			
Program/Cluster/Section: Core programs <ul style="list-style-type: none"> • Cyber-Human Systems (CHS) • Information Integration and Informatics (III) • Robust Intelligence (RI) Cross-cutting programs <ul style="list-style-type: none"> • Big Data Regional Innovation Hubs (BD Hubs) • Big Data Regional Innovation Hubs: Establishing Spokes to Advance Big Data (BD Spokes) • Collaborative Research in Computational Neuroscience (CRCNS) • Critical Techniques and Technologies for Advancing Big Data Science and Engineering (BIGDATA) • Cyberlearning and Future Learning Technologies (CFLT) • Integrative Strategies for Understanding Neural and Cognitive Systems (NCS) • National Robotics Initiative (NRI) • NSF/Intel Partnership on Visual and Experiential Computing (VEC) • Smart and Autonomous Systems (S&AS) • Smart and Connected Health (SCH) 			
Division: Information and Intelligent Systems (IIS)			
Directorate: Computer and Information Science and Engineering (CISE)			
Number of actions reviewed: 192			
Awards: 93			
Declinations: 92			
Other: 7			
Total number of actions within Program/Cluster/Division during period under review:			
	# of Awards	# of Declines	Total
FY 2014	488	2,503	2,991
FY 2015	544	2,646	3,190
FY 2016	575	2,657	3,232
FY 2017	521	2,647	3,168
FY 2018	590	2,500	3,090
Manner in which reviewed actions were selected:			
<ol style="list-style-type: none"> (1) Random selection from bins for programs over the covered period, evenly divided between awards and declinations; (2) Random selection of EAGER, RAPID, travel grant, and workshop awards; and (3) Additional random selection from bins requested by the COV. 			

COV Membership

	Name	Affiliation
COV Chair or Co-Chairs:	Overall COV Chair: Azer Bestavros Subcommittee Chair for IIS: Julia Hirschberg	Boston University Columbia University
COV Members:	Rina Dechter Susan Dumais Chad Jenkins Leslie Pack Kaelbling Ben Shneiderman Venkatramanan Siva Subrahmanian	University of California Irvine Microsoft Research University of Michigan MIT University of Maryland Dartmouth College

MERIT REVIEW CRITERIA

An understanding of NSF's merit review criteria is important in order to answer some of the questions on the template. Reproduced below is the information provided to proposers in the Grant Proposal Guide about the merit review criteria and the principles associated with them. Also included is a description of some examples of broader impacts, provided by the National Science Board

1. Merit Review Principles

These principles are to be given due diligence by PIs and organizations when preparing proposals and managing projects, by reviewers when reading and evaluating proposals, and by NSF program staff when determining whether or not to recommend proposals for funding and while overseeing awards. Given that NSF is the primary federal agency charged with nurturing and supporting excellence in basic research and education, the following three principles apply:

- All NSF projects should be of the highest quality and have the potential to advance, if not transform, the frontiers of knowledge.
- NSF projects, in the aggregate, should contribute more broadly to achieving societal goals. These broader impacts may be accomplished through the research itself, through activities that are directly related to specific research projects, or through activities that are supported by, but are complementary to, the project. The project activities may be based on previously established and/or innovative methods and approaches, but in either case must be well justified.
- Meaningful assessment and evaluation of NSF funded projects should be based on appropriate metrics, keeping in mind the likely correlation between the effect of broader impacts and the resources provided to implement projects. If the size of the activity is limited, evaluation of that activity in isolation is not likely to be meaningful. Thus, assessing the effectiveness of these activities may best be done at a higher, more aggregated, level than the individual project.

With respect to the third principle, even if assessment of Broader Impacts outcomes for particular projects is done at an aggregated level, PIs are expected to be accountable for carrying out the activities described in the funded project. Thus, individual projects should include clearly stated goals, specific descriptions of the activities that the PI intends to do, and a plan in place to document the outputs of those activities. These three merit review principles provide the basis for the merit review criteria, as well as a context within which the users of the criteria can better understand their intent.

2. Merit Review Criteria

All NSF proposals are evaluated through use of two National Science Board approved merit review criteria. In some instances, however, NSF will employ additional criteria as required to highlight the specific objectives of certain programs and activities.

The two merit review criteria are listed below. Both criteria are to be given full consideration during the review and decision-making processes; each criterion is necessary but neither, by itself, is sufficient. Therefore, proposers must fully address both criteria. ([PAPPG Chapter II.C.2.d.\(i\)](#) contains additional information for use by proposers in development of the Project Description section of the proposal.) Reviewers are strongly encouraged to review the criteria, including [PAPPG Chapter II.C.2.d.\(i\)](#), prior to the review of a proposal.

When evaluating NSF proposals, reviewers will be asked to consider what the proposers want to do, why they want to do it, how they plan to do it, how they will know if they succeed, and what benefits could accrue if the project is successful. These issues apply both to the technical aspects of the proposal and the way in which the project may make broader contributions. To that end, reviewers will be asked to evaluate all proposals against two criteria:

- **Intellectual Merit:** The Intellectual Merit criterion encompasses the potential to advance knowledge; and
- **Broader Impacts:** The Broader Impacts criterion encompasses the potential to benefit society and contribute to the achievement of specific, desired societal outcomes.

The following elements should be considered in the review for both criteria:

1. What is the potential for the proposed activity to:
 - a. Advance knowledge and understanding within its own field or across different fields (Intellectual Merit); and
 - b. Benefit society or advance desired societal outcomes (Broader Impacts)?
2. To what extent do the proposed activities suggest and explore creative, original, or potentially transformative concepts?
3. Is the plan for carrying out the proposed activities well-reasoned, well-organized, and based on a sound rationale? Does the plan incorporate a mechanism to assess success?
4. How well qualified is the individual, team, or organization to conduct the proposed activities?
5. Are there adequate resources available to the PI (either at the home organization or through collaborations) to carry out the proposed activities?

3. Examples of Broader Impacts

The National Science Board described some examples of broader impacts of research, beyond the intrinsic importance of advancing knowledge.² “These outcomes include (but are not limited to) increased participation of women, persons with disabilities, and underrepresented minorities in science, technology, engineering, and mathematics (STEM); improved STEM education at all levels; increased public scientific literacy and public engagement with science and technology; improved well-being of individuals in society; development of a globally competitive STEM workforce; increased partnerships between academia, industry, and others; increased national security; increased economic competitiveness of the United States; and enhanced infrastructure for research and education. These examples of societally relevant outcomes should not be considered either comprehensive or prescriptive. Investigators may include appropriate outcomes not covered by these examples.”

² [NSB-MR-11-22](#)

INTEGRITY AND EFFICIENCY OF THE PROGRAM'S PROCESSES AND MANAGEMENT

Briefly discuss and provide comments for *each* relevant aspect of the program's review process and management. Comments should be based on a review of proposal actions (awards, declinations, returns without review, and withdrawals) that were *completed within the past four fiscal years*. Provide comments for *each* program being reviewed and for those questions that are relevant to the program(s) under review. Quantitative information may be required for some questions. Constructive comments noting areas in need of improvement are encouraged.

I. Questions about the quality and effectiveness of the program's use of merit review process. Please answer the following questions about the effectiveness of the merit review process and provide comments or concerns in the space below the question.

QUALITY AND EFFECTIVENESS OF MERIT REVIEW PROCESS	YES, NO, DATA NOT AVAILABLE, or NOT APPLICABLE
<p>1. Are the review methods (for example, panel, ad hoc, site visits) appropriate?</p> <p>The IIS COV found that the reviewing methods were appropriate. In-person panels remain a strong mechanism for proposal review, providing an opportunity for reviewers to share wisdom and perspective in order to arrive at more robust decisions.</p> <p>There were a few cases in which it seemed that the aggregate expertise of the reviewers did not cover the proposal area well and that it might have been useful to use more ad hoc reviewers. The COV was pleased to learn that in some cases in which the lack of expertise among reviewers only becomes clear during the panel meeting, the program director can commission an ad hoc review before the final decisions are made.</p> <p>"Virtual" panels seem to be functioning effectively, and although they don't afford the opportunities for rich interaction among the panelists and NSF personnel that in-person meetings do, they do make it possible for panelists with restricted travel opportunities, e.g. due to teaching or childcare responsibilities, to participate. Innovations in tele-conferencing software and equipment have supported virtual panels well.</p> <p>The COV noted that, on average, proposals that were rejected had more reviews than those that were funded. That may indicate that extra care is being taken that proposals are not rejected without substantial consideration.</p>	Yes.
<p>2. Are both merit review criteria addressed?</p> <p>a) In individual reviews?</p> <p>Both merit review criteria were addressed explicitly in almost all individual reviews, although in some cases the discussion of the broader impact criterion</p>	Yes.

<p>was somewhat pro forma. It appears that reviewers are sometimes not completely clear about how to evaluate broader impact and/or weigh the two criteria in their ratings.</p> <p>b) In panel summaries: Yes, very thoroughly.</p> <p>c) In Program Director review analyses? Program Director review analysis (as well as panel summaries) consistently addressed both criteria.</p>	
---	--

<p>3. Do the individual reviewers giving written reviews provide substantive comments to explain their assessment of the proposals?</p> <p>The IIS COV found that the reviews were in general done with care and consideration. In a few cases, reviews were terse or not directly helpful to the proposers. It might be useful to remind panelists that one of their objectives should be to provide concrete advice to proposers with weak proposals, in order to help them improve the quality of future proposals.</p>	<p>Yes.</p>
--	-------------

<p>4. Do the panel summaries provide the rationale for the panel consensus (or reasons consensus was not reached)?</p> <p>Panel summaries were usually good syntheses of the panel discussion and provided insight into the rationale for the decision. Occasionally they were simply a restatement of the individual reviews. More concerning is that the summaries occasionally seemed to be somewhat at odds with particular reviews (e.g., not acknowledging what seemed to be substantial weaknesses surfaced in individual reviews). None seemed to address the question of why panel consensus was not reached, specifically.</p>	<p>Yes</p>
--	------------

<p>5. Does the documentation in the jacket provide the rationale for the award/decline decision?</p> <p>The documentation in the jacket usually provided very good information about why proposals were awarded or declined. The Review Analysis gave a good summary of the panel deliberations and overall rationale for award decisions. Most walked through each of the steps of the review process for the proposal, adding insights along the way about the reviews (more than just scores), COI resolution, panel deliberation, and their own perspective. One suggestion is whether more information for some of the declined proposals that were not discussed in the panels could be provided. Another addition the group considered was whether the Review Analysis could include information on the process by which panelists were considered, selected, and recruited. For example, how diverse was the pool considered? We understand from breakout sessions with PDs that it is much more difficult to recruit diversity panelists but that allowing reviewers to participate virtually is a very good way to improve this.</p>	<p>Yes.</p>
--	-------------

<p>6. Does the documentation to the PI provide the rationale for the award/decline decision?</p> <p>Yes, for proposals that were discussed in panels, the reviews and summaries provide useful information. We wonder if, for proposals that the panel decides not to discuss, NSF could provide these PIs with information about a PD they could contact if they have questions. Perhaps this could be added to the “context” statement.</p>	<p>Yes.</p>
<p>7. Additional comments on the quality and effectiveness of the program’s use of merit review process:</p> <p>The process and the way PDs and staff are conducting this seems really excellent, especially given the challenges they face in recruiting diverse, non-COI panelists and dealing with many different software systems and many manual tasks. The triage of low-rated proposals also seems to be working very well, as it allows more time to discuss the more competitive proposals. While PDs seem to be adept at finding panelists and ad hoc reviewers, this process did differ from PD to PD. It might be useful to develop some “best practices” guidelines to pass along to new PDs in reviewer selection and how reviewers are asked to rate their expertise in different areas.</p>	<p>.</p>

II. Questions concerning the selection of reviewers. Please answer the following questions about the selection of reviewers and provide comments or concerns in the space below the question.

SELECTION OF REVIEWERS	YES, NO, DATA NOT AVAILABLE, or NOT APPLICABLE
<p>1. Did the program make use of reviewers having appropriate expertise and/or qualifications?</p> <p>The IIS CoV was generally pleased with the quality of reviews and the expertise of reviewers, especially in view of the well-understood challenge of recruiting reviewers. The review of eJackets showed that collectively reviewers had appropriate expertise, including senior personnel and recognized leaders.</p> <p>Two changes seem to have helped Program Directors to identify reviews with appropriate expertise. There is increasing flexibility in constructing panels, including the use of virtual panels and the ability to recruit ad hoc reviews as needed. In addition, new tools have been developed to help PDs identify reviewers with appropriate expertise. These changes increase the number of expert reviewers who can take part in the panels.</p> <p>Identifying expert reviewers for the wide-range of perspectives needed for cross-cutting programs is still a challenge, and the ability to draw on the expertise of ad hoc reviews is especially important here.</p>	Yes.
<p>2. Did the program recognize and resolve conflicts of interest when appropriate?</p> <p>There were no issues identified in the jackets we reviewed. However, we understand the guidelines for COI continue to evolve. We fear that the current guidelines (no panelist from the same institution) will increase the burden on PDs, making it more difficult to find qualified panelists.</p>	Yes.
<p>3. Additional comments on reviewer selection:</p> <p>NSF Program Directors are asked to ensure reviewer diversity along many different dimensions -- experience, geography, type of institution, gender, ethnicity, disability, and institution (industry, academia, government). CISE IIS PDs are committed to these goals and work very hard to put together diverse panels. We encourage continuation of those effort, but it is a very challenging task.</p> <p>We encourage NSF to support a more uniform process across panels to identify and recruit more diverse researchers along all of these dimensions. For underrepresented populations the challenge is especially acute, and we recommend that NSF continues its efforts to engage underrepresented groups</p>	Yes.

and organizations to enable more equitable and effective merit review. These efforts will not only contribute to the diversity of panels, but can also increase awareness of NSF processes. The result will likely contribute to greater proposal submissions from groups currently underrepresented in computing, and improve the quality of scientific advances across the computing research community as a whole.

Explicitly recognizing the work that panel reviewers do (e.g., by certificates or email to department heads) could also increase participation by more junior researchers and researchers from different types of institutions on panels – if this is possible to do annually.

Finally, it is important to systematically capture information about the experience, geography, type of institution, gender, ethnicity, disability, and institution of reviewers and panelists to ensure broad representation and perspectives are taken into account in all aspects of reviewing and decision making.

III. Questions concerning the management of the program under review. Please comment on the following:

MANAGEMENT OF THE PROGRAM UNDER REVIEW
<p>1. Management of the program.</p> <p>According to a line chart in the New York Times (N. Singer, “The Hard Part of Computer Science? Getting Into Class”, Jan 24 2019 - https://www.nytimes.com/2019/01/24/technology/computer-science-courses-college.html retrieved Nov 7 2019), U.S. enrollment in computer science has skyrocketed almost 4-fold from approximately 30,000 in 2007 to over 100,000 in 2017. Moreover, the breadth of computer science research has also widened dramatically – today, researchers in fields ranging from epidemiology to automotive technology to wildlife conservation are all leveraging computational methods in their work. Much of this excitement about computing in recent years has focused on artificial intelligence. There has been growing demand, employment and salaries for students with expertise in Machine Learning and AI.</p> <p>These factors are also seen in the statistics associated with the IIS division, which receives one of the highest numbers of proposals in NSF, yet has one of the lowest funding rates of all NSF divisions. This has led to a significant increase in workload for IIS staff as well as an effective reduction in support for core AI research, given that the allocated AI funds cover research in the application of AI to other areas.</p> <p>Given these factors, the panelists felt that IIS’s management has been stellar. New programs which have been started in the last few years and support for cross-disciplinary programs is promising and should be encouraged, Additionally, the panel notes that support for core Artificial Intelligence cannot be reduced – especially in a climate where the US’s competitors are significantly increasing their own investments in both core and applied AI. NSF needs to address this head on and create a plan for maintaining US leadership in AI.</p>
<p>2. Responsiveness of the program to emerging research and education opportunities.</p> <p>Research in AI has seen a dramatic upsurge in recent years with computer vision, deep learning, and natural language understanding taking giant steps forward. NSF/IIS has introduced several innovative new research programs to respond to these developments, often in partnership with other units at NSF, as well as external stakeholders. While many of these programs are responsive to exciting recent developments in AI, the panel noted that it was sustained NSF support for sub-areas of AI (e.g. machine learning) during times when they made slow progress that had made a major impact. The panel recommended that the increase in funding for these exciting new subareas should continue, but that mechanisms need to be found in order to continue funding other core areas of AI that may, in coming years, generate the same type of significant leaps recently made by machine learning.</p> <p>In addition, IIS has programs with foreign science agencies including ones from Israel, France, and Germany, as well as selected joint fundamental programs with industry – including Intel. Because these engagements can often involve customized partnership agreements, the implementation of these agreements may further add to NSF’s workload.</p>

We also note that some proposals were not particularly strong with respect to educational goals, plans, and aspirations. We recommend continuing emphasis on resources such as the new BPCnet portal.

3. Program planning and prioritization process (internal and external) that guided the development of the portfolio.

New NSF programs are created according to both a bottom up and a top down process. Bottom up programs are created through ideas that originated from the program managers, while top down programs are created using fund reserves held by the NSF top management. These two mechanisms enable NSF to respond quickly to trends and advances on the ground that program managers in specific areas are likely to know, while the top down process enables NSF to respond in an agile fashion to national priorities and new cross-cutting, visionary, multidisciplinary programs.

NSF may consider assessing the role of humans in what will increasingly be an AI-enhanced world. Issues relating to fairness, bias, and error made by AI programs can have deep adverse societal impacts. The question of how humans and AI software and hardware components should interact is a topic worthy of further study.

We recommend that both the top-down and bottom-up processes include workshops where the research community provides input to those in charge of shaping the new program. It would be useful to know more about how NSF decides what new programs to try out and how they assess them.

4. Responsiveness of program to previous COV comments and recommendations.

CISE's response to the comments of the last COV has been exemplary. In particular,

1. The use of *bluejeans* software for virtual panels has now been adopted, although we understand that this may be replaced by Zoom, which does have some unfortunately audio problems.
2. The use of triage to avoid lengthy discussion of proposals that receive poor scores by reviewers is a welcome change and saves much panel time, allowing the group to focus on the best proposals.
3. Most programs seem to be doing a good job in balancing funding by gender and minority status, despite differences in the number of submissions.
4. Having workshops in emerging areas like human-AI interaction is a great way to encourage transformational proposals and also a great way to identify potential reviewers.

We also would recommend that NSF develop some methods to help proposers and reviewers reach a better understanding of "broader impacts."

IV. Questions about Portfolio. Please answer the following about the portfolio of awards made by the program under review.

<p align="center">RESULTING PORTFOLIO OF AWARDS</p>	<p align="center">APPROPRIATE, NOT APPROPRIATE, OR DATA NOT AVAILABLE</p>
<p>1. Does the program portfolio have an appropriate balance of awards across disciplines and sub-disciplines of the activity?</p> <p>The distribution among programs seems well-balanced. We were not clear though what “sub-disciplines” we should be checking. Some of the cross-cutting programs had a low level of awards (e.g., BIGDATA). Specifically we observed the following distribution of awards between CORE and Cross-Cutting.</p> <p>2014: 488 awards, 314 were awarded in the CORE programs (64%) and the rest cross-cutting,</p> <p>In 2015: 544 awards, 330 for CORE (60%),</p> <p>in 2016 575 awards, 325 for CORE (56%),</p> <p>in 2017 521 awards, 295 for CORE (56%) and</p> <p>in 2018, 590 awards, 377 for CORE (63%).</p> <p>Overall CORE awards were ~60% of all awards.</p> <p>This seems like a reasonable balance.</p>	<p>Yes.</p>
<p>2. Are awards appropriate in size and duration for the scope of the projects?</p> <p>It was observed that NRI and S&AS as well as CRCNS and SCH seemed to have a lower success rates. Large proposals were quite rare in jackets we saw.</p>	<p>Yes.</p>
<p>3. Does the program portfolio include awards for projects that are innovative or potentially transformative?</p> <p>Our impression is that most awarded proposals were quite innovative and transformational, although It is hard to tell from a small sample; however, those funded which are not so transformative were definitely going to be beneficial to the larger research community.</p>	<p>Yes.</p>

<p>4. Does the program portfolio include inter- and multi-disciplinary projects?</p> <p>Yes. There were quite a few multi-disciplinary proposals both in the CORE and in the cross-cutting programs. These sometimes seem more difficult to review but are good opportunities for some ad hoc reviewers.</p>	<p>Yes.</p>
<p>5. Does the program portfolio have an appropriate geographical distribution of Principal Investigators?</p> <p>The funding rate and amount generally follow the distribution of population and universities, with a high number and amount of awards in California, Texas, Pennsylvania, Massachusetts, and New York. The EPSCoR program also encourages funding of research of meritorious proposals throughout EPSCoR eligible states and territories.</p>	<p>Yes.</p>
<p>6. Does the program portfolio have an appropriate balance of awards to different types of institutions?</p> <p>The IIS COV found a reasonable distribution of awards across different types of institutions. This distribution of awards includes institutions across the spectrum of types considered. The data suggests that institutions across institution types have a reasonable opportunity for funding in the proposal competitions. As would be expected, this distribution is skewed more heavily towards institutions that grant doctoral degrees, with even greater proportions of awards going to research intensive institutions. We also wondered how Minority Serving Institutions were considered as part of the data collected.</p>	<p>Yes.</p>
<p>7. Does the program portfolio have an appropriate balance of awards to new and early-career investigators?</p> <p>The IIS COV saw a reasonable balance of awards to new and early-career investigators. The funding rates for new versus prior proposers have been nearly at parity. The rate of successful awards to proposals is relatively constant with respect to the proposer's career stage. IIS has granted proportionally more awards to early-career proposers, such as through robust participation in the CAREER and CRII programs. This proportion is aligned with IIS's encouragement of cultivating scholars at every career stage.</p> <p>However, the proportion of funding for CAREER awards has dropped, when we compare 2014 to 2018. There was also a disparity between the IIS programs. Each program saw a very large increase in submissions but two of them (III and RI) saw a decline in percent funded.</p> <p>This decrease in funding submitted proposals is a very serious concern, since these awards are extremely important for young academics. In many institutions, the CAREER program has become an implied requirement for successful (i.e. tenurable) faculty. The IIS COV also found a major challenge that IIS has seen decreases in funding available per capita to proposing</p>	<p>Yes.</p>

<p>PIs. The global competitiveness of the nation, such as for the growing area of AI, is particularly affected by per capita funding available for IIS programs. Greater availability of per capita funding to proposers would better enable the field to better advance both short-term innovation and longer-term exploration of ideas with promise.</p>	
<p>8. Does the program portfolio include projects that integrate research and education?</p> <p>The IIS COV agreed with the criteria used by IIS for projects that integrate research and education. The effectiveness of the assessment of these criteria in the merit review process would make the successes of IIS that much greater. The IIS COV saw in their eJackets that clearer connections could be made between the proposed research and its engagement with education. The expectations set forth by this criterion should be made clearer to proposers and reviewers.</p>	Yes.
<p>9. Does the program portfolio have appropriate participation of underrepresented groups³?</p> <p>The IIS COV was impressed by the efforts and responsiveness of IIS towards improving the inclusion of underrepresented groups in the portfolio. The problem of equitable representation in computing is a primary challenge for the nation, especially with the rapid growth of impact and opportunities in our field. IIS has instituted a number of promising procedures for more effective inclusion and equitability by bridging gaps between core programs and broadening participation in computing. Most notably, selected programs now request a “BPC plan” for submitting proposals that describe how the project will specifically and measurably conduct activities to better include underrepresented groups. The IIS COV was encouraged by the sense of dedication to broadening participation by IIS and is optimistic about the potential for positive effects resulting from their efforts. A looming issue, however, is how these efforts towards broadening participation in computing will become codified into the ethos of IIS beyond the terms of the current division administration. Such a codification of broadening participation would greatly benefit from greater consideration of proposal merit review criteria at levels beyond IIS.</p>	Yes.
<p>10. Is the program relevant to national priorities, agency mission, relevant fields and other constituent needs? Include citations of relevant external reports.</p> <p>CISE staff show a strong commitment to aligning with national priorities and other needs. CISE is well-positioned to contribute to the growing importance of computing & information research and technologies to the economy, national security, community safety, cybersecurity, environmental protection, and other topics.</p>	Yes.

³ NSF does not have the legal authority to require principal investigators or reviewers to provide demographic data. Since provision of such data is voluntary, the demographic data available are incomplete. This may make it difficult to answer this question for small programs. However, experience suggests that even with the limited data available, COVs are able to provide a meaningful response to this question for most programs.

<p>Recent White House reports, e.g. Executive Order: Maintaining American Leadership In Artificial Intelligence (February 2019) and the National Artificial Intelligence Research and Development Strategic Plan: 2019 Update (June 2019) provide a good foundation for NSF future priorities. A useful direction would be to adjust future funding to deal with emerging national concerns such as disinformation, deepfakes, online bullying, surveillance, privacy violation, cyber-attacks/scams, and so on. IIS can be a leader in these important human-centered themes, by raising the profile of human-computer interaction research, including safe and effective use of robots and automation in life critical applications such as driving, aviation, and medical devices.</p>	
<p>11. Additional comments on the quality of the projects or the balance of the portfolio:</p> <p>It is difficult to assess the aggregate quality of projects, but we felt that the proposals we read were thoughtful and well-conceived. There are so many balancing factors to consider, but NSF staff are deeply engaged and eager to do what they can to promote balance across many dimensions. Modern visual analytics tools would help in monitoring these multiple balances to make adjustments over time. These visual tools would also clarify that NSF CISE is doing well in achieving realistic balance.</p>	

OTHER TOPICS

1. Please comment on any program areas in need of improvement or gaps (if any) within program areas.

We noted that there were a few highly-rated proposals that were not funded due to funding constraints of IIS. Perhaps IIS might want to explore ways of addressing this issue in the future. The major issue seems to be a lack of new funding for the IIS programs.

Getting great reviewers is key to any peer review system. It seems like this is a very difficult task. Are there ways to improve this – e.g., identifying a larger pool of potential reviewers, letting people know when reviews will be needed (it is very hard to be asked for reviews within a short timeline) or asking people who submit proposals to review at least a certain number of proposals? This will be particularly important for the more interdisciplinary proposals.

NSF should explore strategies for increasing recognition for reviewers, possibly something simple such as posting an annual web page listing (with consent) the names of those who did reviews, and possibly featuring those who did a large number of them – without disclosing the panels that those reviewers served on.

Data centers today consume about 4% of the world’s electricity production. Large-scale machine learning algorithms run on specialized GPU-based hardware tend to be particularly energy intensive. There is much interest in reducing the energy footprint of computing for both environmental and cost reasons. CISE could play a pivotal role by promoting more research on (1) energy efficient computing, especially for machine learning, (2) the use of machine learning to optimize data center energy use; (3) well-designed interfaces and visual data analytics that enable consumers, businesses and governments to understand current energy consumption and make better decisions that reduce energy use.

2. Please provide comments as appropriate on the program's performance in meeting program-specific goals and objectives that are not covered by the above questions.

Modern widely-used visual analytics tools support exploration and discovery of temporal patterns in flows of proposals, networks of PIs/institutions, emerging topics, regional patterns, and so on. These could be made use of by IIS and the larger community to reveal distributions, clusters, gaps, outliers, and various forms of anomalies to clarify these patterns to the community, journalists, and the general public.

3. Please identify agency-wide issues that should be addressed by NSF to help improve the program's performance.

Improving the Broader Impacts aspect of proposals could lead to better research, clearer societal benefits, and greater awareness of NSF contributions to national priorities. NSF might prepare more specific and concrete guidelines to writing Broader Impact statements, carrying them out, and reporting on them in annual reports, published papers, and public outreach. A suggestion is that CISE might work with other NSF Directorates (Biology and Math/Physics) who have supported the National Association for Broader Impacts (<http://broaderimpacts.net>) for the past seven years. These have suggested: specific agenda/schedule for accomplishing Broader Impacts, a minimum of 10% of budget, citations to previous/related work, and plans to publish on Broader Impacts with their off-campus partners as co-authors. Another strategy might be to require annual and final reports to include a 300-600 word report using language accessible to general audiences. NSF could feature the best of these on the NSF web site.

4. Please provide comments on any other issues the COV feels are relevant.

We have a number of suggestions about how the review process of the COV could be improved:

1. COV group members should be chosen based on the research areas the members are going to be asked to review: e.g. the IIS reviewing group had no one with neuroscience or health experience and 15 Jackets in this area.
2. To facilitate weekly calls within each group, NSF should provide conference call capabilities for each group.
3. It would be helpful if the groups were provided with instructions on how to navigate the SharePoint website since it was not clear how to access the group's annual reports initially (i.e. click on the name of the group's division at the top left).
4. The COV could provide a standard template spreadsheet for review of jackets in which some boilerplate information about PIs (such as institution, gender, minority status, demographic information, whether funded or not) can be pre-populated. (A sample of what IIS developed is attached.) This would enable COV panelists to focus on the key aspects of the eJacket reviewing process. It would also be very useful if the reviewing questions in the report template could also point to the particular documents in the eJackets that could be used to answer each question. While some do, others do not. For example, it was not clear where to look in the eJackets to identify issues of how COI was handled (II:2) or how to check for diversity of reviewers (II:3).
5. The jackets made available to the PIs could include a link to each reviewer's webpage or DBLP or Google Scholar page so that COV panelists can get a better idea of whether the reviewer is well qualified for the reviews they are tasked to write. COV panelists had to do this manually.

6. The COV asks whether PIs conflicts have been appropriately managed – but this poses a challenge for COV panelists who may not be familiar with the PI and/or the reviewers involved.
7. Every jacket contains a mass of information that COV panelists need to process. Many panelists spent a significant amount of time on the first few jackets that they read in order to develop a strategy to follow in order to assess the jackets. Panelists would be more efficient if NSF could show a list of common questions panelists are required to answer when analyzing a jacket and show a short clickable workflow for answering that question.
8. It would be **extremely** useful to a) allow panelists to download entire eJackets in one pdf, especially for panelists who did not always have access to wifi when reviewing.
9. The team also felt that tables summarizing the tables in the annual reports would be **very** useful to identify different trends in the various areas of proposer (and please add reviewer) diversity, funding, number of proposals submitted to the different programs, and so on. We had to create these ourselves in order to identify trends.
10. It would be very useful if team members could be encouraged to gain access to the SharePoint and eJacket websites as early as possible (and preferably early in the week) since four members of our team had problems with access that had to be addressed (often multiple times) by the NSF staff. We do not know how this could be fixed in advance but any ideas on that would be useful

The Committee of Visitors is part of a Federal advisory committee. The function of Federal advisory committees is advisory only. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the Advisory Committee, and do not necessarily reflect the views of the National Science Foundation.

SIGNATURE BLOCK:



For the IIS CoV Sub-committee
Julia Hirschberg
Chair