Check for updates

Holistic Methodologies for Fusing the Material and the Ideological: Tracing the Sociotechnical Imaginaries that Enable the World's Fastest Supercomputer

Sarah Read
Portland State University
USA
read3@pdx.edu

Jordan Frith
Clemson University
USA
jfrith@clemson.edu

Abstract

When the Aurora supercomputer launches at Argonne National Laboratory it will operate at the exascale and be one of the fastest supercomputers in the world. We have been invited by Argonne to do field work to understand both the day-to-day processes that maintain supercomputing (e.g. user documentation, government reporting, user testing) and the sociotechnical national imaginaries [11], –in particular the increasingly tense global competition between the US and China in the supercomputing race–that shaped the development of Aurora. This article proposes our work-in-progress methodological approach that accounts for both higher-level imaginaries and the everyday practices of supercomputing. We also describe how our project plan is designed to answer multiple research questions: 1. What kind of new science will be possible when Aurora comes fully online? 2. What political, economic, technological, human and ideological resources has it required to build the supercomputer, bring it online, and to keep it running? 3. What does answering these questions tell us about the state of science in the US, the geopolitical race for computing power and the discursive and material infrastructures [5, 19] required for success?

CCS Concepts

• **Applied Computing-Document Management and Text Processing**; • **Applied Computing-Physical Sciences and Engineering**;

Keywords

Sociotechnical Imaginaries, Exascale Computing, User Documentation, Discursive Infrastructure

ACM Reference Format:
Sarah Read and Jordan Frith. 2024. Holistic Methodologies for Fusing the Material and the Ideological: Tracing the Sociotechnical Imaginaries that Enable the World's Fastest Supercomputer. In *The 42nd ACM International Conference on Design of Communication (SIGDOC '24), October 20–22, 2024, Fairfax, VA, USA*. ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/3641237.3691653

SIGDOC '24, October 20–22, 2024, Fairfax, VA, USA

ACM ISBN 979-8-4007-0519-9/24/10
https://doi.org/10.1145/3641237.3691653

1 Introduction

1.1 Aurora Finally Comes Online

Argonne National Laboratory will soon launch one of the world's fastest public supercomputers at the Argonne Leadership Computing Facility (ALCF). When it enters into production mode, Aurora [Fig. 1] will become the world's second public exascale-level supercomputer. According to preliminary interviews conducted with ALCF staff for this project, the goal is to have "big science" projects ready to run on Aurora the day it launches. This is a bold mission and big news. To quote the ALCF, Aurora's "high computing speed and artificial intelligence capabilities will enable science that is impossible today" [30]. As the DOE "Exascale fact sheet" argues, Aurora "will allow scientists to create more realistic Earth system and climate models. . .understand the nanoscience behind new materials. . .build fusion power plants. . .[and] power new studies of the universe" [31].

Supercomputing, and the supercomputer Aurora in particular, is not very visible or accessible to the general public or Technical Communication scholars. To help make the impacts and contributions of supercomputing more available and useful, this multi-year research project has been designed around three key lines of inquiry: How Aurora came to exist as a nationally and globally important supercomputer, the role of technical documentation in its building and maintenance, and the impacts that its arrival will have on science. As scholars who have previously theorized infrastructure as an analytical lens for the study of writing, communication design, and technology [5, 8, 19–21], we adopt infrastructure as the primary lens for our methodology, as explained in this article.

1.2 US National Laboratories and the Global Supercomputing Race

The story of all big science projects is one of uncertainty and contingency due to their complexity along multiple axis, including the political, economic and technological [20]. The arrival of Aurora has never been a given, but is integral to the role and reputation of the US as a global scientific leader. The Aurora supercomputer that will launch later this year bears little resemblance to the Aurora supercomputer first publicized in a 2015 DOE press release that announced a $200 million investment in the project. As the Undersecretary of Energy at the time proclaimed, "Argonne National Laboratory's announcement of the Aurora supercomputer will advance low-carbon energy technologies and our fundamental understanding of the universe" [27]. However, in 2016 China launched the Sunway Taihulight supercomputer, which took over

Figure 1: Aurora supercomputer from public viewing room. Photograph by Read.

the top spot on the TOP500 and was nearly three times faster than the top-ranked supercomputer it supplanted. The power of the Sunway Taihulight surprised experts and challenged the position of the United States as the global scientific leader. Media coverage about China's surprising new supercomputer repeatedly evoked metaphors of Sputnik to warn that "China's new supercomputer puts the U.S. further behind" [4]. Then in 2017, Chinese state media announced plans to build the world's first exascale supercomputer, further raising alarm in the U.S. government. Congressional testimony at a hearing held by the Armed Forces Subcommittee on Emerging Threats repeatedly referenced the Sunway Taihulight and China's exascale plans as a major threat to future U.S. dominance [29].

The original plans for Aurora were abandoned in 2017 after the success of the Sunway Taihulight. Instead, the DOE invested an additional $500 million in the project and announced that Aurora would be developed to be the world's first exascale computer rather than a stepping-stone. Aurora then became a microcosm of growing geopolitical tensions, and as a New York Times article put it, Aurora had "become crucial in a high-stakes technology race between the United States and China" [3].

1.3 Why A Methodology for Fusing Discursive and Material Infrastructures

Aurora's success has a great deal riding on it as both a scientific cyberinfrastructure and a part of the U.S. strategic security and scientific agenda. Consequently, this article focuses on the complex nature of supercomputers cyberinfrastructure by introducing a project we have planned to study the launch of Aurora and how supercomputing infrastructure functions in practice. We have been invited to do field work at Argonne National Laboratory by the Aurora project leads. Our project will implement a work-in-progress methodological approach that combines field work that examines the mundane practices of supercomputing (e.g., documentation, human labor) with a higher-level approach that analyzes the sociotechnical imaginaries, a concept we cover in the next section. As we argue, the lower-level everyday practices and grander imaginaries are inextricably linked and can both be accounted for through a fused methodological approach that crosses levels and domains.

Our methodology is designed to pursue the research questions that we bring to this project:

- RQ1: What role does user documentation, as a discursive infrastructure *for* the infrastructure of supercomputing, play in enabling, constraining, and mediating the mission of U.S. National Laboratory supercomputing?
- RQ2: How do technical reporting and documentation processes required by the DOE enable and constrain the mission of the ALCF, and what do they reveal about the 70-year history of linkages between that National Laboratory system and next-generation supercomputing?
- RQ3: What does the national investment in huge scientific projects like Aurora reveal about the increasing tension between the United States and China for global sociotechnical dominance?
- RQ4: How can exascale supercomputing potentially enable novel forms of scientific knowing not possible at the petascale level?

2 Tracing Material and Discursive Infrastructures

Our research methodology is being built from several theoretical concepts and analytical tools from Science and Technology Studies and Technical Communication.

2.1 Opening Infrastructural Black Boxes

Examining the materiality of infrastructure and opening up infrastructural "black boxes" has long been a key STS and Technical Communication research agenda [1, 13–15, 26]. Both literally and metaphorically, a black box is a simplified interface that obscures or disguises the complexity on the inside. Standing in the machine room next to Aurora, it looks to a lay visitor like a very big computer: a collection of large metal cases, copious cables, and other electronics. However, a more informed visitor would know that Aurora is a highly complex black box, a combination of 10,264 computing nodes, two Intel processors developed specifically for the project, a novel HPE Cray software stack, over 230 petabytes of high-performance storage, multiple APIs compatible with its program model, and tens of thousands of Xeon Max CPUs. Many of these elements were built specifically for Aurora through hundreds of millions of dollars of federal funding and a long, uncertain funding and development process. As infrastructural researchers, one of our goals for this project is to explore the complex black box of Aurora as material and discursive infrastructures. The following sections introduce the analytical constructs we will use to open the black box by tracing Aurora's material and discursive infrastructures and their relationship to the sociotechnical imaginaries that shape US science, technology and security policy and investment.

2.2 Sociotechnical Imaginaries

The concept of "imaginaries" has a long history in STS research [16, 2.], but the specific concept of sociotechnical imaginaries was first developed in Jasanoff and Kim's [11] influential article about different understandings of nuclear energy in the United States and South Korea. Jasanoff and Kim's original conceptualization of sociotechnical imaginaries focused almost exclusively on the state's power to guide future development: "Sociotechnical imaginaries as we define them are associated with active exercises of state power, such as the selection of development priorities, the allocation of funds, the investment in material infrastructures" (p. 123) [11]. The authors make clear that these imaginaries are not the same as explicit policy agendas; "They are less explicit, less issue-specific" (p. 123) [11], and are the "collectively imagined forms of social life and social order reflected in the design and fulfillment of nation-specific scientific and/or technological projects" (p. 120) [11]. Despite the more amorphous nature of sociotechnical imaginaries compared to, for example, a specific policy, they are highly consequential because "National imaginations can penetrate the very designs and practices of scientific research and technological development" (p. 124) [11]. The way communities imagine the future potential of scientific and technological development guides resource allocation, shapes public rhetoric, and can play a formative role in making those futures possible. As Mager and Katzenbach put it, "future visions are performative" (p. 223) [18].

Jasanoff and Kim's original definition has been widely cited in STS research but has also been critiqued for focusing too much on the power of the state [9, 18]. Jasanoff altered the definition to account for the role non-state actors play in shaping sociotechnical imaginaries, arguing that the concept needed to be "refined and extended in order to do justice to the myriad ways in which scientific and technological visions enter into the assemblages of materiality, meaning, and morality that constitute robust forms of social life" (p. 4) [10]. She then redefined sociotechnical imaginaries as the "collectively held, institutionally stabilized, and publicly performed visions of desirable futures, animated by shared understandings of forms of social life and social order attainable through, and supportive of, advances in science and technology" (p. 4) [10]. More recent work has gone further in its critique of the original concept, arguing that tech corporations now play a stronger role in driving sociotechnical imaginaries than the state. An entire special issue of the journal *New Media & Society*, for example, was devoted to detailing the powerful role corporate actors play in shaping sociotechnical imaginaries.

The National Laboratory system has long played an important role in the sociotechnical imaginary of the United States [17]. The highest level of this project aims to trace the infrastructures that enabled and operationalized Aurora's shift from a stepping-stone toward exascale computing with plans to launch in 2018 to an exascale supercomputer launching in 2024. The project changed in part because of the influence of the sociotechnical imaginary of the United States as a dominant player in the global race for technological and scientific supremacy.

2.3 Relationality Of Scientific Infrastructure

As discussed above, sociotechnical imaginaries examine how stabilized institutional visions guide scientific and technological development, and we will examine Aurora's role in the U.S. government's imaginary through archival work on supercomputing discourses and interviews with ALCF leadership. However, this project will also leverage our access to Aurora to examine how supercomputers are built and maintained and how they overcome the headwinds of the risks and uncertainties that must be mitigated. We will do so through a framework that retains the higher-level framework of imaginaries but also incorporates the prominent concept of relational infrastructure. Infrastructures are not an ontological state of being. Instead, material objects and other practices become infrastructural when they do infrastructural work that shapes higher-level developments. As Star and Ruhleder noted, we need to ask "when—not what—is an infrastructure" (p. 113) [25]. We argue that a relational approach that examines the everyday practices of supercomputing—which include human and nonhuman actors—is not only compatible with the higher-level focus on geopolitics and imaginaries, but crucial to deepening understandings of how sociotechnical imaginaries of cyberinfrastructure are materially and discursively constituted and shaped. As we detail, for a national "grand strategy" to become concrete reality requires an immense amount of human labor to document and maintain large cyberinfrastructures.

Aurora itself certainly does infrastructural work. Scientific projects that run on the supercomputer will ideally produce new forms of scientific knowledge, even as the supercomputer itself remains invisible to the vast majority of people. In that sense, Aurora is what Sandvig [23], drawing from Kling's [12] work, calls a "hard" infrastructure, a material instantiation of thousands of nodes that take up a huge room in the Chicago suburbs and cost well over $500 million. The materiality of Aurora alone, however, will not determine the success of the project or the value of the investment. Returning to the concept of relationality, while Aurora does significant infrastructural work, it is also a primary object of focus for the people who maintain it and who run projects on it. In other words, Aurora does the opposite of sinking "into an invisible background" (p. 115) [25] for the ALCF staff who manage Aurora, the scientists who use Aurora, and the people responsible for documenting Aurora based on DOE requirements. By adopting a relational view that examines the infrastructural work Aurora will do as a key piece of the national sociotechnical imaginary and analyzes the sociotechnical practices during which Aurora is a primary focus and not infrastructural, this project will open the black box of supercomputing through interviews, document analyses, and field work at ALCF.

Another key part of how this project will use relationality is inspired by Star's argument that "it's infrastructure all the way down" (p. 380) [24]. Or to put it differently, infrastructures have infrastructures. The massive material assemblage that is the Aurora supercomputer is shaped by "soft" infrastructural practices that make supercomputing possible.

2.4 Discursive Infrastructures

This project will examine how discursive infrastructures—a concept the two project leads have developed through multiple publications [5, 6, 19, 20, 22]—shape the sociotechnical practices of Aurora. Discursive infrastructures include the many documents and rhetorical strategies almost no one ever sees that play an agential role in shaping higher-level infrastructural practices. Our concept of discursive infrastructures recognizes that documents do have materiality, but

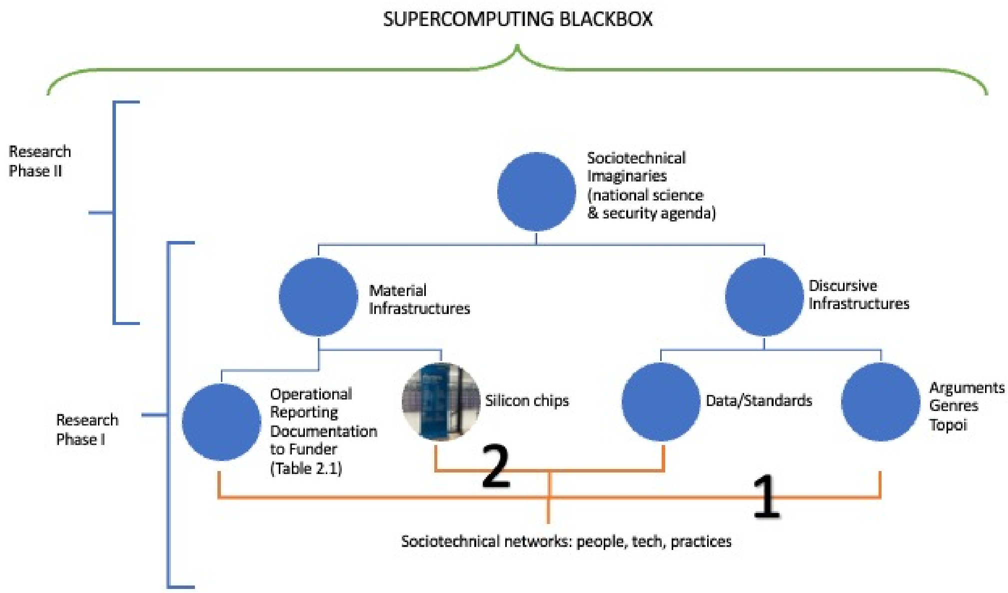

Figure 2: Conceptual Diagram of relationships among elements (blue dots) of the Supercomputing Blackbox. Methodology Illustration 1 and 2 are marked with numbers.

the focus on the discursive is intended to examine the rhetorical infrastructural work documentation does to support higher-level practices. Various documents, which for this project include user documentation designed for scientists and reporting procedures to the DOE, do infrastructural work that "stands up," to use an industry term, a supercomputer.

2.4.1 Infrastructural Role of Documentation. Beneath the large material structure of Aurora lies countless documents that do rhetorical work that makes the material assemblage of nodes meaningful. In other words, for Aurora to successfully further big science and enact national imaginaries, a variety of mundane documentation practices are necessary that do infrastructural work and remain mostly invisible. In fact, the United States has been uniquely successful in building and maintaining big science projects because of the development of a national documentation infrastructure that mitigates the political, economic, and rhetorical uncertainties that can derail big, long-term projects [20]. This project will examine the multiple layers of infrastructure involved in huge supercomputing projects through a relational lens that analyzes the practices necessary to make a $500+ million supercomputer meaningful.

3 Fusing Discursive and Material Infrastructures

Holistically tracing material and discursive infrastructures and how they shape the sociotechnical imaginaries that helped spur massive federal investment in the Aurora project requires a research methodology that can link higher-level sociotechnical entanglements to the concrete things and practices—the wires, computing racks, documents, staff members—of supercomputing.

3.1 Two-Phase Research Plan

Our project will proceed in two overlapping phases, as illustrated in Figure 2. Data collection and analysis during each phase will focus on different elements (blue circles) of the Aurora black box. Elements include sociotechnical imaginaries and material and discursive infrastructures. Linking the material and discursive infrastructures are the sociotechnical networks of people, technologies and practices (orange lines; Illustrations 1 and 2, below).

3.1.1 Phase I: The Infrastructural Function of Exascale Supercomputing User Documentation. The first phase, which we call "The infrastructural function of user documentation for exascale supercomputing," will examine the sociotechnical practices involved in supercomputing. This phase will draw from a relational approach to infrastructure and collects data from users, staff, and documents for whom Aurora is not an embedded, pseudo-invisible infrastructural object. Instead, it is their primary area of focus, and we will build upon research on what we call *discursive infrastructure* to examine the layers of documents and social practices that do the lower-level infrastructural work that makes the higher-level infrastructural work of Aurora possible.

3.1.2 Phase II: Exascale Computing as Infrastructure for U.S. Science: Stable and Emerging Sociotechnical Imaginaries and Beyond. The second phase—"Exascale computing as infrastructure for U.S. science: Stable and emerging sociotechnical imaginaries and beyond"—will move from analysis of the everyday practices of supercomputing to an examination of what our data reveals about the sociotechnical imaginaries of supercomputing. We will explore the multiple phases of Aurora's development and how ALCF staff collect and interpret daily, monthly, and annual operational metrics to examine how Aurora shapes the scientific agenda of the United States. The second phase will also demystify the complex assemblages of supercomputing for a broader audience by opening the black box of exascale supercomputing. Through our interviews and analysis of documents we will develop a deeper understanding of just why exascale computing matters as scientific infrastructure and how exascale operations can produce new ways of knowing, which is a key part of why supercomputing has long been a source of global scientific competition.

3.2 Methodology Illustration 1: From Scientific "Grand Challenges" to Operational Assessment Report

This brief illustration of fusing material and discursive infrastructures traces how a humble table of data (Figure 3) in the latter pages of a glossy professional technical report is the outcome of and operationalizes the US imaginary and strategy as a world leader in addressing the "grand challenges" of scientific discovery and technological innovation. Every year the Argonne Leadership Computing Facility (ALCF) must author and submit a report to the Department of Energy (DOE) that formally reports on whether the operations of the supercomputer have met the operational metrics required to renew the facility's funding. Gathering the data from the supercomputer operations, analyzing the data, narrating the data and designing the report is a year-long process that touches every team in the facility—it consumes a lot of staff time, but is also integral to the work of running the supercomputer, most essentially because its continued funding depends upon meeting and reporting on the metrics set by the DOE. Tracing how the generation of this annual report, and the key metrics reporting table in particular (Fig. 3), is the both the outcome of and the operationalization of strategic sociotechnical imaginaries for US science and technology as a leader in addressing scientific "grand challenges" makes visible the value that reporting processes and documentation have for the facility and the National Laboratory system as a whole. Technical documentation is often invisible and undervalued in engineering and industry, often considered low priority work that is just a write-up of technical operations. For the supercomputer, however, and the US's position as a leader in scientific research, everything does actually depend on the table in Figure 3. Operating the supercomputer and authoring the report at not separate activities, temporally or rhetorically.
Material Infrastructures (from systemic to localized):

- US National Laboratory system
- Leadership Computing Facilities (Oakridge and Argonne)
- Public supercomputers available for scientific and industrial research (Mira or Aurora)
- Technical documentation
- Document cycling processes for authoring technical documentation
- Annual Operational Assessment Review (OAR) report,
- Table (2.1) from Operational Assessment Review report (Fig. 3).

Discursive Infrastructures (from systemic to localized):

	(Blue Gene/Q): 48K-node, 768K-core, 1.6 GHz, 768 TB RAM			
	2013 (Apr 9–Dec 31)		CY 2014	
	Target	Actual	Target	Actual
Scheduled Availability	85.0%	95.5%	95.0%	98.7%
Overall Availability	80.0%	90.6%	90.0%	95.7%
System MTTI	N/A[a]	4.23 days	N/A	8.98 days
System MTTF	N/A	11.29 days	N/A	25.80 days
INCITE Usage	2.1B	2.4B	3.5B	3.9B
Total Usage	N/A	3.6B	N/A	5.8B
System Utilization	N/A	79.4%	N/A	87.6%
Overall Capability[b]	20.0%	60.7%	30.0%	64.5%[d]
High Capability[c]	5.0%	33.3%	10.0%	33.1%

Figure 3: Table (2.1) from the Operational Assessment Review report displaying a summary of the metrics data used to evaluate the performance of the supercomputer for a single funding cycle.

- The rhetorical topoi of efficiency and value for federal taxpayers for funding federal facilities and programs
- Technical documentation and reporting genres
- Genre of the Department of Energy Operational Assessment Report
- Conventions for reporting data in tables in professional reports
- Machine data analysis and business intelligence decision-making practices
- Operational metrics expressed in Table (2.1) (Fig. 3) in terms of machine availability (how often the supercomputer was running) and utilization (how much of the available computing power was used)

Sociotechnical Imaginary: The US is a world leader in addressing the "grand challenges" of scientific discovery and technological innovation. Textual instantiations of the imaginary:

- Text from "About" page of Department of Energy:
- "The mission of the Energy Department is to ensure America's security and prosperity by addressing its energy, environmental and nuclear challenges through transformative science and technology solutions." [https://www.energy.gov/about-us]
- Description of the grant program scientists apply to gain access to the supercomputer:
  - "Open to researchers from academia, government laboratories, and industry, the INCITE program aims to accelerate scientific discoveries and technological innovations by awarding, on a competitive basis, time on supercomputers to researchers with large-scale, computationally intensive projects that address "grand challenges" in science and engineering." [https://www.alcf.anl.gov/science/incite-allocation-program]

3.3 Methodology Illustration 2: From Global Sociotechnical Imaginaries to Material Infrastructures

Sociotechnical imaginaries shape material infrastructures in various ways, most commonly through standard setting. As an example

Figure 4: A label on a can of artichokes that shows the product's barcode. Photo by Frith.

from one of the author's research [7], our global supply chains, retail networks, shipping networks and so on are shaped by the complex standard and international imaginaries that control identification processes through barcode technology. Every number in a barcode (including the barcodes found in grocery stores, the 2D barcodes on packages, and so on) is determined through a standards organization called GS1. GS1 exists in large part because of the sociotechnical imaginary of global cooperation and the free flow of global goods and capital. The image below of a can of artichokes (Figure 4) shows a basic instantiation of that phenomenon.

The barcode on that can of artichokes is highly standardized; each digit describes something about the product, and those digits are managed by GS1. Importantly, the massive investment in creating and maintaining a global body that standardizes identification data is the result of a sociotechnical imaginary of global cooperation designed to ensure the seamless flow of products throughout the global supply chain. Obviously, a sociotechnical imaginary of global collaboration is different from the sociotechnical imaginaries of supercomputing described above; however, we use this example to show how we will adopt this approach in our work to show how imaginaries shape the materiality of Aurora. In the barcode example, the imaginary of a seamless global flow of goods has led to extensive material infrastructural development, including

- The standardization of tens of billions of barcodes with line patterns controlled by GS1
- The distribution of billions of laser scanners to access standardized barcode data
- A huge materials GS1 headquarters built in Brussels
- Over 110 other GS1 buildings in 110 other countries that monitor and control identification standards in different regions
- Thousands of human employees all over the world doing invisible infrastructural work to maintain the databases that enable the global flow of goods and ensure identification numbers are strictly controlled.

3.4 What These Illustrations Show

The report table and barcode examples illustrate our methodological approach that will link sociotechnical imaginaries to the materiality of Aurora as a supercomputer and an ecology of documentation and other discursive infrastructures.

Everything from the silicon chips to the rooms full of nodes to the reams of technical reporting documentation to the Argonne employees exist in that specific place at that specific time because of sociotechnical imaginaries driving the supercomputing race, just as a random barcode at a supermarket only works because of a massive bureaucratic infrastructure driven by an imaginary of cooperation. Consequently, our use of imaginaries combined with our field work will enable our project to link the everyday practices that enable Aurora and the materiality of the project to the higher level national imaginary that led to the development of Aurora in the first place.

And as a final point, we approach these different levels of our research as co-constitutive, not causal. Without the table (2.1) (Fig. 3) in the Operational Assessment Report, continued funding of the supercomputer would be at risk. Returning to barcodes, the imaginary that pushed GS1's global standardization does shape the material practices of identification; however, those material practices then maintain the imaginary because, to put it simply, they work. If the materiality of barcodes and the discursive infrastructures of standardization weren't successful, the sociotechnical imaginary would be significantly impacted (and maybe abandoned). Similarly, while national supercomputing imaginaries drove the specific constellation of discursive and material infrastructures that are Aurora, the supercomputer must succeed to further strengthen that imaginary. Consequently, our project will move back and forth between higher level imaginaries and everyday practice while recognizing both levels as intertwined rather than directly causal. Our work will build upon previous studies that have worked to articulate the materiality of discourse and how it shapes science and technology [16, 28, 29].

4 Conclusion

The research questions and methodology proposed in this article are ambitious and assume a scale of research accessible only to a funded, multi-year, collaborative research endeavor. The opportunities of this kind of project also present limitations: due to the complexity and dynamic nature of bringing a supercomputer fully online, we won't have full knowledge of the people, data, documents and technologies that we will have access to until we begin research. While we have general confidence in our access to the major narratives at the facility, we will need to adapt as our study unfolds to opportunities and limitations that come up. As a result, this paper is a preliminary sketch of our research methodology, which we fully expect to dynamically evolve as the study progresses.

At the point of submitting the final draft of this article have just received confirmation that our Science and Technology Studies National Science Foundation (NSF) grant is funded. Beyond the work ahead of us on this study, the methodology sketched in this article provides an initial blueprint to researchers asking questions that aim to trace how even the most banal daily practices and material and discursive tools of workers are co-constitutive of national agendas, geopolitical tensions and the outcomes of big science. As our research unfolds, we will refine the methodology. We also invite researchers to adopt and refine aspects of this methodology that inform their own research.

Acknowledgments

This material is based upon work supported by the National Science Foundation NSF 22-629 Science and Technology Studies under Grant No. SES-2419815.

References

[1] Nikhil Anand, Akhil Gupta, and Hannah Appel (eds.). 2018. *The promise of infrastructure*. Duke University Press.
[2.] Benedict Anderson. 1991. *Imagined communities: reflections on the origin and spread of nationalism*. New York, London.
[3] Don Clark. 2019. Racing against China, U.S. reveals details of $500 million supercomputer. *The New York Times*. Retrieved July 24, 2023 from https://www.nytimes.com/2019/03/18/technology/china-us-500-million-supercomputer.html
[4] Jack Dungara. 2016. China's New Supercomputer Puts the US Even Further Behind | WIRED. *Wired*. Retrieved July 24, 2023 from https://www.wired.com/2016/06/fastest-supercomputer-sunway-taihulight/
[5] Jordan Frith. 2020. Technical standards and a theory of writing as infrastructure. *Written Communication* 37, 3: 401–427. https://doi.org/10.1177/0741088320916553
[6] Jordan Frith. 2020. The Pedagogical Opportunities of Technical Standards: Learning from the Electronic Product Code. *Technical Communication* 67, 2: 42–53.
[7] Jordan Frith. 2023. *Barcode*. Bloomsbury Academic.
[8] Jordan Frith and Sarah Read. 2022. Introduction: Communication and design infrastructures. *Communication Design Quarterly* 10, 2: 4–10.
[9] Joachim Haupt. 2021. Facebook futures: Mark Zuckerberg's discursive construction of a better world. *New Media & Society* 23, 2: 237–257. https://doi.org/10.1177/1461444820929315
[10] Sheila Jasanoff. 2015. Imagined and invented worlds. In *Dreamscapes of modernity: Sociotechnical imaginaries and the fabrication of power.*, Sheila Jasanoff and SH Kim (eds.). University Of Chicago Press, 321–346.
[11] Sheila Jasanoff and Sang-Hyun Kim. 2009. Containing the Atom: Sociotechnical Imaginaries and Nuclear Power in the United States and South Korea. *Minerva* 47, 2: 119–146. https://doi.org/10.1007/s11024-009-9124-4
[12] Rob Kling. 1991. Computerization and social transformations. *Science, Technology & Human Values* 16, 3: 342–367. https://doi.org/10.1177/016224399101600304
[13] Brian Larkin. 2013. The politics and poetics of infrastructure. *Annual Review of Anthropology* 42, 1: 327–343. https://doi.org/10.1146/annurev-anthro-092412-155522
[14] Bruno Latour. 2004. Why has critique run out of steam? From matters of fact to matters of concern. *Critical Inquiry* 30, 2: 225–248.
[15] Bruno Latour and Steve Woolgar. 1986. *Laboratory life: The construction of scientific facts*. Princeton University Press. Princeton.
[16] Harro van Lente and Arie Rip. 1998. Expectations in technological developments. an example of prospective structures to be filled in by agency. In *Getting new technologies together. Studies in making sociotechnical order*, C. Disco and BJR van der Meulen (eds.). Walter de Gruyter, 195–220.
[17] Donald MacKenzie. 1991. Notes toward a sociology of supercomputing. In *Social Responses to Large Technical Systems: Control or Anticipation*, Todd R. La Porte (ed.). Springer Netherlands, Dordrecht, 159–175. https://doi.org/10.1007/978-94-011-3400-2_8
[18] Astrid Mager and Christian Katzenbach. 2021. Future imaginaries in the making and governing of digital technology: Multiple, contested, commodified. *New Media & Society* 23, 2: 223–236. https://doi.org/10.1177/1461444820929321
[19] Sarah Read. 2019. The Infrastructural Function: A Relational Theory of Infrastructure for Writing studies. *Journal of Business and Technical Communication* 33, 3: 233–267.
[20] Sarah Read. 2020. How to build a supercomputer: U.S. research infrastructure and the documents that mitigate the uncertainties of big science. *Written Communication* 37, 4: 536–571. https://doi.org/10.1177/0741088320939541
[21] Sarah Read and Jordan Frith. 2022. Introduction: Writing infrastructure. *Communication Design Quarterly* 10, 3: 5–10.
[22] Sarah Read and Michael Papka. 2014. Genre Cycling: The Infrastructural Function of an Operational Assessment Review and Reporting Process at a Federal Scientific Supercomputing User Facility. In *Proceedings of the 32nd ACM International Conference on the Design of Communication*.
[23] Christian Sandvig. 2013. The internet as infrastructure. In *The Oxford Handbook of Internet Studies*, William H. Dutton (ed.). Oxford University Press, Oxford, UK, 86–108.
[24] Susan Leigh Star. 1999. The ethnography of infrastructure. *American Behavioral Scientist* 43, 3: 377–391. https://doi.org/10.1177/00027649921955326
[25] Susan Leigh Star and Karen Ruhleder. 1996. Steps toward an ecology of infrastructure: Design and access for large information spaces. *Information Systems Research* 7, 1: 111–134. https://doi.org/10.1287/isre.7.1.111
[26] Langdon Winner. 1986. *The whale and the reactor: a search for limits in an age of high technology*. University of Chicago Press, Chicago.
[27] 2015. U.S. Department of Energy awards $200 million for next-generation supercomputer at Argonne National Laboratory. University of Chicago. Retrieved July 24, 2023 from https://news.uchicago.edu/story/us-department-energy-awards-200-million-next-generation-supercomputer-argonne-national
[28] Ryan Weber. 2022. Making infrastructure into nature: how documents embed themselves into the bodies of oysters. *Communication Design Quarterly Review, 10*, 3): 33-45.
[29] Dean Cheng. 2018. China's pursuit of emerging/exponential technologies.
[30] Aurora Exascale Supercomputer. *Argonne National Laboratory*. Retrieved July 24, 2023 from https://www.anl.gov/aurora
[31] DOE Explains…Exascale Computing. *Department of Energy*. Retrieved July 24, 2023 from https://www.energy.gov/science/doe-explainsexascale-computing