NEEDS AND OPPORTUNITIES IN COSMOLOGY IN THE ERA OF ULTRA-DEEP SURVEYS, Michael L. Norman, UCSD

1. *Research needs/opportunities*
    a. What are some of the open problems in your field that require large-scale simulation to solve? Which might lead to fundamental or foundational advances? Why are these problems not being solved today?

        Origin and evolution of galaxies; nature of "first light objects" and the end of the cosmic dark ages; cosmological constraints on the nature of dark matter and dark energy; star formation across cosmic time; origin of supermassive black holes and their co-evolution with galaxies, decoding the cosmic microwave background polarization signal.

    b. What are some of the open problems in your field that require data-intensive computing, such as large-scale data analytics and data mining? Why are these problems not being solved today?

        The above-mentioned problems are addressed observationally through increasingly broad and deep astronomical surveys that will yield billions of galaxies and quasars (e.g., LSST [1]) and increasingly precise maps of the cosmic microwave background polarization (e.g., PolarBear [2]). Quasar absorption line surveys provide detailed information about the intergalactic medium and dark matter properties. Statistical analysis, including 2-pt and higher correlations, is fundamental to extracting the science. Increasingly large cosmological simulations (currently petascale, future exascale) attempt to create "mock catalogs" of equal size to the survey volumes to extract physical insights about the nature of dark matter and dark energy. This requires multiple realizations, each of the size of the observable universe, for accurate parameter estimation. This is an extreme scale data- and compute-intensive grand challenge that is being attempted rather unsuccessfully at present due to the lack of powerful enough data analytics and access platforms, and long-term data preservation. Such large data sets, whether from observation or simulation, could benefit from advanced data mining capabilities to find rare objects ("needle in haystack") that could result in important and unexpected discoveries.

    c. Are there plans or roadmaps that characterize future computing needs in your field?

        Yes. Two chapters from the DOE report "Scientific Grand Challenges: Challenges for Understanding the Quantum Universe and the Role of Computing at the Extreme Scale". One chapter addresses simulation [3], the other data handling, analysis, and archiving [4]. The projections and recommendations for high end capabilities are as relevant today as they were in 2008.

d. What types of new workflows are emerging that requires complex access pathways between data sets, computation, and storage?

   In terms of time and effort, analysis of cosmological data derived from simulations now far outweighs that devoted to producing it. Analysis of cosmological data, whether from observations or simulations, use similar and in some cases identical workflows and tools: object extraction, cataloging in a database, followed by searching, filtering, and cross-correlating? Other typical operations are imaging, power spectrum analysis, time series analysis, and complex spatial queries. Cosmological parameter estimation requires inverting very large covariance matrices in high dimensional spaces. Because dataset sizes vastly exceed the memory of the HPC machine and IO is the limiting factor, new workflows are needed to efficiently manage the data movement between parallel processing and storage elements.

2. *Advanced computing capabilities, facilities, requirements*
   a. What forms of computing are used in your field? For example, how does your field make use of laptop/desktops, research group clusters, department or campus commodity cluster systems, mid- to large-scale, shared capacity systems such as XSEDE, leadership-class capability systems such as Blue Waters (NSF) or Mira (Department of Energy), or commercial cloud services such as Amazon EC2? How would you characterize the importance of access to each type—required, desirable, or unnecessary? How might these needs change in the future, and why?

      HPC clusters and supercomputers are used for cosmological simulations (e.g., N-body, AMR and SPH hydro), including the very largest: NSF Blue Waters, DOE Mira, and Titan. N-body simulations using more than 1 trillion particles have been carried out in the past few years using the entirety of the very largest supercomputers. Analysis of such data is bogged down because these same HPC systems are ill-suited to the analysis phase [4]. Pioneering data-intensive supercomputer architectures like JHU's DataScope [5] and SDSC's Gordon [6] are extensively used for cosmological data analysis, and point the way to more balanced systems in the future. Commercial cloud services are a negligible factor for the simulation side which requires low latency interconnects between large numbers of processing nodes. Commercial clouds represent a superior architecture for the data analysis side, but adoption has been minimal due to end-user cost (IO and storage charges.)

   b. How are datasets evolving in terms of variety and distribution? Do you access tens to hundreds of near-real-time data sets? Do you rely on a few large repositories?

      Current large cosmological surveys and simulation campaigns generate a data corpus of about 1 PB each. Next generation experiments will increase this number 10-20 fold [3, 4]. Individual file sizes range from kilobytes to terabytes. There is no need for near real time datasets on the simulation side. However observation is moving into the time-based domain, where transient events will signal alerts which may trigger downstream analyses and conceivably simulations.

Large observational surveys (e.g., LSST) have sufficient budget to create their own curated persistent data archives and science access centers. Mid-scale observational surveys (e.g., DES, PanSTARRS) struggle to create similar capabilities, and typically leverage externally funded resources (e.g., XSEDE, NERSC) to support them. The simulation side is left to fend for itself as far as data preservation and data repositories are concerned. Both the DOE INCITE and NSF Blue Waters programs delete the data within a year after the computer allocation expires, and do not provide infrastructure for the creation of community facing repositories that serves up data and tools to exploit the unique simulation data over the longer term. An emerging worldwide phenomenon is the creation of Virtual Cosmology Observatories, such as the one pioneered by the Millenium Simulation [7] in Germany, where researchers can access data products from high value simulations over the web.

c. With computer hardware and software evolving more rapidly than in the recent past, what impacts do you see for your field? For example, what role will new hardware such as accelerators (GPUs or Intel Xeon Phi), FPGAs, new memory systems, or new I/O systems play? Are there barriers to their adoption, such as challenges making necessary modifications to software?

   I expect increasing use of hybrid CPU/GPU and many core CPU architectures in the future. Users will be insulated from hardware complexities through layered software design and abstraction such as object-oriented programming provides. The HACC [8] and 2HOT [9] cosmology N-body codes have already been developed that exploit GPGPUs across the full 18K node TITAN machine. Neither of these codes is publically available, but their technical approach is published in the open literature and readily implemented in community codes. The ENZO hydrodynamic cosmology code [10] uses GPGPUs to accelerate part of the computation, but not all. The ENZO community is developing an extreme scale follow-on called ENZO-P [11] that will port with minimum effort to any future high end system provided CHARM++ is ported to those platforms. This code is publically available, and insulates users and application developers from the details of the underlying parallel architecture through the CELLO AMR infrastructure layer [11]. CHARM++ supports in-RAM checkpoint/restart, and could easily be extended to exploit burst buffers.

d. What software does your field depend on? Who develops and maintains this code, and how is this work supported?

   The cosmology simulation workhorse codes (N-body and hydro) are ENZO, GADGET, ART, RAMSES, NYX, HACC, HOT2, AREPO, GASOLINE, and CHANGA. ENZO and GADGET have large, robust developer communities. The rest are developed by smaller collaborative groups and typically not widely shared. ENZO has evolved into a self-supporting open source community code, with new features being developed "for free" by dev-users. Initial development of ENZO-P/CELLO has been supported by NSF grants. NYX, HACC, and HOT2 have been developed with DOE institutional funds and grants.

A suite of CMB data analysis software and tools [12] is developed and maintained by the C3 group at LBNL.

e. Is your field keeping up the technical skills needed to use new technical capabilities?

Not uniformly. Some are at the leading edge, some are following, the majority is being left behind. Only software developers really need to know what is going on at the high end.

3. *Challenges and suggestions*

a. What are the biggest challenges that your field faces in using computation? Consider access to systems with sufficient capability and capacity; productivity of environments; algorithms; workforce; stability of software and hardware; and the ability to use systems efficiently, including parallelism and scalability.

- Inability to derive full value from massive simulation data sets due to short archive retention times and the lack of appropriate data access platforms
- Capacity and capability of HPC systems
- Access to community codes developed for the high end of simulation
- Availability of scalable, extensible, data analysis tools
- Porting legacy workhorse codes to new HPC platforms

b. What investments would have the greatest positive impact on your research field? For example, this could be more computer systems to increase access, different kinds of systems with a different balance of capability, support for community software, development of new algorithms, or a workforce with better training in computational science.

Coordinated investments in two complementary kinds of scalable high end systems: compute-intensive (traditional HPC), and data-intensive (e.g., DISC concept described [4]). These are fundamentally different architectures with different balance points, but would be used together, as discussed in [3]. Figure is reproduced here for convenience.

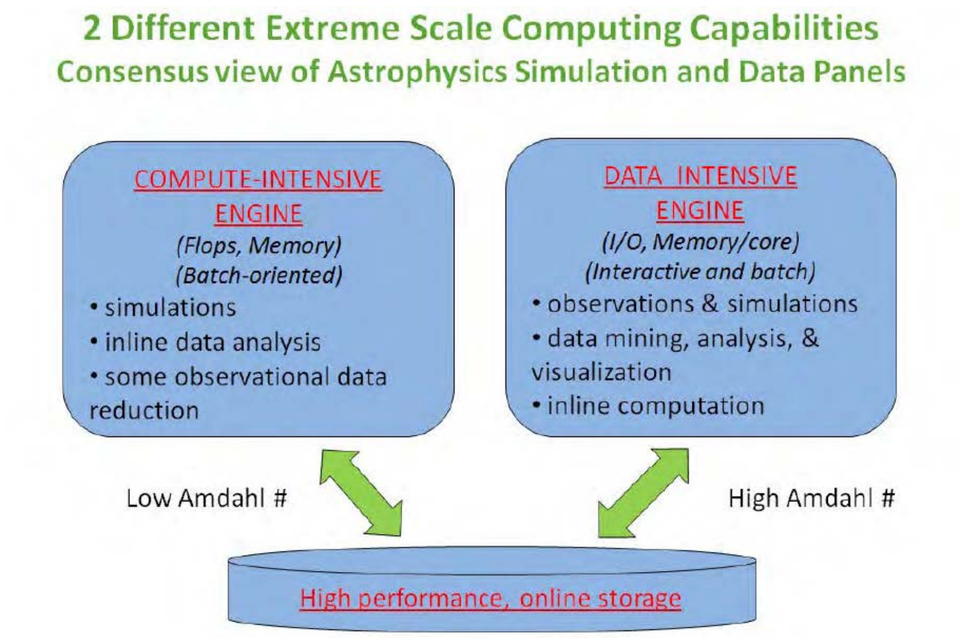

Figure 5. Two different extreme scale computing capabilities' needs for extreme scale astrophysics simulations, simulations data analysis, and observational data analysis.

c. What other elements of national cyber infrastructure would significantly advance the pace of discovery or expand participation? Examples might include shared file systems or standard services and application program interfaces.

- Data discovery/sharing/analysis/publishing/preservation environments for datasets of all size, including current and coming multi-PB datasets
- User- and developer-friendly community application software that masks underlying hardware complexity
- Gateways and portals to simulation and data tools

[1] LSST homepage, http://www.lsst.org/lsst/
[2] PolarBear homepage, http://bolo.berkeley.edu/polarbear/
[3] "Astrophysics and Cosmology Simulation", a chapter in DOE report: Scientific Grand Challenges: Challenges for Understanding the Quantum Universe and the Role of Computing at Extreme Scale
http://extremecomputing.labworks.org/highenergyphysics/reports/HEPreport101609_final.pdf
[4] "Astrophysics Data Handling, Analysis, and Mining", ibid
[5] "New JHU Computer To Enable Data Analysis Not Possible Today", Johns Hopkins University press release
[6] SDSC Gordon homepage, http://gordon.sdsc.edu
[7] Millenium Simulation homepage, http://www.mpa-garching.mpg.de/galform/millennium-II/index.html

[8] S. Habib, V. Morozov, N. Frontiere, H. Finkel, A. Pope, and K. Heitmann. HACC: extreme scaling and performance across diverse architectures. http://doi.acm.org/10.1145/2503210.2504566.
[9] M. S. Warren. 2HOT: an improved parallel hashed oct-tree n-body algorithm for cosmological simulation http://doi.acm.org/10.1145/2503210.2503220.
[10] Enzo homepage, http://enzo-project.org
[11] Enzo-P/Cello homepage, http://cello-project.org
[12] CMB Data Analysis at NERSC, http://crd.lbl.gov/departments/computational-science/computational-cosmology-center/c3-research/cosmic-microwave-background/cmb-data-analysis-at-nersc/