2023 RCEC Annual Meeting

Monday, May 22

Chris Marianetti, *Chair of SRCPAC*
Alexander Urban, *Co-Chair of SRCPAC*
Marc Spiegelman, Chair of the Foundations for Research Computing Advisory Committee

RCEC Agenda

Welcome & Introductions

- Chris Marianetti, Chair of SRCPAC
- Alexander Urban, Co-Chair of SRCPAC

High-Performance Computing Update

- Chris Marianetti, Chair of SRCPAC

Long Term Strategic Thinking about University Needs for Computing and Storage

- Jeannette Wing, EVP for Research

Foundations for Research Computing Update

- Marc Spiegelman, Chair of the Foundations for Research Computing Advisory Committee

High Performance Computing Updates

Chris Marianetti, *Chair of SRCPAC*
Alexander Urban, *Co-Chair of SRCPAC*

HPC Summary

- The Shared HPC service at Columbia continues to provide a consistent service.
- Personnel changes during the pandemic have been challenging, now the HPC team is fully staffed.
- We continue to manage and monitor Data Center capacity for the near future, as the strategic plan initiative continues.

Shared High Performance Computing

Providing Shared Compute
Since 2012

More than

- 20 Million jobs run
- 400 Million core hours of compute provided

***Currently** more than*

- 570 Compute Nodes
- 15,968 Cores
- 518 TFlops
- 2.3 Petabytes of Storage

Faculty-led Governance

Currently more than 4900 users.

More than
180 Group and Department shares

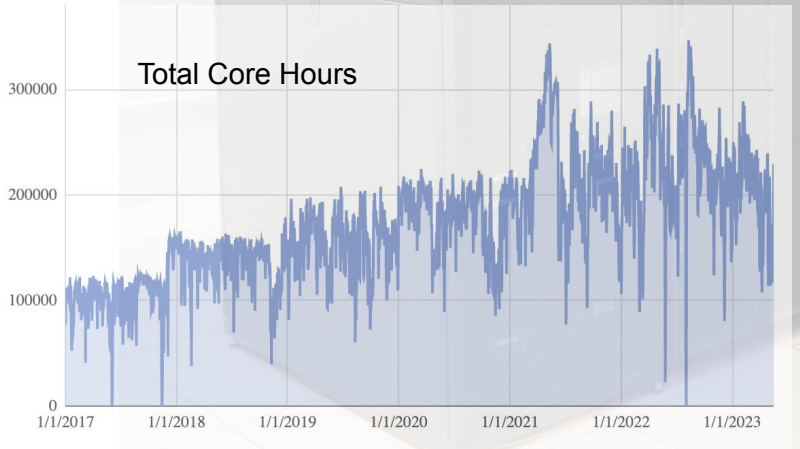

Since 2017

Edu Tier

Total Users: 530 students

Total Use: 1,744,956 core hours

Free Tier

Total Users: 305

Total Use: 1,846,479 core hours

Introductory training offered

COLUMBIA UNIVERSITY
INFORMATION TECHNOLOGY

Current HPC Footprint

	Status	Nodes	Cores	Total $	Comments
Habanero	Retired June 1 2023				Maintaining a portion as free/edu tier
Terremoto Phase I	Active until Dec 2023	110	2640	$1.3M	
Terremoto Phase II	Active until Dec 2024	27	648	$318K	
Ginsburg Phase I	Active until Dec 2025	139	4448	$1.4M	
Ginsburg Phase II	Active until Dec 2026	99	3168	$1.07M	
Ginsburg Phase III	Active until Dec 2027	48	1536	$637K	
Manitou (GPU)	About to go live	13	1248	$754K	
Insomnia					Currently taking orders

HPC 2023 Purchase Round - Schedule

- Buy-in period to join new cluster open through **June 16, 2023**
  - **Prices are significantly higher than last year due to supply chain issues and upgrade to the next generation processor**
- Purchase Order to be issued in July/August 2023
- Go-live of new equipment planned for late Fall 2023
- Significant delivery delays for GPU orders are expected

HPC 2023 Purchase Round - Pricing Menu

	LAST YEAR	2023 estimate
Standard Server (512 GB)	$7,404	$14,105
High Memory Server (1 TB)	$14,922	$16,892
GPU server with 2 x A40	$16,808	$21,340
GPU server with 2 x A100	$25,661	$29,774

Servers Feature

Dual Xeon Platinum 8640Y+ processors (2 GHz, 40 cores each, 80 cores per server), 512 GB Memory

This is a significant increase in cores and memory over last year's model (80 cores vs 32 cores)

Prices Include

- Infrastructure-related costs
- Networking
- Scheduling software
- 5-year support and maintenance

RFQ Committee (confidential)

RFQ Committee convened early March
Members: Kyle Mandli, Julia Hirschberg, Bob Mawhinney, Rob Lane

- Specifications were agreed upon and the RFQ was sent to 15 vendors
- The finalists for the final round have been narrowed down to
  - Penguin Systems
  - Dell
  - Lenovo
- Open order period should start May 22 - June 16
- Menu and final pricing to be distributed early this week

Data Center Capacity

Rack	Loc	2022	2023	2024	2025	2026	2027	2028	AND SO ON...
1	L18	HABA1	GPU Manitou	GPU Manitou	GPU Manitou	GPU Manitou	GPU Manitou		
2	L19	HABA1	GPU Manitou	GPU Manitou	GPU Manitou	GPU Manitou	GPU Manitou		
3	L21	HABA1		*insomnia (estimate)*	*insomnia*	*insomnia*	*insomnia*	*insomnia*	
4	L22	HABA2		*insomnia*	*insomnia*	*insomnia*	*insomnia*	*insomnia*	
5	L24	HABA2		*insomnia*	*insomnia*	*insomnia*	*insomnia*	*insomnia*	
6	i18	MOTO1	MOTO1	*insomnia*	*insomnia*	*insomnia*	*insomnia*	*insomnia*	
7	i19	MOTO2	MOTO2	MOTO2					
8	i21	MOTO2	MOTO2	MOTO2					
9	i22	BURG1	BURG1	BURG1	BURG1				
10	i24	BURG1	BURG1	BURG1	BURG1				
11	i25	BURG1	BURG1	BURG1	BURG1				
12	i27	BURG2	BURG2	BURG2	BURG2	BURG2			
13	i28	BURG2	BURG2	BURG2	BURG2	BURG2			
14	L25		BURG3	BURG3	BURG3	BURG3	BURG3		
15	L27				*insomnia 2 (estimate)*	*insomnia 2*	*insomnia 2*	*insomnia 2*	*insomnia 2*
16	L28				*insomnia 2*	*insomnia 2*	*insomnia 2*	*insomnia 2*	*insomnia 2*
	U7	FREE TIER nodes will be maintained outside of the HD Racks							
	U8	FREE TiER							
	U9	FREE TiER							
	U10	FREE TiER							

Data Center Capacity

- INSOMNIA is a placeholder name for the next cluster. Reference in the rack diagram are estimates only.
- With CPUs packing in more cores, we are seeing increases in power draw per node, which may require more rack space to house the same number of nodes.
- For the short term (next few years), we should able to manage by
  - Rotating out old equipment
  - Strategically managing and consolidating racks

Long-Term Strategic Thinking about University Needs for Computing and Storage

Jeannette Wing, *Executive Vice President for Research*

Charge

The Office of the EVP for Research and CUIT charges this committee: **to recommend a strategic plan for the University's future computational and data infrastructure for research**. The committee should consider all the major elements of this plan, including but not limited to:

- Compute resources
- Data resources for analysis, sharing, storage, archiving, privacy and security
- Technology skills required
- Policy impacts
- High level cost implications

The committee should provide its recommendations by the end of Academic Year 2023-24 (May 31, 2024).

Faculty Committee

The faculty committee represents a diversity of users and needs from schools and institutes across the University:

- ***Hod Lipson, co-chair***, James and Sally Scapa Professor of Innovation in the Department of Mechanical Engineering; Co-Director, Maker-Space Facility; School of Engineering and Applied Sciences
- ***Darcy Peterka, co-chair***, Senior Research Scientist; Scientific Director of Cellular Imaging; Mortimer B. Zuckerman Mind Brain Behavior Institute
- **Timothy Berkelbach**, Associate Professor of Chemistry; Faculty of Arts and Sciences
- **Roisin Commane**, Assistant Professor of Earth and Environmental Sciences, Atmospheric Composition Group, Lamont Doherty Earth Observatory; Faculty of Arts and Sciences
- **Wojciech Kopczuk**, Professor of Economics and of International and Public Affairs; Faculty of Arts and Sciences and School of International and Public Affairs
- **Ciamac Moallemi**, William von Mueffling Professor of Business; Columbia Business School
- **Muredach Reilly** (interim), Florence and Herbert Irving Endowed Professor of Medicine; Director, Irving Institute for Clinical and Translational Research; Associate Dean for Clinical and Translational Research; Columbia University Irving Medical Center

Faculty Committee

The committee is now expanding. Recently added members include:

- **Seth Cluett,** Lecturer; Faculty of Arts and Sciences
- **Siddhartha Dalal,** Professor of Professional Practice in Applied Analytics; Faculty of Professional Studies
- **Frantz Merine,** Chief Information Officer, Information Technology; School of Law
- **Robert Pincus,** Lamont Research; Lamont-Doherty Earth Observatory
- **Julien Teitler**, Professor of Social Work; School of Social Work

Faculty Committee (cont.)

Ex-officio members are:

- **Robert Cartolano**, Associate Vice President for Technology and Preservation, Columbia University Libraries
- **Gaspare LoDuca**, Chief Information Officer and Vice President for Information Technology, Columbia University Information and Technology
- **Alexander Urban**, Assistant Professor of Chemical Engineering, School of Engineering and Applied Sciences and Chair of the Columbia Shared Research Computing Policy Advisory Committee (SRCPAC)
- **Jeannette M Wing**, Executive Vice President for Research; Professor of Computer Science, Office of Research

Staffing the committee will be:

- **Maneesha Aggarwal**, Assistant Vice President, Academic, Emerging Technologies & Research Services, Columbia University Information and Technology
- **Sophie Thuault-Restituito**, Chief of Staff and Executive Director for Special Projects, Office of Research

Foundations Update

Marc Spiegelman, Chair of the Foundations for Research Computing Advisory Committee
Chris Marianetti, *Chair of SRCPAC*
Alexander Urban, *Co-Chair of SRCPAC*

Foundations Mission

Foundations for Research Computing provides **informal training** for Columbia University graduate students and postdoctoral scholars to develop fundamental skills for harnessing computation: core languages and libraries, software development tools, best practices, and computational problem-solving.

Purpose: to provide the investment in people and computational skills required to compliment our investment in hardware, software and systems administration

Foundations Primary Activities

- **Novice trainings**: 2 day training based on Software Carpentry curriculum for novice learners, learning Git, UNIX, and either R or Python
- **Data Club**: revamping of Python Users Group: twice-monthly meeting for those using computation in their research or interest about specific, more advanced topics
- **Intermediate intensives**: 1 day training for intermediate learners
- **Workshops**: 1.5 - 2 hour training opportunity to advance computational skills in a group setting. Workshops are often led by partners including CUIT and the Libraries

Novice Training Bootcamps

- 12 Bootcamps since Aug 2018 (2-3/year)
- Half were remote due to Covid – remote format presented challenges, particularly at Novice level
- Return to in-person, January 2023

Novice Training Data

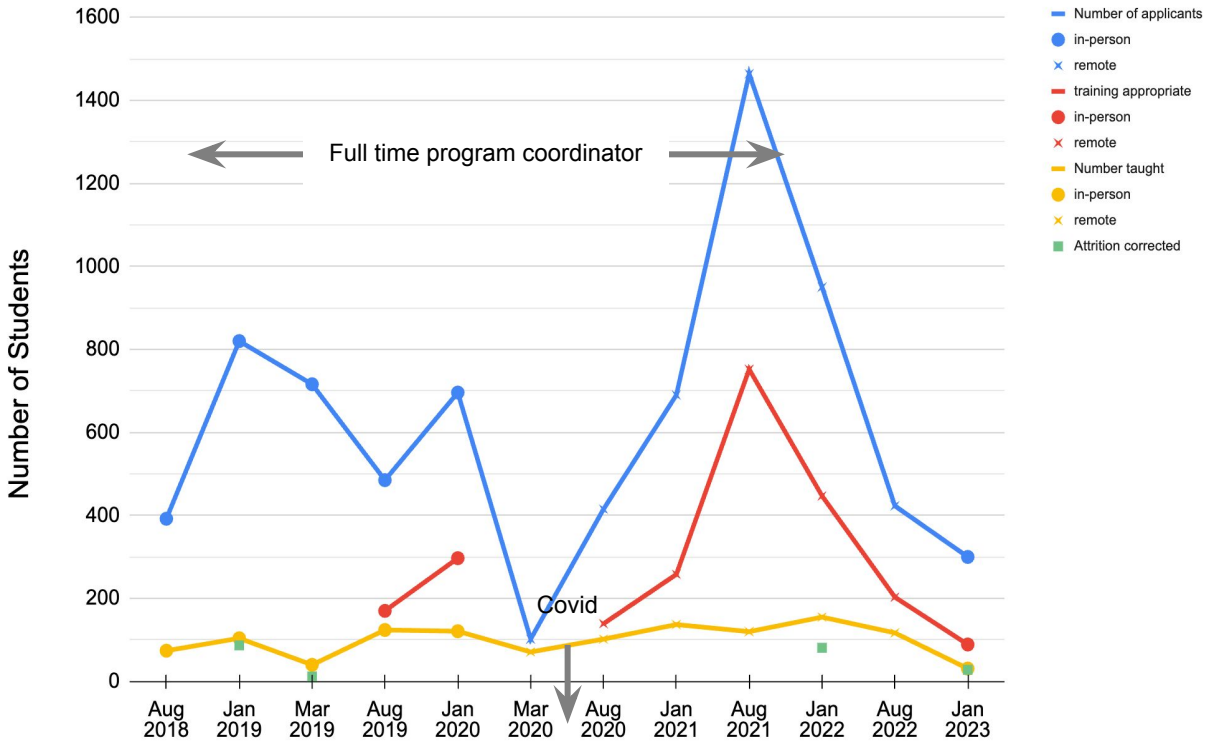

Some Observations

- Demand always exceeds supply
- Even when filtered for background.
- Novice training is extremely labor intensive – challenging to scale
- Identifies considerable demand for more advanced training
- All of this requires a full-time program coordinator

Spring 2023 Novice Training (31 participants)

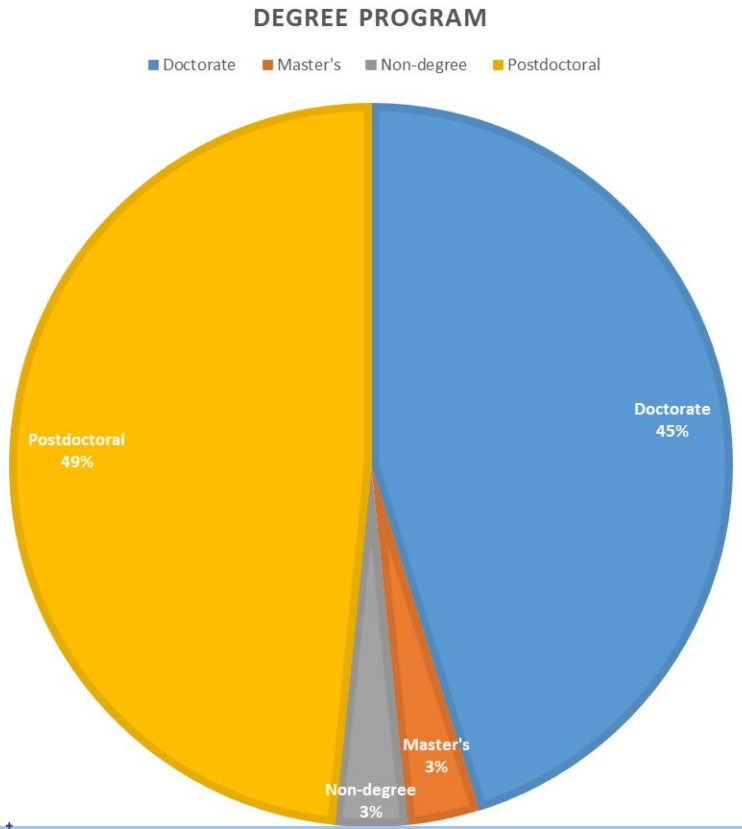

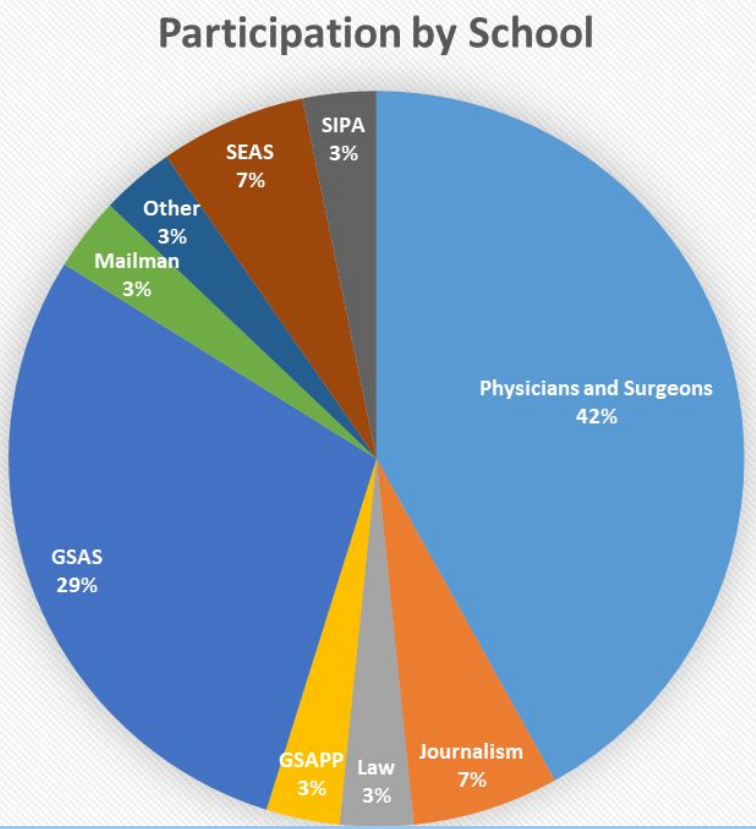

Summary/Conclusions

- The need and rationale for Foundations has not changed
- But the mechanics/structure requires review with all stakeholders
- Now is particularly timely, given new potential hires
- SRCPAC should be a natural place to seek new leadership
- Happy to take any questions

SRCPAC Chair position

We would like to thank **Chris Marianetti** for his eight years of service as Chair of the Shared Research Computing Policy Advisory Committee. He has been an active and engaged Chair, providing strong leadership and guidance. He will be missed.

We would like to welcome **Alex Urban** as the incoming Chair of SRCPAC. As a long-time user of the Shared High Performance Computing Service at Columbia, we are looking forward to benefiting from his experience and knowledge.

Thank you