Transition Presentation

ASCAC Meeting 9-23-19

Roscoe Giles, Boston University
roscoe@bu.edu

Outline

- Charge
- Subcommittee Activities
- Preliminary Findings
- Comments on possible recommendations
- Discussion

Charge: Subcommittee Charge Elements

- **Examine**:
  - ECP lessons learned for **managing large collaborations**,
  - ASCR's historic **fundamental research investments** in applied mathematics, computer science and computational partnerships at the National Labs,
  - **new Research and Development priorities** in artificial intelligence, quantum information systems and strategic computing.
- **Recommendations (R):**
  - for **capturing the lessons learned** from ECP
  - **supporting** the software and hardware technologies and application development from ECP activities
  - informing ASCR's **future investment strategy** for its basic research programs

Subcommittee Members

Jay Bardhan, GlaxoSmithKline	Richard Lethin, Reservoir Labs (ASCAC)
Alan Edelman, MIT	David Levermore, UMD, (ASCAC)
Roscoe Giles, Boston U. (chair)	Juan C. Meza, NSF (UC Merced)
Fred Johnson, DOE/Retired	Dan Reed, U. Utah (ASCAC)
Alexandra Landsberg, DoD HPCMP	

Subcommittee Activities

- Stakeholder Interviews
- Community Meetings
- Reports
- Subcommittee meetings

Overview

Budget timeline

Beyond ECP

Timeline

ASCR budgets have increased steadily during the exascale era.

Figure reflects enacted budget except for FY20. SBIR not included.

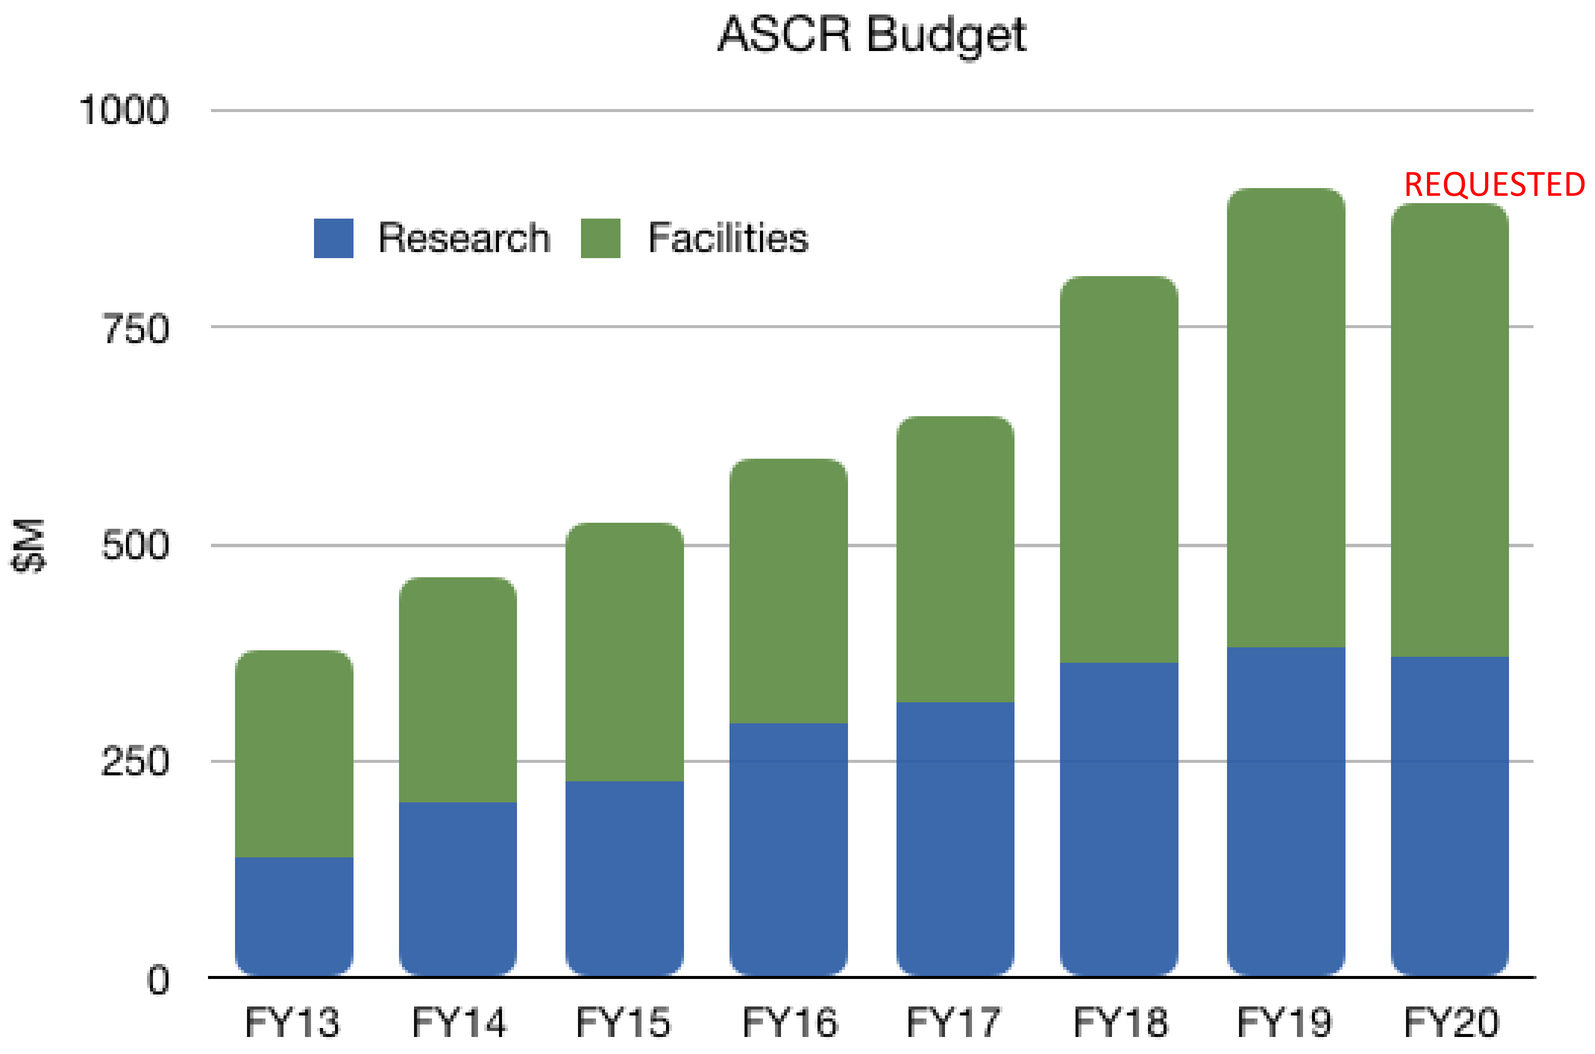

- nonECP = base research (AM, CS) + REP+...
- Funds were spent on exascale before ECP and grew during ECP

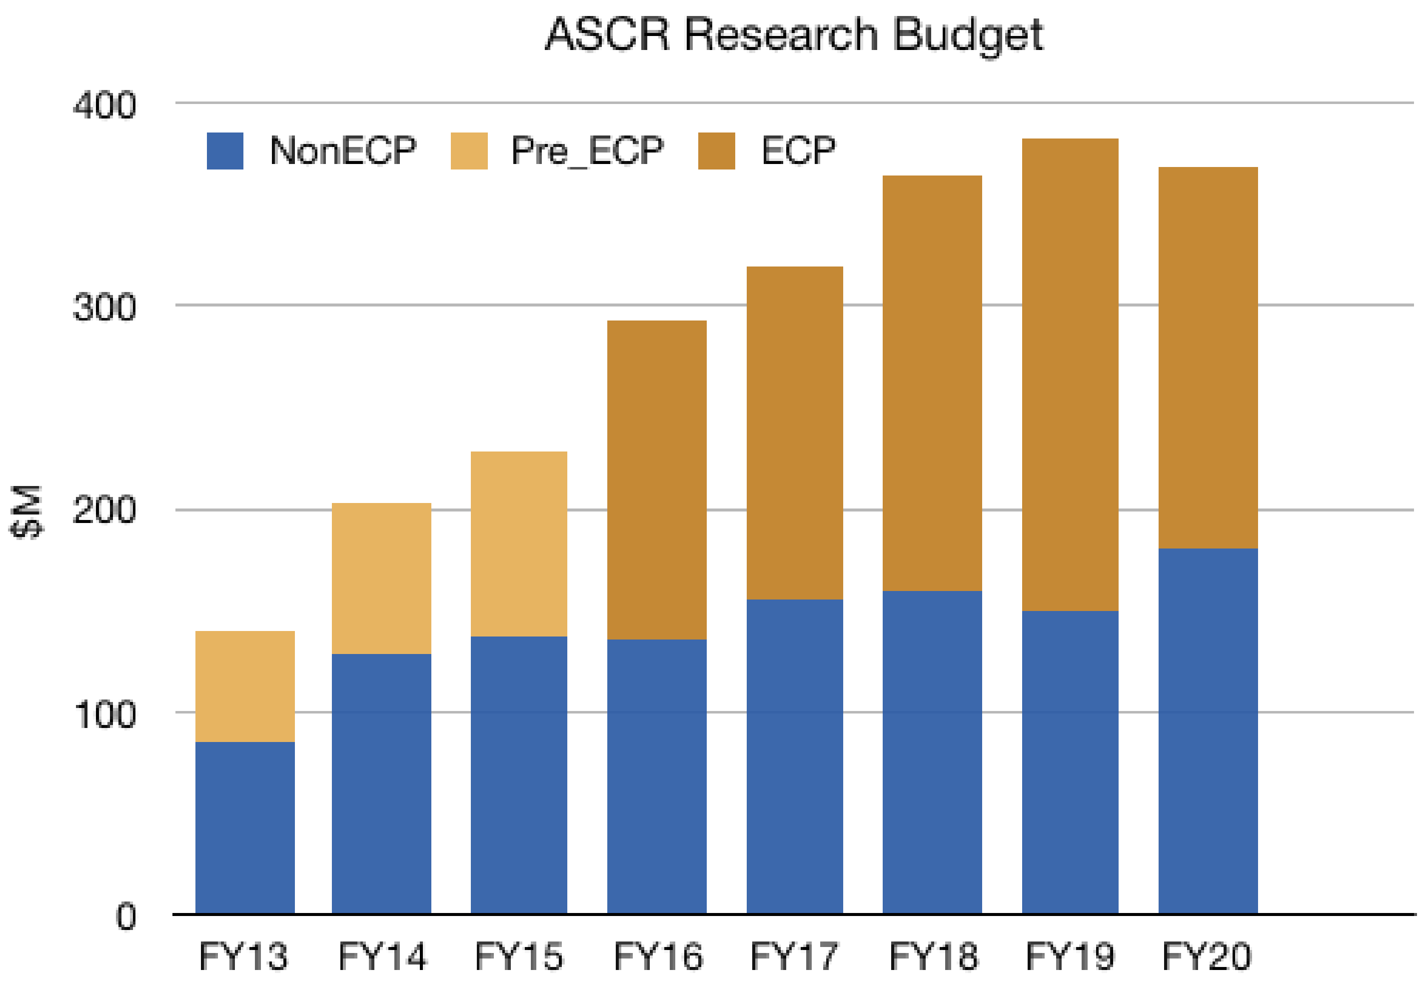

- The recent facilities increases reflect ECI activity.

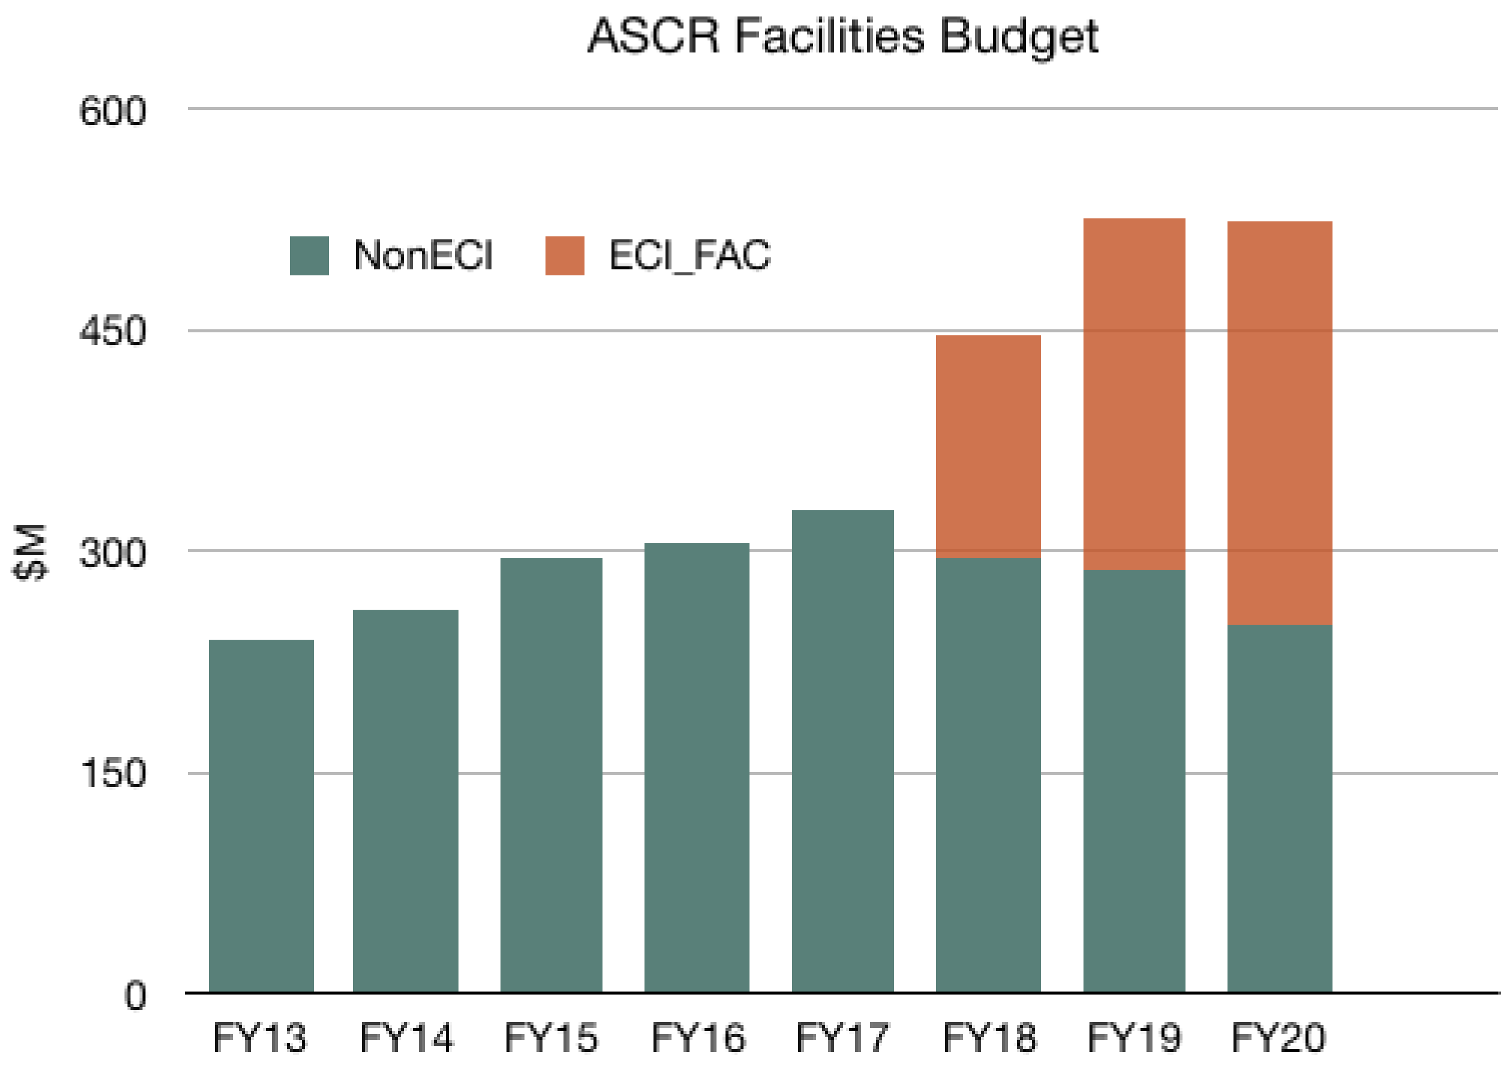

ASCR Research in the era beyond ECP

- Historically: AM + CS + CP core research programs
  - Programs handled start of ECP differently
- During ECP: strong focus on ECP related development (r/D)
- Post ECP: opportunity to: Reinvest in long term research
  - Research questions previously identified, de-emphasized, but still relevant
  - Research questions emerging from exascale computing experience and needs
  - Research questions emerging from external opportunities.
- Opportunities to foster collaborative interdisciplinary research within ASCR portfolio

Exascale in the era beyond ECP

- Continue grow application success
- Continue nurturing software ecosystem
- Building the workforce
- Continued high connectivity of ASCR supported efforts

Preliminary Findings

Overview

Examine: ECP lessons for managing large collaborations

- Complex, multi-institute, multi-stakeholder development can be successful according to all stakeholders
- When DOE labs identify shared needs, they are able to
  - leverage their buying power to improve mission capability (GitLab+CI)
  - advance existing platforms to enhance both DOE and community (Jupyter)
- When DOE labs agree on shared standards, that boosts external adoption (E4S/Spack)
- Value of Agile project management: empowering experts and stakeholders to shape Sprints and execution in agile manner
- Showcased opportunities for increasing impact of future research investments (transitioning to production)
- Higher level of connectivity (cross-X) can move the needle

Examine: ASCR’s historic fundamental research investments

- During ECP, other AM+CS+CP investments were significantly decreased, with very disparate effects that need to be harmonized during transition.
- ECP successes were critically enabled by “seed corn” of fundamental research prior to ECP; seed corn stocks now depleted.
- ASCR’s cross-cutting investments (e.g., SciDAC, CAMERA) have been very successful and there is enthusiasm for more, similar activities.
- ASCR’s investments used to be sizable fraction of R&D in fields stewarded by ASCR; situation has changed (data science, GPU, AI/ML).
- Level and stability of ASCR’s historic investments yielded a strong, highly skilled workforce; multiple factors have depleted key segments of the workforce and workforce pipeline.

Examine: new R&D priorities in artificial intelligence, quantum information systems, and strategic computing

- For NSCI, DOE is charged with delivering exascale capability
- For multiple new priorities, interests are driven not only by science mission but also economic considerations--hence, large investments from industry
- Significant classes of opportunities for ASCR, in basic research and:
  - Co-design expertise in era of Extreme Heterogeneity
  - Breadth and uniqueness of user facilities (used by both science+industry)
  - A trusted, non-commercially-driven partner (standards, security, stability)
- Large industry investments change the operating landscape
  - Workforce recruiting, retaining, and attracting mid-career professionals
  - Nature of leadership will be different from ASCR’s leadership in HPC
  - ASCR must be more agile and responsive to external landscape

Preliminary Observations

Applications, Software Technology, Hardware Integration, Management

ECP applications target national problems in 6 strategic areas

National security

Stockpile stewardship

Next-generation electromagnetics simulation of hostile environment and virtual flight testing for hypersonic re-entry vehicles

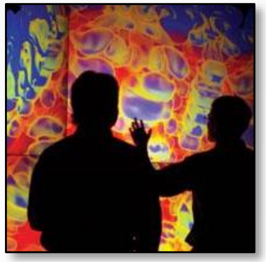

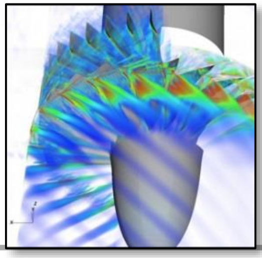

Energy security

Turbine wind plant efficiency

High-efficiency, low-emission combustion engine and gas turbine design

Materials design for extreme environments of nuclear fission and fusion reactors

Design and commercialization of Small Modular Reactors

Subsurface use for carbon capture, petroleum extraction, waste disposal

Scale-up of clean fossil fuel combustion

Biofuel catalyst design

Economic security

Additive manufacturing of qualifiable metal parts

Reliable and efficient planning of the power grid

Seismic hazard risk assessment

Urban planning

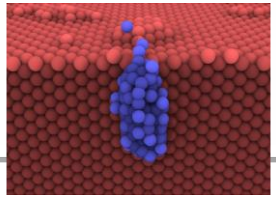

Scientific discovery

Find, predict, and control materials and properties

Cosmological probe of the standard model of particle physics

Validate fundamental laws of nature

Demystify origin of chemical elements

Light source-enabled analysis of protein and molecular structure and design

Whole-device model of magnetically confined fusion plasmas

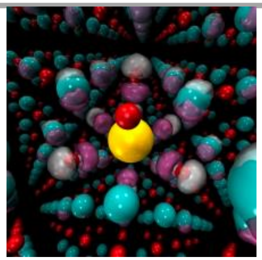

Earth system

Accurate regional impact assessments in Earth system models

Stress-resistant crop analysis and catalytic conversion of biomass-derived alcohols

Metagenomics for analysis of biogeochemical cycles, climate change, environmental remediation

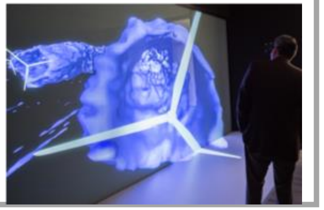

Health care

Accelerate and translate cancer research

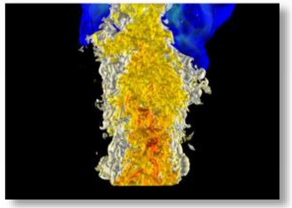

Source: DK @ ECP All Hands Meeting

ECP Impacts: Application Development

- AD teams are producing measurable advances in applications effectiveness
- AD teams experiences impact software and systems design
- Teams are interdisciplinary and link researchers/ algorithms, software, hardware context, in co-design (development).
- Progress is visible both inside ECP/ASCR and to science collaborators in DOE/SC and beyond.

ECP Lessons: Application Development

- Highly coordinated and integrated application development with software co-design is a powerful paradigm.
    - Shared goals from applications and developers $\Rightarrow$ shared responsibilities.
    - Strong team/leader awareness of overall progress.

ECP Impacts & Lessons: Software Technology

- Early success of software development kit (SDK) for math libraries resulted in robust libraries that are usable by greater community, e.g., Extreme-Scale Scientific Software Stack (E4S)
- SDK approach extended to all software technology areas which will result in long-term acceptance and use by HPC community
- These will need to be sustained beyond ECP.
- ECP software technology products primarily based on past ASCR research; new medium/long term research needed to address needs of future hardware generations

ECP Impacts and Lessons: Hardware & Integration

- ECP hardware activities collaborated with industry and facilities.
- Increased industry attention and collaboration through Xforward programs.
- Developed workable NDA processes for protecting critical vendor IP
- Extreme Heterogeneity and architectures require continuous integration.
- Close coordination of software technology with facilities yielding significant benefits

Project Management

- O 413.3B framework for management was skillfully adapted to ECP.
- Comment: L3 leadership world class, exert technical and management leadership.

ECP Impacts and Lessons: Management, Collaboration & Culture

- Agile (Scrum) very successful for ECP, for software and management
- Transparency and trust successful and valuable
- Software engineering (continuous integration, containers, …) also successful
- Strong sense that this culture, methodology, for research should be continued and expanded.
- Need to expand participation to the broader community (academics, hardware vendors, ISV, small business)

ECP Impacts and Lessons: Ecosystem

- ECP created an ecosystem:
    - Work with industry: Xforward programs
    - Application Teams for co-design
    - Systematic Software Technology infrastructure
    - Culture of collaboration
    - Value of interdisciplinary collaboration
    - Importance of trust among arms of ECP activities (ST,HI,AD)

ASCR Research

Status of AM Research

- Previously identified, still relevant
  - Exascale Applied Math (Hittinger/Dongarra, 2014), minus components funded under ECP
  - Challenges underlying MMICCS, SciDAC, RESS, UQMEESS
- Identified during ECP
  - Are there any reports/notes on this? We have discussed our expectation that there are important pieces here…
  - Exascale Applied Math report highlighted systems opportunities for AM
- Emerging independently of ECP
  - Data-driven approaches (e.g. in UQ and for modeling)
  - SciML identifies numerous AM challenges (Baker, 2019)
  - Math / software stack for post-Moore extreme heterogeneity
  - Scalable algorithms

Status of CS Research

- Diverse performers in exascale (X-Stack, systems, …) halted and defunded to take funds for ECP (“hollowed out” “butt end of difficult conversations”)
- Forward looking reports
  - Extreme heterogeneity
  - Correctness
  - Other DOE panel reports here
  - Networking and large scale distributed workflows
- Recognized need by ECP and broader CS community for critical long term research investment

Reinvigorating Long Term CS Research

- Recent reports on CS research needs and priorities
    - Report of the HPC Correctness Summit, DOE, Jan 25-26, 2017
    - Extreme Heterogeneity 2018 Workshop, January 23-25, 2018
    - Future HPC Capabilities, ASCAC, March 20, 2019
- Example areas/topics identified
  - Improve programmer productivity: expressive programming models, intelligent, domain-aware compilers, certified software
  - Enable reproducible science: validation methods for non-deterministic architectures, detection and mitigation of faults/errors
  - Facilitate data management, analytics and workflows: adapting workflows through machine learning approaches

Examples of key new future research questions

- Success of CAMERA and other edge to LCFs/NERSC examples, e.g. high energy physics, point to spectrum of computing from edge (sensors) to exascale and beyond - these come with a multitude of research challenges including mathematics and computer science research, algorithms, software technologies, networking, and workflows.

New Major Cross Disciplinary Areas

- Quantum Computing
  - Collaboration across SC (and beyond)
- AI/ML for Science
  - Active medium/long term reasearch

Workforce

Developing the Workforce: ECP impacts

Positive impacts center on building ties and trust between DOE/SC/ASCR communities, in large part through collaborative delivery of high-quality results.

Challenges center on pipeline strength at multiple stages:

- Reduced direct university engagement
- Some hiring has centered on dev-focused, not research-focused, individuals
- Limited opportunities for early career researchers to build independence/experience with blue sky, high-risk high-reward research

Developing the Workforce: Keeping ASCR at the forefront

- DOE mission is compelling and a source of pride to workforce. Keeping this highly visible will be important.
- University->lab pipeline is important and has diminished during the current era.

Vision: the labs are able to capitalize on the overall R&D&D ecosystem, in an agile way, by bringing in industry/university talent at all levels as are needed, whether for project leadership or domain expertise.

Observations on the structure of programs and projects

- Labs working in common purpose - gitlab. Jupyter
- Transition plans for projects (near term research/Dev)
- Interdisciplinary
- Tighter research/facility integration
- SciDAC like models, Centers (NNSA like PSAP?)
- Planting Seeds vs Garden vs Farm?
- Diversity of performers (academia, vendors)
- Foster high level of connections
- Collaboration - National and International

Additional Questions and Discussion

What about recommendations?

Future Events

- Continued telecoms seeking community input
- October 17: Discussion with ECP Industry Council
- October 24-25: Workshop 2 (https://www.orau.gov/ascac-tes2019/)
- November @ SC: Community discussion @BoF involving some subcommittee members
- December ASCAC meeting: final report