

**Center for Advanced Research Computing
Internal Advisory Board
Tuesday, February 28, 2017, 3 p.m.**

MINUTES

Present: **Patrick Bridges**, Interim Director of CARC; **Gregory Taylor**, Professor, Physics and Astronomy; **Bridget Wilson**, Professor, Pathology; **Hua Guo**, Professor, Chemistry; **Brian Petrowicz**; Interim Deputy CIO/Information Technologies; **Thomas Turner**, Arts & Sciences Associate Dean for Research and Professor, Biology; **Jane Lehr**, Professor, Electrical Computer Engineering; **Patricia Henning**, Interim Vice President for Research and Professor, Physics and Astronomy; **Barbara McCrady**, Director of CASAA and Professor, Psychology; **Tracy Wenzl**, CARC Unit Administrator

Present by phone: **Jeremy Edwards**, Professor, Chemistry

1. CARC Overview/Status – discussion of data presented under the following categories:
 - New projects, users, and publications
 - Personnel status
 - Budget
 - Computer systems
 - Hardware/software infrastructure

2. Technical System and Support Plans – discussion of data presented under the following categories:
 - Personnel – CARC Champions, Student Employees
 - i. Suggestion: online user forums for support concerns/Piazza
 - ii. Suggestion: adopt a cost center model (who could be billed, how would charges be assessed?) UNM says there has to be a policy on cost (rate sheet). Consider calculating CARC expenses as percentage of total staff listed on grant, percentage of equipment used for maintenance charge. Suggestion: create a spreadsheet that automatically calculates cost for PIs (but need formulas to do this.)

ACTION ITEM: Brian Petrowicz to share information about how IT does charge backs
 - iii. Suggestion: consider running grants through CARC instead of home departments. Concerns about F&A sharing, labor expenses.

- iv. Suggestion: See what peer institutions are doing to “charge” for their services
 - Hardware System Architecture
 - System Software Architecture
 - Relationship with Central IT, A&S IT
 - Funding Opportunities
3. Programmatic Strategic Planning – discussion of:
- CSE Program Status, Thoughts
 - i. Suggestion: CSE certificate could be used to train internal UNM staff, possibly through a training grant
 - ii. Suggestion: pursue an MRI grant & create work group
 - Outreach to Campus Units
 - i. Suggestion: Complete a needs assessment to see what campus work isn’t being done. Make business case for funding needed to address needs.
 - 1. Need = export control solutions
 - 2. Need = classified solutions (not much appetite to address this/not consistent with UNM mission)
 - Suggestion: IT Faculty Research Support Officer
 - Subject area focus groups
4. Additional IAB discussions, feedback – covered under previous topics

Meeting adjourned, 4:30 p.m.

**Center for Advanced Research Computing
Internal Advisory Board
Tuesday, February 28, 2017**

AGENDA

1. CARC Overview/Status
 - New projects, users, and publications
 - Personnel status
 - Budget
 - Computer systems
 - Hardware/software infrastructure

2. Technical System and Support Plans
 - Personnel – CARC Champions, Student Employees
 - Hardware System Architecture
 - System Software Architecture
 - Relationship with Central IT, A&S IT
 - Funding Opportunities


3. Programmatic Strategic Planning
 - CSE Program Status, Thoughts
 - Outreach to Campus Units
 - Subject area focus groups

4. Additional IAB discussions, feedback

CARC fact sheet

Budget

- CARC received an allocation of \$668,000 for FY 17. Combined with carry-forward and money left over from SuperComputing conferences, the total budget for FY 17 is **\$739,666.68**.



- The majority of CARC's budget goes to employee salaries. CARC employs:
 - A Director
 - Three administrative staff – a full-time Unit Administrator, a part-time Administrative Assistant 3, and a part-time Marketing Assistant.
 - Four technical staff – an HPC Systems Specialist 3, a Technical Analyst 2, a Systems/Network Analyst 3 and a Systems/Network Analyst 1.
 - 2 graduate student Project Assistants will be hired in the near future and will support technical staff

Users

- In the last six months, CARC has registered over 130 new users working on 45 new research projects and/or participating in one of six academic courses

Projects

CARC new research projects since 2016	
PI	Project
Anil Prinja	A Computationally Efficient Method for High Energy Charged Particle Transport
Irene Salinas	Atlantic salmon microbiome
Michael Andersen	Avian phylogenomics
Brandon Schmandt	Back-projection of large -N array seismic data for microseismicity detection
Shana Judge	Can Emergency Department Diagnoses Identify Women at Risk for Exposure to Unhealthy Sexual Behavior?
John Grey	Charge transfer complexes
Sal Portillo	Charged particle Beam Devices
Thomas Turner	Comparative Transcriptomics of Fishes
Hua Guo	Computational studies of chemical reactions catalyzed by proteins
Mark Gilmore	Coronal Mass Ejection and Spheromak Modeling of the HelCat Plasma Device
Manuel Martinez-Ramon	Deep learning for emotion recognition
Kristina Trujillo	Differential Gene Expression in Cancer
Marianne Berwick	Differential Splicing by Sex in DNA Repair Genes
Tang-Tat Ng	Discrete Element Modeling of Drilled Shafts in Granular Materials
Trilce Estrada	Distributed and in-situ analysis for multidimensional structured data
Martin Kirk	Electronic Structure of Molecular Wires
Arash Mafi	Exploration of optical rogue phenomena in dielectrics as a function of the intrinsic randomness or disorder
Lee Taylor	Fungal Ecology
CM Adema	Gastropod comparative genomics
Jeremy Edwards	Genomic Sequencing and Analysis
Zhen Peng	High-performance integral equation solvers for electromagnetic simulations
Mark Stone	Investigation of floodwave attenuation in the Middle Rio Grande
Adam Hecht	LAMMPS radiation damage simulations of CaF2
Dorian Arnold	Large Scale Systems Research
Robert Miller	Marsupial transcriptome analysis
James Degnan	Measuring imbalance of gene trees
Mousumi Roy	Melt migration in continental interiors
Diana Northup	Microbial Ecology of Subterranean Environments
Peter Vorobieff	Modeling a fully three-dimensional shock interaction with a gaseous density interface
Joseph Cook	molecular evolution of mammals
Vanessa Goss	Monte Carlo Simulation for Weak Neutron Sources
Alejandro Manjavacas Arevalo	Plasmons in low dimensional systems
Lydia Tapia	Protein Folding through Motion Planning
Greg Taylor	Pulsar Searching
Ivan Deutsch	Quantum Information
Melanie Moses	Robot Swarm CPFA Simulation
Davood Tofghi	Robustness Checks for Popular Complex Mediation Tests

CARC new research projects since 2016, continued	
PI	Project
Svetlana Poroseva	Rotocraft Aerodynamics
Timothy Clark	Simulation of Compressible Turbulent Mixing
Robert Rhatigan	Small Area Population Estimates
Huining Kang	Statistical methods for high throughput genomic data analysis
Keith Lidke	Super-Resolution and Single Particle Tracking
Dimiter Petsev	Transport in nanachannels
Donald Natvig	UNM Biology Sevilleta Metagenomics
Wade Wilson	Western Chub

Courses

Courses using CARC resources 2016-2017	
Instructor	Course
Thiruvarangan Ramaraj	BIOL519/CHEM587
CM Adema	BIOL446L/546L
Walter Gerstle	CE598
Dorian Arnold	CS587
Andree Jacobson	CS442/542
Trilce Estrada	CS 591

Publications

- In 2016 and 2017, 44 papers have been published using research done at CARC.

Faculty publications

2016-2017

2017

Dixon, D.S., Mozumder, P., Vásquez, W.F. and Gladwin, H., 2017. Heterogeneity Within and Across Households in Hurricane Evacuation Response. *Networks and Spatial Economics*, pp.1-36.

Stone, M.C., Byrne, C.F. and Morrison, R.R., 2017. Evaluating the Impacts of Hydrologic and Geomorphic Alterations on Floodplain Connectivity. *Ecohydrology*.

Kolb, B., Luo, X., Zhou, X., Jiang, B. and Guo, H., 2017. High-dimensional Atomistic Neural Network Potentials for Molecule-Surface Interactions: HCl Scattering from Au (111). *The Journal of Physical Chemistry Letters*.

Vangara, R., Brown, D.C.R., van Swol, F. and Petsev, D.N., 2017. Electrolyte solution structure and its effect on the properties of electric double layers with surface charge regulation. *Journal of Colloid and Interface Science*, 488, pp.180-189.

2016

- Coutsias, E.A., Lexa, K.W., Wester, M.J., Pollock, S.N. and Jacobson, M.P., 2016. Exhaustive Conformational Sampling of Complex Fused Ring Macrocycles Using Inverse Kinematics. *Journal of Chemical Theory and Computation*, 12(9), pp.4674-4687.
- Xie, C., Jiang, B., Yang, M. and Guo, H., 2016. State-to-State Mode Specificity in $F+CHD_3 \rightarrow HF/DF+CD_3/CHD_2$ Reaction. *The Journal of Physical Chemistry A*, 120(33), pp.6521-6528.
- Manjavacas, A., 2016. Anisotropic optical response of nanostructures with balanced gain and loss. *ACS Photonics*, 3(7), pp.1301-1307
- Benson, B.W., Iyer, S.R., Kemper, K. and Zhao, J., 2016. Director Networks and Credit Ratings.
- Qi, X., Baragiola, B.Q., Jessen, P.S. and Deutsch, I.H., 2016. Dispersive response of atoms trapped near the surface of an optical nanofiber with applications to quantum nondemolition measurement and spin squeezing. *Physical Review A*, 93(2), p.023817.
- Shahbazikhah, P., Kalivas, J.H., Andries, E. and O'Loughlin, T., 2016. Using the L1 norm to select basis set vectors for multivariate calibration and calibration updating. *Journal of Chemometrics*.
- Ard, S.G., Johnson, R.S., Martinez, O., Shuman, N.S., Guo, H., Troe, J. and Viggiano, A.A., 2016. Analysis of the Pressure and Temperature Dependence of the Complex-Forming Bimolecular Reaction $CH_3OCH_3+Fe^+$. *The Journal of Physical Chemistry A*.
- Manavi, K., Jacobson, B., Hoard, B. and Tapia, L., 2016. Influence of model resolution on geometric simulations of antibody aggregation. *Robotica*, 34(08), pp.1754-1776.
- Hoard, B., Jacobson, B., Manavi, K. and Tapia, L., 2016. Extending rule-based methods to model molecular geometry and 3D model resolution. *BMC Systems Biology*, 10(2), p.48.
- Tencate, A., Kalivas, J.H. and Andries, E., 2016. Penalty processes for combining roughness and smoothness in spectral multivariate calibration. *Journal of Chemometrics*, 30(4), pp.144-152.
- Cavanagh, J.F. and Castellanos, J., 2016. Identification of canonical neural events during continuous gameplay of an 8-bit style video game. *NeuroImage*, 133, pp.1-13.
- Chen, T., Lehr, J., Lavrova, O. and Martinez-Ramonz, M., 2016, July. Distribution-level peak load prediction based on Bayesian Additive Regression Trees. In *Power and Energy Society General Meeting (PESGM), 2016* (pp. 1-5). IEEE.
- Manjavacas, A., Rodríguez-Fortuño, F.J., de Abajo, F. and Zayats, A.V., 2016. Lateral Casimir force on a rotating particle near a planar surface. *arXiv preprint arXiv:1612.03733*.
- Lin, J., Wester, M.J., Graus, M.S., Lidke, K.A. and Neumann, A.K., 2016. Nanoscopic cell-wall architecture of an immunogenic ligand in *Candida albicans* during antifungal drug treatment. *Molecular biology of the cell*, 27(6), pp.1002-1014.
- Martinez-Herrero, R. and Manjavacas, A., 2016. Basis for paraxial surface-plasmon-polariton packets. *Physical Review A*, 94(6), p.063829.

- Moradi, C.P., Xie, C., Kaufmann, M., Guo, H. and Douberly, G.E., 2016. Two-center three-electron bonding in ClNH_3 revealed via helium droplet infrared laser Stark spectroscopy: Entrance channel complex along the $\text{Cl} + \text{NH}_3 \rightarrow \text{ClNH}_2 + \text{H}$ reaction. *The Journal of chemical physics*, 144(16), p.164301.
- Abaie, B., Mobini, E., Karbasi, S., Hawkins, T., Ballato, J. and Mafi, A., 2016. Random lasing in an Anderson localizing optical fiber. arXiv preprint arXiv:1612.03835.
- Kolb, B. and Guo, H., 2016. Communication: Energy transfer and reaction dynamics for DCI scattering on Au (111): An ab initio molecular dynamics study.
- Lin, S., Gao, H.W. and Peng, Z., 2016, May. High-fidelity, high-performance full-wave computational algorithms for intra-system EMI analysis of IC and electronics. In *Signal and Power Integrity (SPI), 2016 IEEE 20th Workshop on* (pp. 1-4). IEEE.
- Lushnikov, P.M., 2016, September. Nonlinear combining of multiple laser beams in multimode optical fiber. In *Australian Conference on Optical Fibre Technology* (pp. JM6A-15). Optical Society of America.
- Shi, Y. and Lee, J.H., 2016. Sample size calculations for group randomized trials with unequal group sizes through Monte Carlo simulations. *Statistical Methods in Medical Research*, p.0962280216682775.
- Wheeler, L.B., Galewsky, J., Herold, N. and Huber, M., 2016. Late Cenozoic surface uplift of the southern Sierra Nevada (California, USA): A paleoclimate perspective on lee-side stable isotope paleoaltimetry. *Geology*, 44(6), pp.451-454.
- Winner, K.R.K., Steinkamp, M.P., Lee, R.J., Swat, M., Muller, C.Y., Moses, M.E., Jiang, Y. and Wilson, B.S., 2016. Spatial modeling of drug delivery routes for treatment of disseminated ovarian cancer. *Cancer research*, 76(6), pp.1320-1334.
- Flehart, M. E., van Swol, F., & Petsev, D. N. (2016). Solvent Role in the Formation of Electric Double Layers with Surface Charge Regulation: A Bystander or a Key Participant?. *Physical Review Letters*, 116(4), 048301.
- Han, H., Alday, B., Shuman, N.S., Wiens, J.P., Troe, J., Viggiano, A.A. and Guo, H., 2016. Calculations of the active mode and energetic barrier to electron attachment to CF_3 and comparison with kinetic modeling of experimental results. *Physical Chemistry Chemical Physics*, 18(45), pp.31064-31071.
- Cowen, B.J. and El-Genk, M.S., 2016. Bond-order reactive force fields for molecular dynamics simulations of crystalline silica. *Computational Materials Science*, 111, pp.269-276.
- Kerketta, R., Halász, Á.M., Steinkamp, M.P., Wilson, B.S. and Edwards, J.S., 2016. Effect of spatial inhomogeneities on the membrane surface on receptor dimerization and signal initiation. *Frontiers in Cell and Developmental Biology*, 4.
- Cowen, B.J. and El-Genk, M.S., 2016. Probability-based threshold displacement energies for oxygen and silicon atoms in α -quartz silica. *Computational Materials Science*, 117, pp.164-171.
- Song, H., Lu, Y., Li, J., Yang, M. and Guo, H., 2016. Mode specificity in the $\text{OH} + \text{CHD}_3$ reaction: Reduced-dimensional quantum and quasi-classical studies on an ab initio based full-dimensional potential energy surface. *The Journal of chemical physics*, 144(16), p.164303.

- Martin, T.D., Hill, E.H., Whitten, D.G., Chi, E.Y. and Evans, D.G., 2016. Oligomeric Conjugated Polyelectrolytes Display Site-Preferential Binding to an MS2 Viral Capsid. *Langmuir*, 32(47), pp.12542-12551.
- Shao, Y., Wang, S. and Peng, Z., 2016, March. Hierarchical modeling and scalable algorithms for in-situ characterization of 3D IC packages. In *Wireless Information Technology and Systems (ICWITS) and Applied Computational Electromagnetics (ACES)*, 2016 IEEE/ACES International Conference on (pp. 1-2). IEEE.
- Ox, J., 2016. Color Systems are Categories that Carry Meaning in Visualizations: A Conceptual Metaphor Theory Approach. *Electronic Imaging*, 2016(16), pp.1-9.
- Zhao, B., Sun, Z. and Guo, H., 2016. A reactant-coordinate-based approach to state-to-state differential cross sections for tetratomic reactions. *The Journal of Chemical Physics*, 145(18), p.184106.
- Song, H., Li, A. and Guo, H., 2016. Rotational and Isotopic Effects in the $H_2 + OH \rightarrow H + H_2O$ Reaction. *The Journal of Physical Chemistry A*, 120(27), pp.4742-4748.
- Adair, J., 2016. Reconstructing the historical Albuquerque reach of the Middle Rio Grande to evaluate the influence of river engineering on floodplain inundation (Doctoral dissertation).
- Gustafson, G.T., 2016. Phylogenetics and diversification of whirligig beetles (Coleoptera, Gyridae) (Doctoral dissertation).
- Silantyev, D.A., Lushnikov, P.M. and Rose, H.A., 2016. Langmuir wave filamentation in the kinetic regime. II. Weak and Strong Pumping of Nonlinear Electron Plasma Waves as the Route to Filamentation. arXiv preprint arXiv:1610.10071.
- Zhao, B., Sun, Z. and Guo, H., 2016. State-to-state mode selectivity in the $HD + OH$ reaction: Perspectives from two product channels. *The Journal of chemical physics*, 144(21), p.214303.
- Samedi, V.G. and Bocklage, T., 2016. Breast Cytology. In *Pitfalls in Diagnostic Cytopathology With Key Differentiating Cytologic Features* (pp. 1-41). Springer International Publishing.

UNM CARC Spring 2017 IAB Technical Update

Patrick G. Bridges

Current Technical Computing Support

CARC provides staff, hardware, and software to support computation-based research in the UNM community. Most of these capabilities are in the form of shared resources provided to all of campus, but a number are dedicated resources hosted for specific research groups.

Personnel

- 3.5 Technical Staff (Hussein Al-Azzawi shared with Psychology)
- Staff handle all system installation, maintenance, and support, technical interactions with UNM networking and IT staff, machine room maintenance management, user troubleshooting and code optimization, user training, etc.
- Paid for largely out of baseline budget, some portions covered by grant funds
- User interactions growing with a potential growing support backlog

Hardware

- Wide range of systems, some provided and administered, some hosted
- Newest systems are small, largest systems are aging
- Diversity of systems makes for a significant support load

Administered Systems

Machine Name	Nodes	Cores	RAM/Node	Peak FLOPS (theoretical), in TFLOPS
Poblano	1	64	256 GB shared	0.179
Metropolis	70	140	8 GB	0.728
Nano	35	140	16 GB	1.288
Pequena	22	176	16 GB	2.1
Ulam	120	960	16 GB	4.2
Galles	200	400	2 GB	2.2
Gibbs	24	384	64 GB	3.996
Xena	32	576	64 GB, 512 GB, 1 TB	18
Oort	1	20	1TB	
Total	505	2860		32.691

Administered Storage Systems

- NFS Home Directories (20TB usable)
- Research Storage Consortium Public System (~2-300TB usable)
- RSC Phase 4 (Crandall)
- GlusterFS Temporary File System (~36TB usable)

Hosted Systems

- University Libraries VM systems (VMs)
- LWA Data Archive (Storage)
- Data One Cluster (VMs, Storage)
- CREATE nodes
- CRF cluster (VMs, Storage)
- ARTSLab Skyscan system (defunct)

Software

Most of the system software setup is a legacy setup, including our password/user system (NIS), our project management and ticketing system (NIS), and our infrastructure services (DHCP/DNS/etc.) Little effort or planning has previously been made for modernizing this.

Machine Room

2 full hot/cold aisles with space for 22 standard racks per aisle, multiple AC units and uninterruptible power supplies for the nodes.

Upcoming Changes

- SGI XE1300 cluster system arriving from LANL on Thursday 3/2/17. 26 racks (4 will be spares). When turned on:
 - Wheeler: 504 nodes, 4032 cores, 12TB total RAM, 42TF total compute
 - CARC now: 505 nodes, 2800 cores, 10TB total RAM, 32TF total compute
- Tentatively named Wheeler, will be stood up in stages over coming months, resulting in the near-immediate retirement of Metropolis, Pequena, and Nano and the eventual retirement of Ulam (Retire 8TF of compute capacity for 42TF of new compute capacity)
- Planning new condo-model machine (Gibbs follow-on) tentatively names Taos, to include GPU, storage, and large-memory nodes for users with custom demands, potentially leveraging limited use of VMs for customization.

Ideal Technical Computing Support

Challenge: Current systems running near capacity, current staff at or above capacity, large number of various hosted systems with custom needs due to research, price, etc. Both hardware and software infrastructure is aging.

Personnel Plan

- Short Term: Hire student employees to assist technical staff, one focused on system support and one on user support
- Medium Term: User Workshops and training/reference material – train users in bulk and provide them with high quality materials to help them get started with less one-one-one staff involvement
- Medium Term: CARC Champions – recruit and train main campus IT staff who are the front-line contacts for users in units to provide basic technical support for users using CARC systems. Requires figuring out ticket system integrations
- Long Term: Need research/system grants and/or a new CARC funding model to grow staff

Hardware Plan

CARC should ideally provide 4 system types (only!) to users:

1. Limited-access (but still shared-resource) capability systems acquired through grant and condo-model funding for users with specialized needs. Xena and Gibbs are currently this
2. Open access capacity systems – a single large shared resource that can handle significant computing capacity for the campus community. This is currently a combination of Metropolis, Nano, Pequena, Ulam, and Galles
3. High-performance/high-capacity storage for research projects that supports and manages different data performance, liveness, and safety needs at appropriate cost points. Currently done through a combination of RSC, home directory, and temporary storage.
4. Virtualization for supporting custom research needs and CARC infrastructure (not done!)

Software Plan

Current software architecture is functional but antiquated; limits system reliability, ability to collaborate with users/staff on main campus, ability to support custom needs well

- Deploy basic virtualization infrastructure in collaboration with UNM IT to host infrastructure servers at CARC with appropriate redundancy/failover
- Deploy LDAP/Kerberos for authentication to machine service/login nodes instead of NIS
- Deploy modern cluster management tools for diskless clusters, most likely OpenHPC with Slurm from Intel/LBL/IU/etc. with help of staff at New Mexico Consortium.
- Replace AIRS user/project/ticket system with ticket system that can collaborate/share tickets with main campus and a new user/project management system
- Deploy better node/system monitoring tools as part of switch to new cluster management system

- Expand virtualization support to be able to host custom VMs for PIs, virtualize login nodes to better support science gateways, interact with VM-based cloud services