

Development of Large-Scale Scientific Cyberinfrastructure and the Growing Opportunity to Democratize Access to Platforms and Data

Jakob Luettgau¹[0000-0002-3672-206X], Giorgio Scorzelli²,
Valerio Pascucci²[0000-0002-8877-2042], and Michela Taufer¹[0000-0002-0031-6377]

¹ University of Tennessee, Knoxville, TN, USA
{luettgauj,taufer}@acm.org

² University of Utah, Salt Lake City, UT, USA
valerio.pascucci@utah.edu

Abstract. As researchers across scientific domains rapidly adopt advanced scientific computing methodologies, access to advanced cyberinfrastructure (CI) becomes a critical requirement in scientific discovery. Lowering the entry barriers to CI is a crucial challenge in interdisciplinary sciences requiring frictionless software integration, data sharing from many distributed sites, and access to heterogeneous computing platforms. In this paper, we explore how the challenge is not merely a factor of availability and affordability of computing, network, and storage technologies but rather the result of insufficient interfaces with an increasingly heterogeneous mix of computing technologies and data sources. With more distributed computation and data, scientists, educators, and students must invest their time and effort in coordinating data access and movements, often penalizing their scientific research. Investments in the interfaces' software stack are necessary to help scientists, educators, and students across domains take advantage of advanced computational methods. To this end, we propose developing a science data fabric as the standard scientific discovery interface that seamlessly manages data dependencies within scientific workflows and CI.

Keywords: Scientific data · Cyberinfrastructure · National Science Data Fabric.

1 Introduction

Computational sciences in the last decade experienced rapid adoption throughout a wide range of domain sciences. Many computational methods that only a few decades ago required supercomputers evolved and matured into tools available to the research community and the general public. Two trends, in particular, are contributing to the adoption. First, miniaturization, innovation, and economies-of-scale make vast amounts of computational power available both

in small form factors and at an affordable price. Second, ubiquitous access to fast internet and wireless/cellular services allows sharing of remote computational resources, offering opportunities for better utilization and enabling new usage models such as near real-time integration of supercomputing into field-work. Miniaturization, innovation, and economies-of-scale allowed transferring methodologies that used to require supercomputers and applying them to new problems outside the data center and at affordable costs. What used to be a supercomputer 40 years ago is routinely carried in our pockets as a smartphone. However, miniaturized computing platforms can not overcome all constraints to actual problem size, time to solution, or energy requirements. These constraints are largely lifted by ubiquitous access to fast internet that allows using of custom context as data input and returns the results of computations performed in the cloud or at an HPC site.

The technologies to democratize **access to platforms** exist today, but their adoption is progressing at different paces for different communities and regions. Access to compute and networking infrastructure alone is not sufficient; overcoming barriers to **access to data** and to the use of distributed systems are similarly crucial. While students today can reproduce many computations that made history, it can be not easy to find existing datasets [18, 25] and be able to process the retrieved data to answer scientific questions.

We argue that both finding data and using advanced computational methods, while associated with many technical challenges, today is also a challenge in making the proper application programming interfaces (API) available to the scientific community. A workflow to answer a seemingly simple research question may contain numerous tasks and process large amounts of data requiring a long time for the execution on a single computer. Modern data science libraries allow for hiding a lot of the complexity of scaling to additional resources. Still, the scientist must be comfortable navigating a fragmented landscape of computing providers, best practices, and technical jargon [21, 19, 22]. The navigation process can be a significant barrier for users whose priority is scientific discovery.

In this paper, we systematically analyze the development of two large-scale cyberinfrastructures (i.e., computing facilities and network infrastructure) and how the roll-out of network technologies, public clouds, and new models of providing access to compute resources allow a broader audience to leverage technologies that originated in HPC.

The contributions of this paper are:

- An analysis of compute capabilities over time about the capabilities in modern devices widely available today to understand the constraints for applications with and without remote compute capabilities.
- An analysis of the roll-out of different networking infrastructures with a special focus on the distribution and capabilities of research networks.
- A discussion of the combined perspective to better understand where efforts to improve the software ecosystem and aspects of human-computer interaction (HCI) provide an opportunity to help democratize access to platforms and data.

The remainder of this document is structured as follows: In Section 4, we discuss related work that also develops an overview of networking and computing research infrastructure. In Section 2, we consider the development and distribution of compute capabilities and relate them to applications and capabilities possible on current consumer electronics. In Section 3, we consider the development and distribution of national and education networks (NRENs) and other connecting nations, organizations, and individuals. We discuss and summarize the study’s results in Section 5.

2 Development and Distribution of Compute Capabilities

We analyze the development and the geographic distribution of computing capabilities and break our analysis into two focus areas:

- The history of supercomputing and equivalent of computational power in different modern devices through seven eras characterized by the FLOPS ceiling Section 2.1 to understand the development of computational power in different platforms over time.
- The geographic distribution of computing and data resources to understand which regions have direct access and where centers of advanced computing innovation are located in Section 2.2.

2.1 Milestones and Eras of Computational Power

Computing power increased many orders of magnitude from hundreds of floating-point operations per second (FLOPS) available in the first computers to exaFLOPS (10^{18} FLOPS) in 2022 [27, 10]. Table 1 lists seven milestones or eras of computing in terms of FLOPS, each characterized by a three-order of magnitude increase in computing power. The table presents the first computer of a given era, the year the milestone was reached, and a modern-day equivalent that provides the same computing power where appropriate.

Miniaturization enabled by Moore’s Law and Dennard Scaling [8] allows even standard consumer electronics to out-compete supercomputers from only a few decades ago. To put this into perspective, a low-cost microcontroller available since 2003 used by millions of hobbyists (e.g., Arduino Uno) routinely outperforms the leading supercomputer of 1954, just 49 years after it became the first computer of the kiloFLOPS era [15]. A high-end smartphone CPU from 2015 (e.g., Qualcomm Snapdragon 617) achieves 228.4 MFLOPS in LINPACK for Android just 39 years after supercomputers entered the 100+ MFLOP era. A Raspberry Pi 2 from 2015 is as powerful as a Cray2 that unlocked the GFLOPS era in 1985 [28]. A modern gaming console (e.g., the PlayStation 4), of which over a hundred million units were sold, through their graphic card features enough processing power to rival a supercomputer from the 1+ TFLOPS era [24]. Finally, modern Graphical Processing Units (GPUs) are beginning to rival systems in the 100+ TFLOPS era for single-precision workloads and even over 1+ PF in Tensor Performance (e.g., GeForce RTX 4090 with up to 82 TFLOPS) [20].

	FLOPS Machine	Year	Speed	Same performance available in a ...
	ENIAC	1946	385 FLOPS	
kilo	IBM 704	1954	12 KFLOPS	2003 Low-Cost Microcontroller [15]
mega	IBM 360	1964	1 MFLOPS	
	CDC Cray-1	1976	160 MFLOPS	2015 Smartphone CPU [28]
giga	CDC Cray2	1985	1.9 GFLOPS	2015 Raspberry Pi 2 [28]
	Num. Wind Tunnel	1993	124.0 GFLOPS	2010 GPU (double-precision)
tera	ASCI Red	1997	1.1 TFLOPS	2013 Gaming Console (double-precision) [24]
	IBM Blue Gene/L	2006	136.8 TFLOPS	*2022 GPU (single-precision) [20]
peta	Roadrunner	2008	1.0 PFLOPS	*2022 GPU (tensor-performance) [20]
	Tienhe-2	2013	33.9 PFLOPS	-
	IBM Summit	2018	122.3 PFLOPS	-
	Fujitsu Fugaku	2020	442.0 PFLOPS	-
exa	ORNL Frontier	2022	1.102 EFLOPS	-

Table 1: History of computing systems passing significant milestones concerning the FLOPS achieved by supercomputers compared to modern devices integrating the same computational capabilities in a single device/package. The two modern systems marked with "*" denote that these devices surpass the TFLOPS and PFLOPS barriers for a more specialized operation only. This table is an adaption of [10] that was augmented with modern devices performance information from [15, 28, 24, 20].

As these devices become more powerful and affordable, computing finds its way into more and more aspects of our daily lives for professional and recreational purposes. We now take many once-demanding computational applications, such as route planning, panorama stitching, or speech recognition, for granted to be available on a smartphone. At the same time, applications such as search services or speech recognition are often still augmented by network requests to cloud services for more sophisticated functionality. Notably, billions of people use these technologies without noticing the complex systems making this possible because many insights from human-computer interaction (HCI) are employed to make them as easy to use as possible and for non-technical users. *A consequential question is how can we help researchers, educators, and students to take advantage of this relative wealth of processing power for their research?*

2.2 Geographic Distribution of Supercomputers

Learning to use advanced computing systems requires practice and exposure to actual systems. For decades this required in-person access to a computing facility operated by the university or research institution. Today, most users of cloud and HPC systems never set foot anywhere near the facilities they are running their calculations. While geographic access has become less important to users due to high-speed internet from the direct user perspective, they remain relevant for research both from an operational and educational perspective. For operations, many tasks and processes triggered by users, such as data transfers, are constrained by the underlying physical network topology and geographic distribution. Similarly, research into improving these operations through better




Fig. 1: Top500 Supercomputers plotted on the world map at the end of every year in which the fastest system surpassed one of the milestones denoting a new era: 100+ GFLOPS, 1+ TFLOPS, 100+ TFLOPS, 1+ PFLOPS, 100+ PFLOPS, 1+ EFLOPS. Figures: [17]

system architectures, scheduling, and data management is usually driven by exposure to an operational problem which appears to be also reflected in the authorship of publications related to such improvements.

For this reason, we want to understand the geographic distribution of supercomputing resources better. Figure 1 plots the geographic distribution of the Top500 supercomputers over time for every year in which the fastest system on the list surpassed a critical milestone entering a new era between 1993 and 2022:

100+ GFLOPS, 1+ TFLOPS, 100+ TFLOPS, 1+ PFLOPS, 100+ PFLOPS, 1+ EFLOPS. The Top500 does not include all operating HPC systems but only lists the 500 fastest systems for which benchmark results were submitted. Interestingly, supercomputer adoption is spreading only relatively slowly to new regions. In the 1990s, most supercomputers were located in the US, Europe, and Japan, but in 1993 there were systems also in Brazil and Australia. In the 2000s, China and India start deploying more supercomputers. From 2007 going forward, the Top500 began to record city and country for most sites, with only some opting to remain anonymous. At the end of 2018 and 2022, 45.4% and 32.4%, respectively, of all the supercomputers were deployed in China, many of the systems were not resolved to a city. In Africa, only South Africa and Morocco ever appear on the list.

3 Development and Distribution of Networking Infrastructure

One of the enabling technologies to make advanced computing techniques available to off-site collaborators or research institutions is communication networks. We analyze how networking infrastructure’s performance and geographic distribution evolved, how the increased connectivity provides opportunities for new supercomputing applications, and the democratization of access to platforms and data. Our analysis focuses on the following:

- An assessment of network topologies and capacities of different research and educational networks in the US (i.e., ESnet6, Internet2, and regional NRENs) and the state of high-speed internet coverage through broadband and mobile networks within the US.
- The comparison of the US-based network topologies and capacities to similar efforts in other regions (i.e., Africa, Asia, Europe, and South America).

Many networks report their link capacity in gigabits per second (Gbps) or short just G. We adopt the same convention and will use the shorthand unless it is ambiguous. Some networks do not explicitly report the link capacity but denote that two nodes are connected via ”fiber.” These links deploy so-called open-line fibers for which capacity is primarily constrained by the optical transmitter and receiver technologies, which can be continuously upgraded as better technologies or funding becomes available while not needing to change the fiber connecting two sites.

3.1 High-Speed Research and Education Networks in the US

ESNet The Energy Sciences Network (ESnet) is among the world’s fastest research networks. It was officially formed in 1986. Today ESnet6 connects most Department of Energy (DOE) facilities and other national and international research institutions (e.g., NASA, CERN). Figure 2b shows a topology including link capacities forming ESnet6. Most of the links in ESnet6 are equipped to provide 800G+ network capacity.




Fig. 2: Snapshot of topologies and coverage for different network and internet infrastructure in the US.

Internet2 Research and education communities, industry, and government members jointly lead the Internet2 non-profit networking consortium. First established in 1996, the project achieved 10 Gbit/s connectivity between over 230 member institutions in early 2004. Shortly after, the network was upgraded to 100 Gbit/s and today largely features a fiber network as illustrated in Figure 2a showing the internet’s topology, including the underlying link technologies.

Regional NRENs in the US At the regional level, many states and universities in the US are invested in their own educational and research networks (NREN). Today most of these networks are part of the coalition for Advanced Regional Networking in Support of Research and Education (Quilt) [1]. Quilt currently lists 43 networks that connect over 900 universities as well as tens of thousands of other educational and community anchor institutions. Figure 2c shows the mesh of connections that span the US through regional NRENs. Unfortunately, the maps do only cover connectivity but not link capacities. Many of these NRENs are connected to Internet2, allowing them to accelerate long-distance transfers vastly.

Broadband and Mobile Networks The widespread availability of high-speed broadband and mobile internet is also essential for field applications that integrate advanced computing. Figure 2d and Figure 2e show the coverage for fixed broadband internet and mobile internet, respectively, as collected by the Federal Communications Commission (FCC) [11]. The coverage maps show that high-speed internet is widely available across the US via fixed broadband and cellular networks. This allows considering access to advanced computing methodologies not only for users with fixed internet connection but also enables many mobile applications. Examples of mobile applications that use supercomputing resources include integrating wildfire simulations into team-awareness tools available to firefighters [14].

3.2 GÉANT and Other International and Regional NRENs

Investments in networking infrastructure to connect advanced computing infrastructure and institutions are not limited to the US. Figure 3 shows the network topologies of similar efforts in other regions such as Europe, South-East Asia, South America, and Africa. The GEANT project, in succession to various earlier European efforts (EARN, EUNET, EuropaNET, TEN-34, TEN-155) to connect research institutions dating back to 1983, began in 2000 to build the Gigabit European Academic Network. In its current iteration, GEANT connects 44 national NRENs in Europe [12]. GEANT is connected to ESnet and Internet2, as well as many other regional NREN coalitions. Beginning in 2003, the RedCLARA network connects 15 countries in Latin America with direct connections to Internet2 [23]. Established in 2004, The Trans-Eurasia Information Network (TEIN) connects various institutions across 21 countries in Asia and is also connected to GEANT and Internet2 [26]. In Africa, multiple NREN coalitions coexist: The




Fig. 3: Network topology of different NRENs or coalitions of NRENs in Africa, Asia, Europe, and South America. Figures are modifications of the originals to fit space and enhance contrast [12, 26, 23, 7].

Arab States Research and Education Network (ASREN) connects 22 countries in Asia and North Africa. The West and Central African Research and Education Network (WACREN) connects 14 West and Central African countries. The UbuntuNet Alliance for Research and Education Networking (UbuntuNet) connects 17 countries in Eastern and Southern Africa. Each of the three networks is also connected to GÉANT.

4 Related Work

Development and the assessment of access to computational resources and connectivity of research organizations have been subjected to extensive research and surveying activities[1, 27, 9] and annual assessments of, for example, funding

agencies [2–5, 13]. In the context of HPC, the Top500 list³ started tracking the 500 fastest supercomputers since 1993 on a bi-yearly basis [27]. We take a closer look at the data provided by the Top500 but add the geographic perspective in Section 2.2 and analyze the development over time. For advanced communication and network infrastructure, most network providers and alliances provide current topology maps, including capacity information [6, 12, 16, 1], but not in a standardized or machine-friendly format. For regional NRENs in the US, the Quilt⁴ [1] is a coalition representing 43 educational networks across the United States. This effort also includes publishing up-to-date maps to visualize where research and education networks are deployed see Figure 2c. While these efforts focus on specific categories of cyberinfrastructure, our work performs a combined analysis. It relates it to new opportunities in democratizing access to platforms and data and the role of human-computer interaction therein.

5 Lessons Learned

We analyzed the development of two important cyberinfrastructure (i.e., computing facilities and network infrastructure). We analyzed the computing capabilities of supercomputers over 7 decades and the geographic distribution of the Top500 over the last 3 decades. We also related the computing capability of past supercomputers to modern devices and consumer electronics ubiquitous today. The analysis shows that while we may have a relative wealth of computing power in comparison to previous decades, direct geographical access to supercomputing only slowly evolves and most new supercomputers replace already existing ones.

In a second analysis, we considered the network infrastructure connecting researchers worldwide. We observe that a hierarchy of networks connects research and education networks globally. While the US and the EU networks are accommodating more participants at high link speeds, research networks in Asia, South America, and Africa rapidly connected countries and institutes in the last 20 years. In the US the largest network capacity for research transfers is found in ESnet at 800+G which connects the DOE national labs and other large research facilities. At the regional level universities and states invested in local networking infrastructure. This mesh of regional research and education networks already ensures connectivity across institutions and universities but at potentially high latency and low bandwidth. Here cross-regional efforts such as Internet2 connect even geographically distant NRENs with low latency while reducing contention in local networks. Finally, we also considered the availability of broadband and mobile internet in the US. The analysis using FCC data suggests that internet coverage is available except for remote areas such as the mountains.

This analysis suggests that many of the technologies to democratize access to platforms and data are in place. Yet, our experience working with institutions, universities, domain scientists, and students tells us that adopting these technologies remains challenging for many. Continued investments into expanding

³ <https://www.top500.org/>

⁴ <https://www.thequilt.net/>

the physical infrastructure remain important, but if we want to empower more researchers to leverage advanced computing methodologies, it is absolutely necessary to also invest in the software stack so research can turn to a standard API to define their data dependencies allowing researchers to reclaim time spent on research while also giving cyberinfrastructure providers mechanisms to optimize.

Acknowledgment

This research was supported by the National Science Foundation (NSF) under grant numbers #1841758, #2028923, #2103845, and #2138811; the Advanced Cyberinfrastructure Coordination Ecosystem: Services and Support (ACCESS) program, under allocation TG-CIS210128; Chameleon Cloud under allocation CHI-210923; and IBM through a Shared University Research Award.

References

1. The Quilt - Advanced Regional Networking In Support Of Research And Education. <https://www.thequilt.net/>
2. Scientific and Engineering Research Facilities: 1999. <https://wayback.archive-it.org/5902/20150628160048/http://www.nsf.gov/statistics/nsf04334/pdfstart.htm>
3. Scientific and Engineering Research Facilities: 2001. <https://wayback.archive-it.org/5902/20150629121928/http://www.nsf.gov/statistics/nsf02307/sectb.htm>
4. Scientific and Engineering Research Facilities at Colleges and Universities: 1998. <https://wayback.archive-it.org/5902/20150627201815/http://www.nsf.gov/statistics/nsf01301/>
5. Scientific and Engineering Research Facilities at Colleges and Universities: 1998 - Appendix E. <https://wayback.archive-it.org/5902/20150629135427/http://www.nsf.gov/statistics/nsf01301/appe.htm>
6. ESnet6 Maps. <https://www.es.net/welcome-esnet6/esnet6-maps/> (2022)
7. Banda, T.: Research and Education Networks in Africa (Aug 2020)
8. Bohr, M.: A 30 Year Retrospective on Dennard's MOSFET Scaling Paper. IEEE Solid-State Circuits Society Newsletter **12**(1), 11–13 (2007). <https://doi.org/10.1109/N-SSC.2007.4785534>
9. Chalker, A., Hillegas, C.W., Sill, A., Broude Geva, S., Stewart, C.A.: Cloud and on-premises data center usage, expenditures, and approaches to return on investment: A survey of academic research computing organizations. In: Practice and Experience in Advanced Research Computing. pp. 26–33. ACM, Portland OR USA (Jul 2020). <https://doi.org/10.1145/3311790.3396642>
10. Chen, J., Ghafoor, S., Impagliazzo, J.: Producing competent HPC graduates. Communications of the ACM **65**(12), 56–65 (Dec 2022). <https://doi.org/10.1145/3538878>
11. FCC: FCC National Broadband Map. <https://broadbandmap.fcc.gov/home> (2023)
12. GEANT: GÉANT Connectivity Map. <https://map.geant.org/> (2023)
13. Gibbons, M.: Computing and Networking Capacity Increases at Academic Research Institutions (2013)

14. Holland, T.M.: ATAK Improves Situational Awareness for California Fire Department. <https://insights.samsung.com/2019/10/16/atak-improves-situational-awareness-for-california-fire-department/> (Oct 2019)
15. Ian: Answer to "How computationally powerful is an Arduino Uno board?". <https://robotics.stackexchange.com/a/538> (Nov 2012)
16. Internet2: Operations and Support. <https://internet2.edu/network/operations-and-support/> (2023)
17. Luettgau, J.: Maps of the Top500 Supercomputers over Time (Nov 2022). <https://doi.org/10.5281/zenodo.7606369>
18. Luettgau, J., Kirkpatrick, C.R., Scorzelli, G., Pascucci, V., Tarcea, G., Taufer, M.: NSDF-Catalog: Lightweight Indexing Service for Democratizing Data Delivering. In: IEEE ACM International Conference on Utility and Cloud Computing (UCC2022) (2022)
19. Luettgau, J., Olaya, P., Zhou, N., Scorzelli, G., Pascucci, V., Taufer, M.: NSDF-Cloud: Enabling Ad-Hoc Compute Clusters Across Academic and Commercial Clouds. In: Proceedings of the 31st International Symposium on High-Performance Parallel and Distributed Computing. pp. 279–280. ACM, Minneapolis MN USA (Jun 2022). <https://doi.org/10.1145/3502181.3533710>
20. NVIDIA: ADA GPU Architecture V1.01. <https://images.nvidia.com/aem-dam/Solutions/geforce/ada/nvidia-ada-gpu-architecture.pdf> (2022)
21. Olaya, P., Kennedy, D., Llamas, R., Valera, L., Vargas, R., Lofstead, J., Taufer, M.: Building Trust in Earth Science Findings through Data Traceability and Results Explainability. IEEE Transactions on Parallel and Distributed Systems **34**(2), 704–717 (Feb 2023). <https://doi.org/10.1109/TPDS.2022.3220539>
22. Olaya, P., Luettgau, J., Zhou, N., Lofstead, J., Scorzelli, G., Pascucci, V., Taufer, M.: NSDF-FUSE: A Testbed for Studying Object Storage via FUSE File Systems. In: Proceedings of the 31st International Symposium on High-Performance Parallel and Distributed Computing. pp. 277–278. ACM, Minneapolis MN USA (Jun 2022). <https://doi.org/10.1145/3502181.3533709>
23. RedCLARA: Network Maps. <https://www.redclara.net/index.php/en/recursos/publicaciones-para-difusion/mapas-de-la-red> (Mar 2020)
24. Sony Entertainment: Announcement of the Playstation 4. https://web.archive.org/web/20130424075309/http://scei.co.jp/corporate/release/130221a_e.html (Apr 2013)
25. Tarcea, G., Puchala, B., Berman, T., Scorzelli, G., Pascucci, V., Taufer, M., Allison, J.: The Materials Commons Data Repository. In: 2022 IEEE 18th International Conference on E-Science (e-Science). pp. 405–406. IEEE, Salt Lake City, UT, USA (Oct 2022). <https://doi.org/10.1109/eScience55777.2022.00060>
26. TEIN: Network Maps. <https://www.tein.asia/sub/?mc=2030> (2020)
27. Top500: Top500 Supercomputing Sites. <http://www.top500.org/> (2019)
28. Vince Weaver: The GFLOPS/W of the various machines in the VMW Research Group. https://web.eece.maine.edu/~vweaver/group/green_machines.html (2023)