7. Acknowledgements

Indiana University's Big Red supercomputer was funded in part by the Indiana METACyt Initiative and is supported in part by a Shared University Research grant from IBM Inc. to Indiana University. The Indiana METACyt Initiative is supported by a grant from the Lilly Endowment, Inc. to Indiana University. Big Red is connected to the NSF-funded TeraGrid. IU's participation in the TeraGrid is funded by National Science Foundation grant numbers 0338618, 0504075, and 0451237. The IU Data Capacitor was funded by grant number CNS-0521433. IU's work with LEAD has been supported by NSF grants ATM 0331480 and EIA-0202048. CICC was supported by P20HG003894-01. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation (NSF) or the National Institutes of Health (NIH). The authors wish to thank the people involved in building CTSS software on Big Red: Eric Blau, Charles Bacon, JP Navarro (University of Chicago), and Jason Brechin (National Center for Supercomputer Applications). The MyCluster component of CTSS was built by Evan Turner (Texas Advanced Computing Center) and the SRB client was built by George Kremenek (San Diego Supercomputer Center). Thanks to David Hart and Nancy Wilkins-Diehr (SDSC) and Suresh Marru (IU) for providing information and statistics related to the TeraGrid, LEAD, and Science Gateways generally. We would like to thank the staffs of the Research Technologies Division of UITS, the Pervasive Technology Institute at IU, and ZIH at Technische Universität Dresden, for their contributions and effort that made the implementation of Big Red and the Data Capacitor, their integration within the TeraGrid, and the performance analysis possible. The authors would particularly like to thank Ray Sheppard and Malinda Lingwall for writing and editing. Gratitude goes particularly to IU President Michael A. McRobbie, whose leadership made Big Red's acquisition and implementation possible.

8. References

[1] IBM, Inc. IBM BladeCenter JS21: The power of blade innovation. http://www.redbooks.ibm.com/redbooks/pdfs/sg247273.pdf

[2] IBM, Inc. IBM BladeCenter. http://www-03.ibm.com/systems/bladecenter/

[3] Accelerating Growth: Indiana's Strategic Economic Development Plan. 2006. http://www.in.gov/iedc/files/Strategic_Plan.pdf

[4] University of Tennessee at Knoxville. HPC Challenge. http://icl.cs.utk.edu/hpcc/

[5] National Science Foundation. Benchmarking Information Referenced in the NSF 05-625 "High Performance Computing System Acquisition: Towards a Petascale Computing Environment for Science and Engineering". 2005. http://www.nsf.gov/pubs/2006/nsf0605/nsf0605.pdf

[6] The Weather Research & Forecasting Model (WRF). Home page. http://www.wrf-model.org/

[7] Cable, SB and E D'Azevedo. Benchmarking OOCORE, an Out-of-Core Matrix Solver. 2001. http://www.ornl.gov/~webworks/cppr/y2001/misc/123564.pdf

[8] GAMESS. Home Page. http://www.msg.ameslab.gov/GAMESS/

[9] MIMD Lattice Computation (MILC) Collaboration. Home page. http://www.physics.indiana.edu/~sg/milc.html

[10] National Energy Research Scientific Computing Center. PARAllel Total Energy Code (PARATEC). Home page. http://www.nersc.gov/projects/paratec/

[11] University Corporation for Atmospheric Research. High-Order Methods Modeling Environment (HOMME). Home page. http://www.image.ucar.edu/CMG/Software/HOMME/