Advanced LIGO MREFC Project – Status Update and No-cost Extension

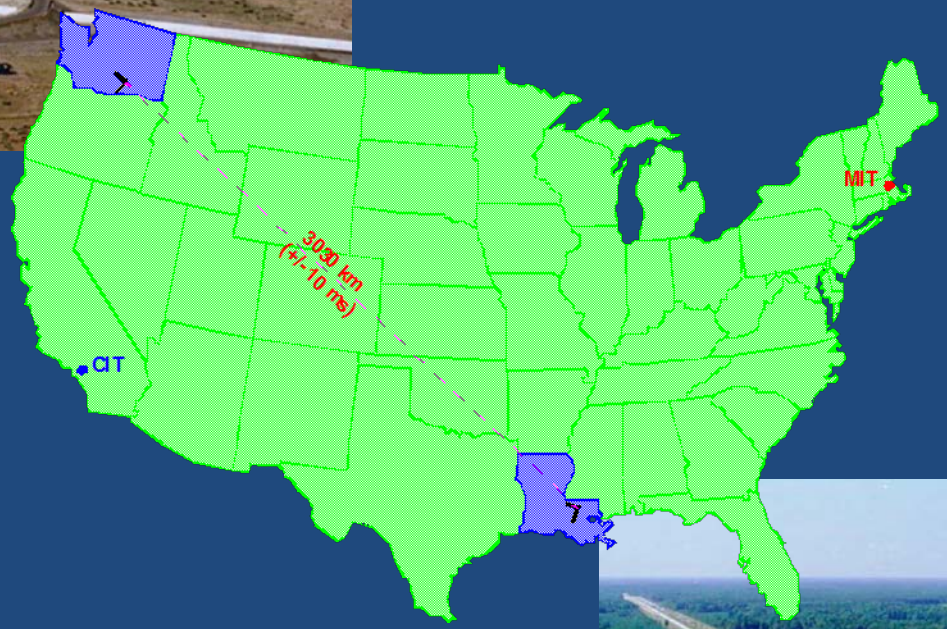

Mark Coles
Physics Division
MPS

Advanced LIGO Computing Plan

- **2008 MREFC Proposal:** Estimated computing needs by scaling initial LIGO operation to Adv LIGO bandwidth, with plan to buy hardware (HW) prior to end of construction – *March 2015* – subject to NSF review and approval
- **Current view**: Additional computing capacity not needed until Fall 2017 science run

No Cost Extension of MREFC award March 31, 2015 → Sept. 30, 2017

- Delays major purchases of computing hardware until they are needed, in Fall 2017, to benefit from continued vendor performance/price improvements (Moore's Law)
- Funds continued activity to optimize software (SW), evaluate HW accelerators, investigate shared computing platforms
- Provides for continued brief monthly status reports and annual reviews of computing scope

About the 2014 Computing Review

- **Background:**
  - Data analysis plans expanded greatly after 2008, far exceeding planned computing capacity
  - In 2013, PHY encouraged LIGO to assess use of ACI-funded resources for some long latency computing & report evaluation at 2014 NSF review
  - LIGO/XSEDE tests, SW optimizing produced >3X speed-up!
- **Review:**
  - Panel included ACI-recommended supercomputing experts
  - Recommended LIGO/ACI interactions continue, just-in-time HW purchases, other action
- **LIGO actions following Review:**
  - Prioritized data analysis goals + HW requirements
  - Hired supercomputing expert to lead SW optimization
  - Investigating hardware accelerators for further gains

Advanced LIGO MREFC Status

- ***ALL THESE MREFC ACTIVITIES ARE COMPLETE!***
  - Interferometer HW increases strain sensitivity ~10X and reduces low frequency cutoff from ~50 Hz → ~5 Hz, + storage of 3rd interferometer (~$195M)

- **NCE data analysis hardware + software activity:**
  - Optimize short latency software for EM triggers + long latency searches for complex waveforms. Just-in-time hardware buys. (~$10M)

- **PHY oversight actions maximizing Science/$:**
  - Pushed LIGO evaluation of XSEDE
  - Partnered with ACI/CISE to evaluate LIGO computing
  - Directed LIGO prioritization of science goals + computing needs
  - Directed LIGO develop 2015-17 NCE workplan
  - Implemented oversight framework for on-going assessment

Supporting material

LIGO: Laser Interferometer Gravitational-Wave Observatory

- Two Michelson interferometers, 2000 miles apart, that detect ripples in fabric of space-time due to stellar collisions, supernovae, etc.
- Proposed to NSF in 1989 as a two-step program:
  - **Initial LIGO:** Construction of infrastructure and initial interferometer elements that could plausibly detect GWs
  - **Advanced LIGO:** Upgrade based on experience + additional R&D with sensitivity that makes detections likely

Data and Computing Scope in Advanced LIGO Proposal

- Scaled detection bandwidth of LIGO → Adv LIGO, ~10X computing capacity
- Estimates simulations at 3X detection search computing needs
- Applied Moore's law to budget estimates for purchases in 2014, scaling from 2007
- Assumed 50% of computing resources provided by LIGO Science Collaboration
- Assumed 400 TB/yr data production, 3 copies

Current picture of LIGO computing

- 2017 computing requirements are ~10X higher than 2015 or 2016, motivating NCE.
- Deferring purchases until needed, while aggressively optimizing code and investigating alternate architectures.
  - 4000 cores to support site clusters and meet short and long latency needs in 2015 and 2016
  - ~9000 cores additional in 2017

Panel recommendations for alternate architecture investigations

- Double precision NVIDIA GPUs (in 2014: K20 or K40): While these may not be optimal for the project (see below) they may be installed in general purpose clusters (like TACC's Stampede) where they serve a broader community that requires double precision performance.
- Single precision NVIDIA GPUs (in 2014: K10). These may be optimal if dedicated hardware is purchased.
- Intel Xeon Phi: These may deliver lower performance than GPUs but are ubiquitous in the Stampede cluster and can be used 'free of charge' if Stampede resources are utilized. Experience in the community has shown that performance optimization targeting the Xeon Phi often results in enhanced performance on Xeon CPUs as well.