



Seam Work and Simulacra of Societal Impact in Networking Research: A Critical Technical Practice Approach

Gloire Rubambiza

Cornell University

Department of Computer Science

Ithaca, New York, USA

gloire@cs.cornell.edu

Hakim Weatherspoon

Cornell University

Department of Computer Science

Ithaca, New York, USA

hweather@cs.cornell.edu

Phoebe Sengers

Cornell University

Departments of Information Science and Science &

Technology Studies

Ithaca, New York, USA

sengers@infosci.cornell.edu

Jen Liu

Cornell University

Department of Information Science

Ithaca, New York, USA

jl3835@cornell.edu

ABSTRACT

This paper explores how conceptions of societal impact are produced and performed during academic computer science research, by leveraging critical technical practice while building a digital agriculture networking platform. Our findings reveal how everyday practices of envisioning and building infrastructure require working across disciplinary and institutional seams, leading us as computer scientists to continuously reconceptualize the intended societal impact. By self-reflectively analyzing how we accrue resources for projects, produce research systems, write about them, and maintain alignments with stakeholders, we demonstrate that this seam work produces shifting simulacra of societal impact around which the system's success is narrated. HCI researchers frequently suggest that technical systems' impact could be improved by motivating computer scientists to consider impact in system-building. Our findings show that institutional and disciplinary structures significantly shape how computer scientists can enact societal impact in their work. This work suggests opportunities for structural interventions to shape the impact of computing systems.

CCS CONCEPTS

• **Networks** → *Network design principles*; • **Human-centered computing** → *Empirical studies in collaborative and social computing*; • **Applied computing** → *Agriculture*.

KEYWORDS

critical technical practice, systems research, networking, anticipation, design for societal impact, values in design, seamfulness

ACM Reference Format:

Gloire Rubambiza, Phoebe Sengers, Hakim Weatherspoon, and Jen Liu. 2024. Seam Work and Simulacra of Societal Impact in Networking Research: A Critical Technical Practice Approach. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '24)*, May 11–16, 2024, Honolulu, HI, USA. ACM, New York, NY, USA, 19 pages. <https://doi.org/10.1145/3613904.3642337>

1 INTRODUCTION

HCI has long been concerned with finding ways to address the societal impact of computing systems [13, 47, 49, 64, 77, 87, 92]. While engineers and computer scientists building systems often ground them in visions of positive societal impact, these systems often have unintended consequences (e.g., [8, 99]). As computing technologies have become ever more ubiquitous and their unintended effects more obviously problematic, HCI researchers continue to explore ethical issues in computing systems development and how to better integrate considerations of societal impact into their technical design [11, 14, 69, 95, 105]. Fundamentally, this work is grounded in the idea that making computer scientists aware of the possibility of particular societal consequences of technical choices will enable them to make better decisions about what to build.

This paper examines what happens to ideas about societal impact when they become caught up in the everyday realities of producing computer science (CS) research. How do CS researchers envision societal impact through the process of envisioning, building and publishing about research systems? How are they orienting towards these ideas of societal impact in the decisions they take in their work? What factors shape the forms of societal impact they take into account, and that they claim for their systems? How and to what degree do these ideas shape what is actually built? This work aims at uncovering technical, institutional, and disciplinary factors that shape how particular research systems are understood to “have” societal impact and the work these claims do to produce particular emerging sociotechnical worlds.

Our work suggests that such structural factors torque how societal impact is imagined, how those imaginations become embodied in concrete systems, and how these systems become enrolled in claims of potential societal impact that motivate publication and



This work is licensed under a Creative Commons Attribution-NonCommercial-ShareAlike International 4.0 License.

CHI '24, May 11–16, 2024, Honolulu, HI, USA

© 2024 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0330-0/24/05

<https://doi.org/10.1145/3613904.3642337>

further research. These structural dimensions complicate the idea that changing computer scientists' thinking about and valuing of societal impact is adequate to produce systems that make a positive impact on the world. They also suggest new avenues for both technical and HCI researchers to shape societal impact by making topical and structural interventions into how technical research is funded, organized, and reviewed.

To make this argument, we integrate social observation and analysis with hands-on technical development. Our work builds on Agre's critical technical practice [3, 4]: a unified practice of technical system-building paired with critical reflection on sociotechnical factors that shape a system's development. We build on four years of critical reflection paired with ongoing technical development of a networking research system to articulate how structural factors shape the ways technical researchers orient to societal impact, and their material consequences. This paper, then, presents an insider's perspective on the day-to-day practice of building computing systems with societal impact in mind.

The context for this paper is an area of CS research called "systems," which aims to improve the performance and reliability of computer systems and networking. Systems research focuses on computational infrastructure, including operating systems, mobile computing, and interprocessing management; the research area we focus on in this paper is networking.¹ Specifically, we investigate the context of building a technical platform for digital agriculture (or DA), which entails data collection, transmission, analysis, and actuation on rural American farms [12, 66, 89, 98]. The platform, dubbed the Software-Defined Farm (or SDF), was originally envisioned as supporting societal impact by addressing the need to feed a growing world population. As the platform has developed, our understanding of societal impact issues in digital agriculture has refined to encompass addressing infrastructural inequalities between rural and urban spaces and reducing risk of vendor lock-in, whereby consumers become locked in by design to particular hard- and software suppliers. Systematic reflection as we build this emerging infrastructure has revealed aspects that enable, shape, constrain, or reorient how we envision societal impacts.

Our study of the SDF relies primarily on critical technical practice because it allows an extraordinary level of access to detailed technical decision-making, while viewing those decisions through a critical lens which can bring powerful new insights. As critical maker Ratto argues,

[C]urrently the distance between [technical and social] areas of knowledge remains vast — often even within the geography of the university campus. However, there is an important need for critical makers that can reintegrate technical and social work and thereby innovate both. [80, p. 258]

To ensure that we are not blinded by assumptions inherent in a personal narrative, we supplement the critical technical practice with an outside perspective through ethnographic observations of SDF technical work. Building on both perspectives, our analysis examines how ideas and claims about societal impact are enacted

and oriented towards within the everyday practices of envisioning, building, and publishing about research systems.

Our key argument is that, even when computing systems researchers envision and work toward pre-determined societal impacts, they work towards these envisioned impacts within institutional, infrastructural, and disciplinary structures that gradually rework this envisionment in both direct and subtle ways. The results section presents a detailed, reflexive account of processes of envisioning societal impacts with differently situated stakeholders, and describes how those envisioned impacts are transformed as we set technical goals and build infrastructure, disseminate our findings and navigate CS peer review, and produce accounts of the success of our work for different stakeholders. We show how, throughout these activities, we must continuously work across technical, institutional, and disciplinary seams. We characterize such maneuvering across sociotechnical structures as 'seam work'. Further, we argue that navigating these seams continuously refactors the way the system and its envisioned impacts are represented, narrated, and oriented towards, in response to structural pressures at each seam. Thus, we argue that, while researchers do have agency to work towards societal impact, our agency is compromised [63] by our placement within institutional structures that subtly guide research agendas and outputs toward outcomes that benefit or reinforce existing financial, legal, intellectual property, or other structures. As critical technical practitioners, we are far from being innocent outsiders to such pressures; we must also work with(in) these institutional and disciplinary powers to attempt to maintain course toward envisioned societal impacts. We articulate how critical technical practice supports such navigation.

We argue that convincing computer scientists to care about different values may be ineffective, given the technical, institutional, and disciplinary constraints on how such care can be enacted in technical work. We propose three alternative ways that the HCI and computer systems communities can work together to better support technical research aimed at improved societal impact: by redesigning our institutional and disciplinary systems, by systematically deploying critical technical practice as an effective tool for investigating infrastructure development, and by fostering what we term trilinguals of the future — i.e., individuals bringing together computing technology, domain understanding, and critical reflection — as a way to improve the broader impacts of computing.

2 RELATED WORK

This paper is particularly concerned with reflections on societal impact within academic CS research. A key HCI approach to exploring how computer scientists do and could orient to the societal impact of computational systems frames these issues as questions of values and design [39, 73, 75, 76, 90, 93]. This work focuses on how technology developers often unconsciously build their own values into systems, long before they become user-facing end products. Often, these approaches pair social scientists with expertise in ethical and social issues with developers with technical expertise, with the aim of collaboratively identifying and altering assumptions in early system design that may have negative consequences downstream. For example, Cheon and Su worked with robotics researchers to

¹We acknowledge that, for CS researchers, systems and networking fall under different professional organizations - SIGOPS and SIGCOMM, respectively. However, SIGOPS considers its topics of interest to also include networking [2]. We explicitly use the term 'networking' where the distinction between networking and systems is needed.

envision their own future autobiographies as a way of exploring values issues in robotics [22].

As this work has evolved, researchers have identified recurrent challenges in engaging computing researchers in discussions around values and unintended societal impacts that arise from their disciplinary and institutional situation. For example, Shilton et al. describe challenges in engaging CS researchers in values discussions that arise because these researchers frame the goal of building general-purpose systems as requiring values neutrality [91]. More generally, as Ribes et al. have described [81], CS researchers often frame their work in terms of a separation between more general-purpose technology and specific domains in which that technology would be applied. In this framing, computer scientists oversee the general-purpose technology, while domain experts are in charge of domain-specific details. This separation would tend to suggest to practitioners that societal impacts are part of an application domain, and thus not in computer scientists' purview.

Recently, Do et al. explored how academic computer scientists orient towards unintended consequences of their work [28]. They showed that while computer scientists did care about unintended consequences, they rarely took action to address them during technical development. Causes included a lack of structured proactive methods, faith in their own estimation of the societal benefit of their work, the siloing of CS from community and other disciplinary experts who may be better placed to anticipate unintended consequences, and the pressure to publish quickly.

In this paper, we build on this prior work to explore the factors that enable, shape, constrain, and often blinker the consideration of eventual societal impacts in the everyday practices of academic CS. While Do et al.'s work relied on retrospective accounts by computer scientists, we here focus on how these pressures manifest within the everyday, mundane work involved in CS research, in ways that often escape computer scientists' notice as they focus on the difficult task of producing publishable research.

We particularly focus on the *envisionment* of impact, rather than the eventual actual impact of the systems being built. By "societal impact", we refer to broader impacts of technical systems in end users' lives beyond the simple technical functioning of the system. We understand such societal impacts to be produced by an unfolding, long-term interaction between technical possibilities, political-economic forces, and the actions of stakeholders in navigating each of these (e.g., [68]). This process is incompletely predictable at design time; while engineers often aim towards specific possible impacts for their work, its actual impact will not be known until long after the work has been taken up in society [35].

Thus, our paper engages with the *prospective* imagining of societal impact in technoscientific research, also termed speculation, expectation, envisionment, or sociotechnical imaginaries [61]. Such speculation plays a central role in technoscientific research, since imaginaries guide engineering ideas about what can and should be built [65] and are used to enroll funding, labor, and other resources to ongoing projects [52, 57]. Researchers can and must use visions of societal impact to argue for the legitimacy of their project, in ways that are constrained by the present-day logics of the stakeholders to whom they are accountable [52]. Visions of societal impact matter, not because they predict the future, but

because they performatively shape it, though not always in the ways that researchers intend [61].

Our work focuses specifically on reflection on societal impact within academic networking research. Here, we build on a small but significant literature in HCI and CSCW exploring social dimensions of computing from the perspective of networking researchers. Such researchers have explored how communities develop and maintain their own networks [30, 32–34, 42], cultural and societal challenges in efforts to establish new networking infrastructure [15–17, 31], how networking design choices do and should land in users' lives [23–26, 36, 43, 48] and considerations for the design of new networking infrastructure from communities that are structurally underserved [18].

Taken as a whole, this body of work showcases a deeper engagement with social science among (some) academic networking researchers that run against the previously reported narrative of disengagement with questions of societal impact [28]. Further, it illustrates the complex sociotechnical dimensions and consequences of networking infrastructure. In so doing, it demonstrates the power of integrating deep familiarity with the technical dimensions of networking systems-building with social-scientific consideration for HCI/CSCW. As Dourish demonstrates, the material specificities of networking systems are deeply consequential [29]. As Edwards points out, these "lower layers" of infrastructure are crucial for HCI to reflect on, yet access to them is mostly outside the HCI methodological toolkit [37]. Our work adds to this interdisciplinary conversation by developing approaches to critically explore the structural possibilities and limitations of academic networking research itself during the everyday practice of doing such research.

To do so, we build on Agre's notion of critical technical practice [3, 4]. Critical technical practice is a methodology that integrates computational development with critical reflection such that "rigorous reflection upon technical ideas and practices becomes an integral part of day-to-day technical work itself" [3, p. 3]. Agre used critical technical practice as a tool for identifying inherent conceptual limits in how Artificial Intelligence (AI) was framing the problem of producing intelligent action, identified alternative conceptualizations of such action, and developed technical innovations which embodied these alternatives [5, 6]. Similarly, critical technical practice has been drawn on in HCI to explore and alter values and assumptions built into computational systems [1, 38, 58, 60, 86, 87]. But Agre did not just use critical technical practice to reflect on and improve technical systems; he also used it to explore the structure of AI as a field, as everyday practice, and as a discourse which enables certain forms of work and sidelines others [4]. It is in this vein that we use critical technical practice here as a tool to reflectively explore what happens when academic CS researchers aim to develop technical research systems with the goal of achieving positive societal impact.

3 APPROACH

In this work, we examine how societal impact is envisioned and worked toward in the everyday practices of academic computer systems research, with a specific focus on networking research. We do so by engaging in networking research and system-building

based on specific visions of societal impact, while reflexively analyzing the process. Our point of departure was to enrich technical design for societal impact by critically reflecting on how definitions and practical orientations to societal impact are shaped by assumptions, institutional and disciplinary pressures, and work processes throughout the process of system-building. In this section, we will first provide an overview of the case study used in this paper: the development of a new networking infrastructure for digital agriculture. Then, we will describe the approach to critical technical practice that we use in this paper.

3.1 Case Study: The Software-Defined Farm (SDF)

Our work is oriented around a case study of the Software-Defined Farm (SDF) (Figure 1). The SDF is a technical platform intended to support digital agriculture, which refers to data-driven techniques to collect, process, and actuate on farm data [12, 66, 89, 98]. Prior to our work on the SDF, our research group had no experience with the agricultural domain. The conceptualization of the SDF arose from our university's investment in new interdisciplinary collaborations in data-driven agriculture. Through this initiative, we became aware of the problem of feeding a growing world population, and eager to apply our networking expertise to it.

Our project came to focus on two main challenges in the production of digital agriculture. The first challenge arises because digital agriculture as currently envisioned is premised on the collection of massive amounts of on-farm data [21, 102]. This vision is challenged by the reality of sparse Internet connectivity and unreliable power in the rural regions where the vast majority of farms are located, which historically face infrastructural deficits produced by urban-centric technology development [19, 50, 56, 71]. Thus, digital agriculture faces challenges in collecting the envisioned volumes of data, processing them, and communicating with a not-always-available cloud, the paradigm under which most commercial data processing services currently operate.

The second challenge we came to center in our project is vendor lock-in. In the agricultural context, vendor lock-in arises when commercial sensor vendors limit consumer choice by allowing data access only through their approved mechanisms and application programming interfaces (APIs). A similar issue can be observed in cloud computing services [78]. Vendor lock-in renders customers reliant on a single data provider, imposing interoperability, data sovereignty, and data privacy issues [41].

The SDF is therefore designed to provide a plug-and-play architecture which will allow the modules that handle data sensing, actuation, and the cloud-based processing of data to reliably - if intermittently - interconnect. To address the paucity of rural Internet, the SDF relies on existing high-bandwidth technologies, such as 4G LTE, and new networking techniques such as long-range radio (LoRa) [7, 102] and TV White Spaces (TVWS) [10, 82, 84]. The core engineering challenge of the SDF, then, is to disentangle the hardware devices from the overlaying software infrastructures in order to enable plug-and-play recombination of different sensors, networking devices, and software modules. For instance, a consumer could interconnect TVWS networking, data storage from the

Google Cloud Platform (GCP) [46], model storage from the Azure Cloud [72], and computing devices at the farm.

3.2 Methodology

We approach the development of the SDF through an analytic lens drawn from infrastructure studies as pursued in CSCW [54, 97]. Sandvig [85] describes this school of infrastructure studies as focusing on infrastructure as always emerging within human practices and contexts, with five main methodological commitments: (1) to make visible normally invisible aspects of infrastructure; (2) to focus on the human practices and hidden labor underlying infrastructure; (3) to “find and make comprehensible the invisible negotiations that are producing the infrastructure” (p. 98); (4) to study the boundaries and gateways through which infrastructures interact; and (5) to pay special attention to the early days of infrastructural development when future path dependencies are often built in. Here, we embody this approach by *making visible* the processes by which new infrastructure is envisioned in the earliest stages of its conceptualization. We describe the nature and consequence of the *tacit labor* of developing new infrastructural concepts. We show how these concepts are produced through *negotiations* between differently placed institutional actors, in ways that lay the groundwork for *path dependencies* that may diverge from actors' intentions. Because the infrastructure is not yet standardized, we do not address boundaries and gateways.

Our methodology builds on Agre's articulation of critical technical practice, as previously described [3, 4]. Agre's work is not methodological in the sense of developing a “recipe” that is easily reproduced in other circumstances [38]. But the core commitment of Agre's approach is to continuous technical system-building and experimentation, coupled with rigorous reflection on the assumptions, limits, and consequences that are implicit within the technical process. This work thus involves developing a conversation between critical, social research and conceptualizations emerging from technical work [4].

Our project was likewise based in an integration of computer networking research and critical social science. The first author, Gloire, is formally trained in computer systems, networking, and social science. The third author, Hakim, is a researcher in computer systems and networking. Together, Gloire and Hakim designed, developed, and deployed numerous SDF prototypes on our university's test farms over four years, in collaboration with other students and faculty. They also actively participate in the computer systems and networking communities through conferences and other formal events. This technical work was the medium through which the critical technical practice was pursued.

The second author, Phoebe, is trained in CS, but currently works in social studies of technology. During the technical development, Phoebe was the PI of an ethnographic study of the process of SDF development, in which all the authors participated along with other collaborators. Phoebe and Hakim, who are tenured professors, are Gloire's Ph.D. co-advisors; this detail will matter for the story told below. Ethnographic observations were conducted over three years (2020-2022) and were led by the fourth author, Jen, a PhD student in Information Science; her observations included visionary workshops, technical research meetings, and in-situ deployments of SDF

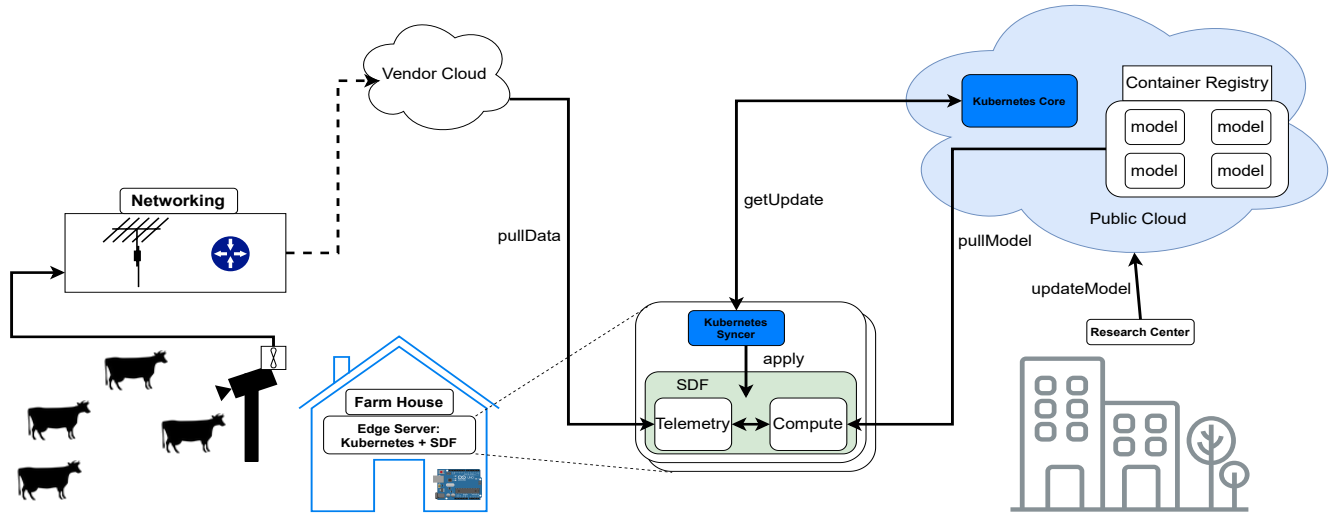


Figure 1: The SDF is envisioned as a general architecture to sense, transmit, analyze, and actuate on rural American farm data.

prototypes. Phoebe’s and Jen’s role on this project was to provide an outsider perspective that would counterbalance the danger in autobiographical projects to be blind-sided by one’s own assumptions [74]. Phoebe also advised on social-science methodology and contributed to the analytic development of the argument ².

Our approach in this paper builds on the way that Agre used reflection on his personal experiences as an AI researcher to explore the shape and limits of the field’s discursive structure and practices [4]. We provide an insider’s perspective on the everyday practices of computer systems research by reporting on Gloire’s experiences and reflections building research systems. The observations and analyses reported here are primarily Gloire’s. However, as our study will underscore, critical technical practice does not happen in an institutional vacuum, and our use of the pronoun “we” respects the role that the co-authors play in informing Gloire’s ideas and their expression.

Critical technical practice as a method relies on a technical practitioner reflecting on their own practice. As such, it bears affinity with autoethnography [70], an established HCI method (e.g. [44]). Like autoethnography, critical technical practice trades off the breadth of multiple cases for deep, long-term engagement in a single case. Like analytic autoethnography [9], we here use reflection on personal experiences to advance theoretical conceptualization of how systems research produces and performs visions of societal impact.

Agre notes the trickiness involved in reporting reflexively as a technical practitioner [4], including the temptation to present one’s own experiences as setting the standard for right action:

I am not interested in in portraying myself as a victim of circumstance or an innocent party in a conflict... Nor... would I wish to portray myself as Jesus among the Pharisees – the virtuous hero who uncovers the

corruption of traditional learning and yet fails to persuade the learned of their errors. Mine is not a tale of virtuous heroism, heaven knows, simply of the historical conditions of a path [4].

Like Agre, our aim is not to argue for our work as the right way to address societal impact in technical work. Rather, our goal is to provide a nuanced, experiential, and (self-)critical account of defining and working towards societal impact within contemporary computing systems research in the US, and thereby to better illuminate its structural conditions and how they shape envisionment of and orientation towards societal impact by academic systems researchers.

3.3 Data Collection and Analysis

The work detailed here builds on prior work reflecting on the early system design, implementation, and deployment phases of the SDF [83]. For this paper, to mitigate the danger of unreliable memories and retrospective bias, we focus our analysis on biweekly diary entries which Gloire systematically recorded for critical reflection for a period of 12 months (May 2022 - April 2023). The diary entries were recorded using markdown in a private GitHub repository with linked evidence to technical deployment for easy retrospective analysis. These notes reflected on software development, hardware deployments, and working across disciplines. Gloire conducted the initial coding and qualitative content analysis of the diary entries, with feedback from the second and third authors.

Approximately nine months into the twelve-month data collection phase, the diary entries narrowed their foci to mostly reflect on working across fields to publish our SDF lessons to date and onboard new ideas, funding, and partnerships for the next iteration of the platform. In particular, the last few months’ entries included subsections on writing papers across fields, the different community reactions to our results, and the tensions in initiating or maintaining collaborations across disciplines and organizations.

²Note that, except for the authors and systems named so far, all system, conference, and stakeholder names henceforth are pseudonyms.

Where appropriate, Gloire still reported on any carry-over software development/debugging and prototype deployments. Besides the technical transitory period, the narrowed focus was also driven by Gloire's initial thematic data analysis which highlighted trends related to the formulation of societal impact in computer systems research through a critical technical practice lens. Therefore, Gloire's analytic work came to focus on the challenges of working and writing within and across domains, navigating the computer systems community reactions to our work, and collaborative tensions to (re)-align on societal impact goals. For each theme, Gloire drafted memos and highlights that facilitated discussions with the other authors. The goal of these discussions was to refine the themes and core argument by counterbalancing them with insights and perspectives derived from the larger, longer-term ethnographic project.

As Gloire further developed his analysis in collaboration with the other authors, we began to realize that an important thread through his observations was a navigation of technical, disciplinary, and institutional 'seams' and how this navigation came to shape the technical team's work. The concept of seams originates in Weiser's vision of ubicomp technologies as seamlessly connecting digital and physical spaces [104], in which seams could be understood as bridging points between digital and physical infrastructure. Since this articulation, both seamlessness and seamfulness — i.e., design to explicitly expose infrastructural seams to user awareness [20] — have been repeatedly explored as design qualities in HCI [53]. Researchers in Science and Technology Studies (STS) have taken aboard the language of seams, seamlessness, and seamfulness that resulted to articulate how users actively and creatively work across the boundaries between infrastructural systems in order to accomplish their goals [96, 103]. Here, we focus on the efforts of networking researchers to work not only across technical, i.e. classically infrastructural, seams, but also across the seams between organizations and disciplines, in order to stitch together technical systems and resulting research contributions intended to produce societal impact in a specific domain.

4 RESULTS

Because the results section was produced by Gloire as a self-reflexive analysis of his work and experiences, this section is written from his first-person perspective. In this section, the pronoun "we" refers to the technical team Gloire was working with, which included Hakim and other actors, but not the other co-authors.

In this section, I describe how we oriented towards envisioned societal impacts in different stages of computer systems and networking research in response to opportunities, pressures, and expectations from ourselves, other stakeholders, and our discipline. The first four subsections explore different stages in research development: from project specification, to technical development, to writing up results, to navigating peer review. The final subsection explores how we came to define success to ourselves and project partners at different stages of the project. Each subsection delves into how and why our envisioned societal impacts were in focus, out of sight, re-framed, or inadvertently forgotten at each stage.

4.1 Pitching a Project

The process of orienting towards societal impact started with the inception of our project. The research question in the abstract of the SDF GitHub repository is: 'how do we sustainably feed 10 billion people in 2050?' This question was adopted from our university's collaborative research thrust for digital agriculture, which had inspired our team to begin working in the agricultural domain. As initially formulated, this vision was untethered from technical possibilities. As our research progressed, our ideas of how to contribute to feeding the world through our research evolved in response to funding opportunities, accessible stakeholders, and the needs and desires of different collaborators.

We began our project by listing societal impact goals, defining technical goals, enrolling institutional stakeholders, and submitting proposals for funding. The initial version of SDF (i.e., SDF 1.0), became funded through grants from the US National Science Foundation (NSF), the US Department of Agriculture (USDA), and CloudAndMore, a large tech company which runs proprietary cloud computing services. We articulated the intended societal impact of our system by developing a pitch that spoke to these stakeholders. We came to envision the SDF 1.0 as providing new networking capabilities while anticipating the societal impacts of high-bandwidth networking on American farms in rural areas. Given the paucity of Internet in these areas, the networking would rely on innovative modalities such as long-range radio (LoRa) [7, 102] and TV White Spaces (TVWS) [10, 82, 84]. The envisioned applications for this first iteration of the platform were developed in collaboration with agricultural researchers at our university; as a result, we intended to support early cow disease detection in dairy farms, early virus detection in vineyards, and measurement of plant water stress at our university's apple orchards and greenhouses. The realization of the platform thus brought together computer scientists and engineers, animal scientists, biomolecular and chemical engineers, plant scientists, and researchers and engineers from CloudAndMore, whose interests were all reflected in the platform pitch.

In order to elicit funding from governmental agencies and industrial partners, visions of the SDF often tout ubiquitous, end-to-end, scalable platforms that speak directly to funders' keywords. For example, one grant pitch described a system grounded in "(1) tech-enhanced innovation, (2) data integration, and (3) data analytics and decision support" (Research Diary, 01/20/2023), with the planned research system producing a data pathway from corn to cow and back:

As an example, it would be good to monitor the conditions under which the corn which is fed to cows is grown. From there, the cows can also be monitored, and any conditions can be traced back to the growing conditions for the corn. Finally, the manure from the cows can also be sent back to the corn fields. This, in essence, is a closed-loop system that is supposed to achieve the [USDA's] farm of the future. - Research Diary, 01/20/2023

The apparent success of SDF 1.0 was a catalyst to enlist new funding and collaborations for SDF 2.0, which in turn reshaped the societal impacts towards which we were orienting. The SDF 2.0 project brings together computer scientists, animal scientists, plant pathologists, and research and engineers from another technology

giant. It is funded through grants from NSF, National Aeronautics and Space Administration (NASA), and USDA. These stakeholders all therefore play a role in determining what the SDF 2.0 should be for. This version of the SDF developed in the context of a new university research center focused around a vision of digital agriculture enabling better plant-human communication. The SDF 2.0 is thus intended to support holistic, two-way communication with plants and animals in an imagined Internet of living things (IoLT). Additionally, based on our experiences managing vendor lock-in during the development of SDF 1.0, SDF 2.0 is intended to address the issue of vendor lock-in, where consumers are locked into one (cloud) provider, and moving to a new provider presents numerous technical, security, scalability, and other issues [78]. In the agricultural context, vendor lock-in leads to challenges in user autonomy and privacy [41]; it also leads to technical difficulties in farm data aggregation and analytics because sensor data are collected and processed in silos [83].

In summary, the evolution of the SDF is predicated upon and shaped by long-term partnerships between computer scientists and stakeholders in other domains to work toward shared visions. These partnerships are formally structured through shared vocabularies (e.g., ‘broader impacts’), principal investigators (PIs) and technical leads, and reporting mechanisms intended to keep stakeholders aligned on particular envisioned societal impacts. While the PIs are in charge of formal reports, the system design, implementation and deployment processes are characterized by continuous negotiations across institutional and domain structures done mostly by technical leads (i.e., graduate students, software engineers, and research scientists). This work shapes how the systems are envisioned and pitched, and which societal impact goals are considered relevant and addressable through project work.

4.2 Building Computer Systems

Envisionment of societal impact is not only shaped by stakeholders’ intentions at the start of the project, however. Even as our team was explicitly working toward specific desired impacts, the possible scope of that eventual societal impact was being reshaped during the mundane, day-to-day work of building computer systems through low-level software, hardware, and infrastructural negotiations.

As is typical in networking research, to produce SDF 1.0 I spent years of effort deep in the technical trenches dealing with low-level technical issues. Such everyday technical work requires persistent trials-and-errors that often take many hours. For example, low-level technical issues derailed me repeatedly as I tried to build an Ubuntu image for the Raspberry Pi 4B (Research Diary, 05/20/2022). Sometimes such trouble-shooting even takes months or years:

After a year of troubleshooting, I think I figured out the bug with resending of the [cow] records... it’s because... the protobuf was reducing the precision of the timestamp to be sent to the compute module... The fix was simply changing a float to a double in the protobuf definition of a sensor update. - Research Diary, 01/06/2023

Caught up in these low-level frustrations, it became increasingly difficult to notice how the infrastructure-building process itself

could be shaping our eventual impacts. For example, we might decide to work with certain cloud providers or operating systems because they are more amenable to our software development or hardware deployments. A simple hardware vendor choice (i.e., Raspberry Pi 4B) influences, in increasing order of importance, the hardware architecture, the operating system image, the programming languages, the middleware packages, and the cloud provider. These choices can all influence how the ensuing system works, or does not work, to address our societal impact goals, for example by putting us in a situation where vendor lock-in is an unavoidable outcome.

Producing a complex, long-term system like the SDF requires the enrollment of many stakeholders, including undergraduate research assistants, graduate students, faculty members, and partnerships within/across institutions (e.g., industrial partners, building managers, carpenters, etc), who then both support and constrain how our project can unfold. Thus, issues arising from technical constraints are compounded by the number of stakeholders involved in our technical work whose concerns need to be addressed. Our technical infrastructure in principle was defined around a set of desired principles and knowledge production processes we believed would support our technical and societal impact goals, but sometimes these principles had to be revised. For example, we planned to maintain all software artifacts emerging from the research as open source, as is often mandated by funding agencies such as the NSF. Maintaining an open source and reproducible knowledge production process across domains is rife with negotiations, especially in teams involving computer scientists, who will need to coordinate on GitHub repository structures, branch naming and pull requests, etc. Moreover, navigating open-source goals is tricky when working with domains such as animal science that engage directly with commercial vendors providing proprietary hardware and software:

We had the meeting regarding [the progress on] open-sourcing the SDF [1.0] code. Giancarlo’s [research] group raised concerns about the embedding of their commercial [sensor] providers’ names in the dairymanager code. Therefore, we have decided to open-source the portions of the code sans dairymanager, and then the plan is to sanitize the dairymanager code of any commercial [vendor] names before adding it back to the rest of the open-source code. - Research Diary, 12/23/2022

Giancarlo is a PI leading an animal science research lab, and their concerns are valid given their direct relationships with commercial sensor vendors. However, removing vendor information from open source code also stripped the software artifacts of any contextual information on commercial vendors. This removal, in turn, subtly inhibit our efforts to address vendor lock-in by dulling the modular power of the SDF’s software abstractions. Further, it complicated our efforts for reproducible research by slowing down the open-source process.

4.3 Writing Across Disciplines and Academic Rungs

So far, we have seen that visions of societal impact are produced as projects are conceptualized and pitched, and that orientations

to societal impact become complicated and reframed during the process of building test systems. But visions undergo a further transformation in the process of producing publications about those systems and their implications. Such publications are an integral part of institutional practices for career success, tied to factors such as student graduations and faculty promotions. They also form the major avenue by which academic systems researchers conceptualize the path to societal impact for our work, with the idea that published insights may be adapted by the community through further research, and potentially adopted by commercial actors to be integrated into widely available products and services.

During the course of data collection, we wrote and submitted two technical papers - one for the systems audience and one for plant pathology. My plant pathology collaborator and fellow PhD student Frank and I co-led this second paper, which presented a vineyard disease application of the SDF and was submitted for review to a plant pathology journal in December 2022. This paper was co-written with our respective advisors, Laura (a tenure-track professor in plant pathology), and Hakim, who is tenured in systems. This paper relied on a parallel paper led by Frank to another plant pathology journal. While his paper presented the plant pathology advancements, our joint paper described how the SDF enabled these advancements.

As Ph.D. students leading the writing efforts, Frank and I found ourselves negotiating across competing domain value systems to produce technical content that is digestible by both disciplinary audiences. The need to navigate multiple domains in writing was further complicated by power differences among authors at different rungs of the academic ladder: Ph.D. students, tenure-track professors, and tenured professors. These aspects came together in an incident that occurred while Frank and I were writing this paper. Upon sharing a draft of the paper with my co-advisor Phoebe, who is tenured in HCI, she challenged my use of the technical term 'abstraction' in the manuscript, arguing the plant pathology audience would be unlikely to know what this term meant. Phoebe's identification of abstraction as an issue in this paper was ironic because we were simultaneously working on a social-scientific paper that was examining the role that abstraction plays in the societal impact of digital agriculture systems, thus I should have been sensitized to trickiness around claims of abstraction. I responded to her email comment as follows:

'Here, again, Hakim rewrote what I had described... I tried to stay away from words like abstraction. [W]hen Hakim made that suggestion... I incorporated it. In retrospect, that sentence is more digestible to a CS audience than a general one. What's ironic is that, as a collective of writers (i.e., Hakim, Laura, me, Frank), we went through the review phase with a goal of writing our paper such that both of our communities can grasp everything... [I realize] certain words are simply seeped so deeply in our CS minds that we take them for granted as part of lingua franca' - Research Diary, 12/09/2022

This incident highlighted several factors. First, it demonstrated how difficult it is to speak comprehensibly across multiple domains, and

how the process of negotiating across domains can lead to blinkering of social scientific insights. This led me to put aside my critical perspective to produce technical content for an interdisciplinary audience. As another example, I failed to notice a questionable claim in the draft that "food security through consistent and scalable agriculture is the invisible foundation of modern society." This discomforted me, since I felt it important to integrate my critical and technical identities and disturbing that I could abandon my critical identity without realizing it.

Second, the blinkering of critical perspectives was exacerbated by the academic power dynamics that led me to defer to Hakim's favoring of the term abstraction, prioritizing faculty assertions and problem/solution framing. And because on this project it was one of the junior researchers who was the biggest advocate for critical reflection, it was relegated to the backseat. This led to criticism from the other faculty member I was accountable to, Phoebe, adding another stakeholder whose views I needed to manage in the writing process.

4.4 Navigating Peer Review in Computer Systems

The previous section demonstrated how the writing process is laden with challenges in finding the appropriate technical depth, taking CS concepts for granted, and avoiding critical reflection on writing claims in deference to faculty seniority. But challenges that reshape how societal impact is framed continue after drafting publications through pressures from peer review. I will describe this by reflecting on a technical paper whose writing I led, which described the SDF architecture under the pseudonym of Chameleon. Prior to data collection, we had already submitted this paper twice (Conference A and B); during the period covered in this paper, we submitted it for review three more times to different systems conferences (Conference X in July 2022, Conference Y in September 2022, and Conference Z in January 2023).

The systems research community operates under a set of implicit principles and value systems which are disciplined through the review process in ways that shape how work can orient to societal impact. One aspect that reviewers emphasize is novelty. Novelty is understood in systems research as producing new configurations of technology. But many systems designed for societal impact involve getting known technologies to work in different, often low-resource contexts, or altering the political-economic functioning of the system, innovations that reviewers may not see as novel. For example, the SDF replicated the functionality of commercial systems to connect sensors with cloud processing, but did so in a way that reduced the risk of vendor lock-in. This led one reviewer to comment that they "see plenty of description of what all technologies you used. But maybe there was a way to highlight the novelty even better" (Anonymous review, Conference A, 2022).

More generally, the networking community currently favors technology development related to running commercial, large-scale data centers, leaving other forms of domain-oriented research out in the cold.

The recent meeting of researchers that dabble in networking + HCI got us thinking that it may be time to push the community toward being more open to

this kind of work. The common issues identified were that the community is too focused on certain kinds of work (ahem datacenters), and this in turn makes publishing and getting funding for other kinds of interdisciplinary [work] difficult. - Research Diary, 5/20/2022

In fact, the apparent hostility in the networking community to interdisciplinary, impact-oriented work was exactly what motivated this group of related researchers to band together for mutual support.

Another value in systems research that makes it challenging to produce research oriented to societal impact is that the community favors generalizability of system ideas over getting a system to work in any particular application domain. This produces challenges for impact-oriented research, which often addresses impact on particular communities or in a particular setting. In our case, systems reviewers characterized our primarily applied work as having insightful deployment experiences for the community, but the “paper itself is a bit under-cooked” (Anonymous review, Conference X, 2022). Comments like these imply that the paper does not contain enough generalizable insight to be worth publishing.

The mismatch between applied research concerned with societal impact and the framings of ‘novelty’ and ‘generalizable findings’ regularly led to negative, sometimes biting reviews. For example, one review of the Chameleon paper from Conference X was completely dismissive of the work and its authors:

One of three reviewers was particularly *vicious*³, to the point of saying none of the authors, including our PIs, had any evident experience submitting to a conference like [X]. I’ve been trained by my advisor [Hakim] to build thick skin as a member of [the] systems/networking [community]. The community is very negative (typical acceptance rates for conferences $\leq 20\%$), but this [Conference X] reviewer hit a new level of negativity! - Research Diary, 09/30/2022

The typical outcome from such reviews is that paper drafts require multiple submissions before getting accepted for publication in systems conferences. In fact, at two separate top-tier systems conferences, a fellow Ph.D. student described submitting work more than three times before getting accepted.

In our case, this led the system paper to go through multiple double-blind submissions that effectively molded it into a beast that is familiar to the community. The continuous recrafting of the paper in response to review is evident from a reflection after receiving reviewer feedback for our Conference X (July 2022) submission:

We received reviewer feedback from... our third [re]-submission. The paper did not get accepted... [the] theme is... reviewers are still having a hard time deciphering our “scientific” contributions... we took [the] comments with a grain of salt... the surgical changes focused on having clearer takeaways at the beginning and end of each section especially the later sections - Research Diary, 09/30/2022

To counteract the emphasis of the community on novelty and generalizability, we started leveraging a different value in the systems community: the importance of deep system building and deployment of ‘real’ prototypes, as Jen had highlighted in her ethnographic work. My Ph.D. co-advisor, Hakim, often refers to this ethic as ‘getting nerdy with it’. In other words, the community appreciates novel systems that have been kept running in what is considered “the real world” for months or years. As Jen pointed out, during our ethnographic research meetings, Hakim routinely instilled this value as ‘don’t pull the wool over people’s eyes’. Ironically, the more times the paper was rejected, the more we could highlight the length of the running SDF deployments as a sign of quality for the system, which could counteract some of the demands of novelty.

A second tactic that we eventually used to counteract the perceived lack of ‘novelty’ and ‘practical insights’ was to follow up on suggestions from reviewers to submit to a different paper track:

We received 2 weak rejects and 1 weak accept from our [Conference Y]’23 reviewers. Overall, they were excited about the case studies and practical experience. In fact, two of them were surprised that the paper was not submitted under the Operational Systems track. This is leading us to cut down the paper to 11 pages and submit to [Conference Z] (Z ’23) under the Deployed Systems track. This would be the *fifth* submission of the work, and its potential acceptance could cement my position as a sociotechnical researcher. In other words, I don’t have to let go of my critical perspective. - Research Diary, 12/23/2022

The ‘Operational Systems’ (also known as ‘Deployed Systems’) track is a separate submission track primarily dedicated to applied work. Submitting there was particularly advantageous because those papers “need not present new ideas or results to be accepted; indeed, new ideas or results will not influence whether the papers are accepted” [101]. However, with a few exceptions [50, 51], this applied track is mostly dominated by large technology giants with billions of users [45] — systems at a scale that far outreaches the scope of what most societal-impact-oriented academic researchers are in a position to produce. As a junior researcher invested in societal impacts, competition against tech giants whose impact can easily reach billions of users, paired with vicious reviews, led to constant imposter syndrome.

While we *were* eventually able to successfully publish this work, an unintended consequence of the incremental content restructuring this required was the gradual sidelining of societal impact considerations as the paper moved closer to acceptance and took on a form more recognizable to the systems community:

The Chameleon paper got accepted to [Conference Z ’23]!!!! This was a big surprise given the last diary entry! The reviewers actually spent time reading our rebuttal, and it was the deciding factor because it outlined some of the ‘practical insights’ that the reviewers had been wanting... It took **FIVE** tries for the paper to be accepted, and the acceptance rate was **18%**! The reviewers were fair in saying that the

³Emphasis in original diary entry

insights need[ed] to be pulled out from the deployments... However, our work with the abstraction paper described in the previous diary entry highlight again how the goal is to make the insights general (i.e., placeless) for other [system] designers. Overall, this paper has taken multiple years to get accepted, and the final product may end up [with] a lot of the context stripped out to accommodate the inclinations of the CS systems/networking community.” - Research Diary, 04/28/2023

The abstraction paper referred to above is a critical paper from our project led by social scientists. The paper, which is under review at the time of writing, examines how the disembedding of agriculture from its contexts associated with the rise of industrial agriculture may be replicated in contemporary digital systems for farming, and outlines the risks and costs of such disembedding. It was therefore ironic that a similar disembedding of our own technology was exactly what we had to do to get our work published. In the Chameleon case, we had gradually stripped back the societal impact considerations in the publication to reveal what reviewers would likely consider the “essential” systems/networking insights. In prioritizing publishing our work, and with the goal of having junior members of the team be recognized as members of the systems community, we had engaged in what we term “contextual disembedding:” i.e., we had decided to forgo reporting contextual information about relationships between places, people, and organizations involved in the research, in the hopes that this would make the work seem less “applied.”

4.5 (Re-)defining success

The prior results have shown how envisioned societal impacts become constructed and reshaped through setting visions, building infrastructure, writing up findings, and getting through peer review. Here, I step back from these individual stages to consider how, across these stages of research, we construct success of our project in moving towards desired societal impacts to different stakeholders. Agreements on milestones that define “success” must at least temporarily align expectations between academics, industrial partners and funding agencies. Here, I describe how developing notions of success towards societal impact goals is intertwined with publishing and funding gymnastics.

The envisioning process brings together different stakeholders (e.g., agricultural, technology, academic, and funding partners). First, for agricultural stakeholders, the envisioning involved conducting interviews and visiting their work sites to guide our system design for SDF 2.0. Even within the same industry (e.g., vineyard management), stakeholders face significantly different pressures (e.g., disease detection versus labor optimization) to be seamlessly addressed by the same system, let alone across application domains (e.g., dairy cow management and vineyard management). Thus, “success” for our system is co-defined with stakeholders, based on their specific domain goals.

Second, as envisioned impacts were being brainstormed with technology partners before developing a legally binding joint study agreement (JSA) which would govern intellectual property on the

project, I faced a level of uncertainty about what they would consider progress:

... It is unclear if/how working with tech giant will dictate the kind of research questions we can ask. In particular, the tech giant would [likely] push the agenda for [their edge computing platform] and its use case in this context. However, it is still unclear to me if this will be more of an engineering or research effort. - Research Diary, 03/05/2023

It was important for me to understand the nature of anticipated efforts between engineering and research tasks because our reviewers care mostly about the systems and networking research questions and consider the engineering questions to be “merely applied.” Moreover, understanding our technology partners’ incentives is important to keep them engaged because they provide technical resources and help in using their platforms. Is it enough for them to claim that their platforms were used in the system implementation? Is the claim that their cloud services were used in a paper publication sufficient? This practice of probing corporate incentives in research partnerships is not new. I had previously observed similar comments in interviews with academic researchers working to realize the potential of digital agriculture in collaboration with another major tech corporation [83].

Third, I observed that some academic researchers engage with the ‘broader impacts’ funders ask for in limited ways. Stated otherwise, the broader impacts appeared to be tacked onto research visions to secure and maintain funding:

Giancarlo presented the vision for the Farm of the Future (FotF) to the [tech giant] visitors... as a site for developing and showcasing future farm technologies. This involved four phases: data aggregation, data analytics, data processing, and assessing societal impact. In the meeting, we all joked at the fact that societal impact is mentioned as a last phase. This is funny when paired with the fact that [a sociologist involved in our project] had given a talk earlier that day on the importance of assessing impact in all phases of development/deployment. - Research Diary, 02/17/2023

During system building, to assess whether we are moving towards our envisioned impact, we need to frame it in terms of specific technical milestones. For instance, one envisioned societal impact in building SDF 2.0 described earlier is to address vendor lock-in. More concretely, our platform should not be tied to a particular cloud service provider. Indeed, our team partially achieved this feat:

Ron [undergrad researcher] developed an awesome web app that replicates the tech giant’s web app... [which] has components for displaying sensor box data... statuses, ... configurations, and the notifications system. At the core of this project is the reuse of existing Azure services... (e.g. Az Tables, IoT Hub) to provide more cloud-agnostic functionality. - Research Diary, 05/20/2022

However, what counts as a sociotechnical win to us may be perceived as re-inventing the wheel and lack of ‘technical novelty’ by

reviewers in systems, since such a system already exists commercially in closed form.

Conversely, what counts as a win to reviewers may not be what we would consider a real success. For example, a workaround I can reliably use to convince peer reviewers is to have long-running system deployments and package our findings from agricultural contexts as networking insights grounded in long-term, real-world experience in the field. But keeping experimental systems running long-term often involves hacks and kludges:

In the 3/31 technical meeting, [we] described our progress made over the last four years to aggregate data from so many [sensor] vendors. What's striking about the progress is that, despite aggregating data successfully, we still find ourselves having "post-processing" scripts before the data can be fed to Machine Learning (ML) models. On a related note, the possibility of retraining ML models (which currently are using dummy models) sounds like an exciting use case in our problem definition for the SDF 2.0. - Research Diary, 04/03/2023

In this case, our implicit argument for validity is four years of use, but this use has been contextually disembedded to create the impression it runs seamlessly. Conversely, the actual lack of seamless operation can motivate a new funding proposal to address this technical limitation. In other words, the same research endeavor - cow disease detection using networked data and ML models - can masquerade as both insightful to a networking audience and forward-looking to a funding review panel. These aspects are symbiotic; if we successfully sell the idea to a networking audience to get published, then we can use the same finding and ensuing publication as evidence that retraining ML models for cow disease is an important issue deserving of a multi-million dollar grant to fund SDF 2.0.

Ultimately, envisioning and system building seem to center around maintaining the illusion of a single seamlessly operating system which is narrated in fundamentally different ways to different audiences. In the above research diary entry, for example, the agricultural context (i.e., identity of the sensor vendors, agricultural feature importance for training ML models) does not matter to systems peer reviewers. Rather, of interest to reviewers are the networking insights from the data aggregation and limitations of the "post-processing" scripts. In contrast, the technical contents (i.e., "post-processing" scripts, ML model training and deployments) do not necessarily matter to the funding agency to whom we are simultaneously pitching the next stage of the project. Rather, the four years of experience running apparently seamless technical systems for agriculture is proof enough that the seemingly performative societal impacts in the research proposals can be achieved with the funding. The key observation is that maintaining this illusion of seamlessness is precisely what simultaneously complicates our work toward the envisioned impacts and shapes what could be its eventual impacts.

5 DISCUSSION

In this section, we analyze the intricate relationship between how we as systems researchers orient to societal impact and the technical, institutional, and disciplinary aspects that enable, shape, and constrain how we can do so. We do so by analyzing the central role of navigating technical, institutional, and disciplinary seams in systems research, and the ways this shapes how societal impact is imagined and worked toward.

In prior work integrating ethnographic and technical approaches to networking research [83], we found that systems researchers' visions of seamlessness for the eventual use of our technology contrasts sharply with the mundane everyday work of networking research, which requires navigating numerous material resistances and challenges in order to construct and maintain apparently stable test deployments. While some of the infrastructural seams we as researchers encounter do eventually disappear in the development process, others are forgotten or blinkered from view because of institutional and disciplinary factors. This means that our claims to seamlessness for such technologies are constructed over a reality of recalcitrant seams which may resurface when the technology is adopted in end-use contexts.

In this work we orient to seams at a higher level, considering them as metaphorical descriptions of misalignments between technical, disciplinary, and institutional systems that have to be navigated in order to successfully produce networking research. Working across technical, disciplinary, and institutional seams is a ubiquitous aspect of networking research, especially when working in other domains where the technology to be built is envisioned to create societal impact. In this section, we will argue that this need to adapt and work across seams continually reworks representations of the technical systems being built and how their societal impact is envisioned. Structural seams are marked by differences such as disparate publication requirements, conflicting motivations, and academic power differences. These seams are brought into alignment to produce temporarily stable images of a technology and its societal impact that can motivate stakeholders and be narrated as successful work deserving of continued funding. These changing representations therefore continuously rework how societal impact is oriented to in response to structural pressures.

In this section, we will characterize the nature of work across structural seams in domain-focused systems research, describe how the mechanisms to work across seams continually produce new representations of technical systems and their envisioned impact, demonstrate how this process torques the way societal impact is oriented to in computer systems research, and describe the role of critical technical practice in navigating the dilemmas that result.

5.1 Seam Work

In this section, we draw on STS work on how people navigate seamful infrastructure [96, 103] to characterize how networking researchers navigate institutional, disciplinary, and organizational seams while building systems oriented towards societal impact in a particular domain. We build particularly on Vertesi's [103] emphasis on the agency people have to work around, with, and through the infrastructural seams they are confronted with, and her articulation of the creative work required to bring messy, overlapping

systems into temporary alignment, which she aptly describes as “patchwork”.

We frame researchers’ efforts to align action across technical, institutional, and disciplinary seams as a form of **seam work**, by which we mean *artful work across structural seams to produce products that can be narrated or experienced as “seamless”*. Our results demonstrate that seam work is ubiquitous when doing domain-focused academic networking research. Most obviously, we must navigate *technical seams* within the emerging infrastructure being built, such as technical incompatibilities between different packages that need to be combined to produce the envisioned software. But we must also navigate many *seams between organizations*. When we set visions of societal impact and start working toward them, we continuously probe whether they are in alignment with university (e.g., student graduation), funding (e.g., NSF/USDA grants), corporate (i.e., our industry partners), and agricultural (e.g., vineyards and dairy) partner priorities. For example, while we may define success as a system that enables interconnection among many different platforms, an industry partner may define success of the project as a public demonstration of what their technical platform in particular can support. In order to continue to access the resources we need to move forward with our project, we need to show each set of stakeholders progress in ways that matter to them.

Because we are doing research in collaboration with academic researchers in other domains, we must also navigate *disciplinary seams* between what our discipline considers to be “CS research” and the concerns and interests of domain researchers. This requires us to actively translate between whatever our partners want us to do and research questions that can be made interesting to the networking/systems community. Because this community defines publication-worthiness as technical novelty independent of a particular domain, our work to build functioning deployments with our project partners must be renarrated before publication as producing such technical novelty.

We must also navigate *seams within the academic hierarchy*. Our interdisciplinary project collaborations, like many in our area, are driven by collaborations between PhD students. This makes us answerable to multiple PI’s with different disciplinary orientations and the power to veto ideas that are unfamiliar in their area, as when Gloire’s write-up of the goals of the team’s work became caught up in a conflicted accountability with Hakim, Laura, and Phoebe.

In each of these cases, these seams are quite visible to us as actors; we cannot move forward with our research successfully without solving the problems produced by working across each of these seams. When we say this seam work is oriented to producing seamlessness, then, what we mean is that the goal of seam work is to remove these seams as a visible problem. Success means rendering them, if not invisible, then at least unremarkable. We find ways, for example, to cobble together incompatible software into apparently smoothly functioning systems; to align with our industry partners into jointly articulated projects that smooth over differences in outlook and rewards; to produce deployments that knit together our own concerns and interests with those of domain researchers; to produce research papers that present coherent representations of our work to differing disciplinary audiences; and to adjudicate

which views can and can’t be included in how we represent our project. This takes a lot of artful work!

In practice, seam work involves not simply working across a single seam, but navigating multiple seams at once. For example, as we build an infrastructure, we run into many material resistances, such as when a particular operating system does not work well with a particular machine-learning package. In navigating these issues and pivoting where needed over the course of years, we are continuously taking into account the many seams across which we are working. We consider institutional seams: if we make a technical pivot to manage this issue, is our work still in alignment with our institutional partners and whatever vision they had of societal impact? We consider disciplinary seams: what about our reviewers, what story can we craft that is ‘nerdy’ and makes it seem like this system has been running for a long time and can provide ‘networking insights’? Could we be painting ourselves into a corner of being too applied? At the same time, we need to maintain our plant pathology and animal science collaborations so that we have concrete case studies that we can use to motivate our system design.

5.2 Seam Work Produces Simulacra of Impact

Seam work, then, involves translation work, as we renarrate the system we envision building, its goals and its impact, based on pressures produced by differentials across various seams. Thus seam work produces shifting representations of what we are doing and why, that align with the needs of different stakeholders as the work develops.

One aspect of the project that becomes re-represented through seam work is the nature of the technical system that is being built. Academic systems research involves not only building concrete systems but using them to argue for a more generalized understanding of how future systems can be usefully produced. These understandings become reshaped and reworked in response to our audiences, as we pitch them to funding agencies, enroll industry partners, collaborate with domain partners, and test arguments through the peer review process. The possibilities we can argue for also become reworked as we run into technical, disciplinary, and other seams that shape how we can implement the technical deployments that provide the empirical grounding for our argumentation.

Another aspect of the project that becomes re-represented is the intended societal impact of the system. The actual societal impact our work will eventually have is, of course, not currently knowable to us. But our envisioned impact acts as an orientation point for our technical decision-making, and enables us to enroll stakeholders in our project. Still, these ideas of societal impact do not remain fixed; they must respond to ongoing contingencies that arise in our seam work. For example, when our funders’ insistence on open-source outcomes clashed with a domain collaborator’s industrial relationships, we had to compromise on our own preference for societal impact. Our understanding of what societal impact is and how to orient towards it, then, shifts repeatedly as we work across seams.

As such, *representations of the system and its intended impact are continuously being refactored as we work across different seams*. Vertesi highlights the temporary, contingent, and local nature of the patchwork involved in working together infrastructural seams

[103]. But while such contingency also characterizes our work across structural seams as networking researchers, our goal is to produce apparently stable representations of the technical system we are producing and its societal impact, even when the image we present belies the actually cobbled-together, temporarily functional nature of our work. We therefore refer to these representations as *simulacra* of the SDF. By “simulacrum” we mean a *clear and coherent representation of a technical system and its societal impact and meaning*. Seam work continually produces new simulacra as we narrate how the system should be understood and argue for its benefits to new audiences, and simultaneously to ourselves.

While these representations are often stated in the form of claims about reality, their relationship to empirical evidence is more complex. For example, our networking publications emphasize the length of time that we have been running SDF deployments, because this provides evidence of robustness. But the longevity of this system relies on so much maintenance and code revision that it is practically debatable whether it should be called the “same” system.

But simulacra are also not fake, in at least two ways. First, they provide the best view that we can currently muster of what we believe our system can be argued to do, based on factors such as empirical work, research insights, intuition, and the demands of the audience. Second, these representations become orientation points that guide ongoing activity, such as in the transformation from SDF 1.0 to SDF 2.0. They thus guide the system into being.

5.3 Reworking Impact Across the Seams

So far, we have argued that seam work is an inherent part of systems research, and that seam work produces shifting representations of a system and its intended impact. How does this matter for how we as academic systems researchers orient to societal impact in our everyday work? In this section, we will describe two examples of how seam work systematically torques the ways in which technical systems and their societal impact are represented and oriented to.

5.3.1 Pressures in bilingual research. The recent proliferation of computing systems, especially AI, into many domains is leading to a policy push in the United States to train the ‘bilinguals of the future’ [67]. These are understood as people in fields like biology, chemistry, politics, history, and linguistics who are also skilled in the techniques of modern computing that can be applied to them [67]. For example, the US National Science Foundation (NSF) has committed millions of dollars to NSF Research Traineeships (NRT) to train such bilingual scholars in various domains across the US (e.g., [27, 40, 59, 94]). In the SDF context, Gloire and his collaborator Frank are bilingual products of such NRT training at our university, bringing together computing, engineering, and plant sciences. In other words, bilingual projects are a policy approach for fostering broader impacts for computing by producing what Ribes et al. call ‘boundary-spanning figures’ [81] who are deeply familiar with both computing techniques and a specific domain.

This institutionalization of an approach to producing design for domain-specific impact produces endemic structural seams that shape how projects can be developed and narrated. Specifically, it (1) builds on interdisciplinary collaborations between CS and other disciplines which are often less powerful and resourced, and (2)

places the conceptual integration of these disciplines in the hands of some of the most junior members of the academic hierarchy.

Gloire had to navigate these tensions during the interdisciplinary paper-writing progress. As described in the findings, for every project, his collaborators write a paper describing the advancements in plant pathology while he writes a paper describing the computer systems and networking advancements. In this writing, it was difficult to find the right technical depth when describing computing advancements. For instance, to a CS audience, ‘object pickling’ is a fairly obvious concept. In contrast, the phrase ‘pickling ML models’ may present a foreign concept to a plant pathology audience. One outcome of this tug-of-war was the unintended dominance of CS parlance (e.g., abstraction, cloud, etc.) in the problem/solution framing of interdisciplinary contributions.

The second factor considers the power dynamics involved in the writing process. By writing with stakeholders across academic rungs, the findings pointed to an implicit reverence for problem/solution framing from (senior) faculty. In a context where societal impact considerations abound, critical reflections inadvertently took the backseat because we were preoccupied with producing technical content that can speak to multiple audiences. This challenge is more pronounced when junior bilingual researchers are most concerned with and/or have the most expertise on societal impact considerations. In other words, this power dynamic can effectively cripple critical reflection.

5.3.2 Publish, perish, or compromise for societal impact. Previously, we described how the HCI community often challenges computer scientists to grapple with the societal implications of their work [28, 91]. A recurring critique is that computer scientists are detached from or ambivalent to the societal impacts of their work. Our findings present a more nuanced account of societal impact considerations from within computer systems research, suggesting that societal impact considerations become structurally marginalized through the current CS peer review process. Recently, Edward Lee, an eminent figure in CS, argued that computing research has developed a ‘toxic culture of rejection’ based on a feedback loop between high submission rates, overwhelmed reviewers, low acceptance rates, and resubmission of previously rejected work [62]. We, too, faced significant challenges getting publications within computer systems, including five revisions and resubmissions of the same paper. A few of the typical reasons for rejections include ‘lack of novelty’ or results are ‘too obvious’ [62]. These reviews were signs of the need for more work across the seam between technical deployments meaningful to domain partners and the intellectual commitments of current systems research.

This “publish or perish” [62, 88] culture has two effects on societal impact considerations. First, we are caught up in gradual, persistent revisions to get publications. This leads us to continually re-represent our work and revise ongoing technical work to better align with the expectations of systems research. As Scott Shenker, an eminent figure in the networking community, argues “the publish-or-perish atmosphere inevitably leads us to alter our research so as to increase the chances of our work getting accepted” [88, p. 28]. In doing so, we easily lose sight of institutional, infrastructural, and other factors that are shaping our research agenda and, by extension, the eventual societal impact. The irony

is that we need those publications in order to achieve our intended societal impacts by making our ideas public. In other words, we will not (re)-consider the impacts until we have reached a twisted version of them, as was the case for the Chameleon paper.

Second, Gloire experienced a repeating cycle as a junior researcher jumping through hoops and dealing with negative reviews that take a mental toll. Lee argues that this “culture of rejection” can cause junior researchers to either leave the field or persist through stubbornness, eventually to haze the next generation [62]. While these factors are relevant for all systems researchers, we note they place particular pressure on researchers interested in societal impact in ways that defy the “domain-independent” norms of computing research.

This cycle effectively presents three paths for young computer scientists interested in societal impact to consider: persistence, defeat, or compromise. The first path is a self-repeating cycle, where Gloire persisted through five submissions, and could therefore be likely to perpetuate the cycle by subjecting other younger researchers to the same ‘hazing’ when he joins PCs. The second path is defeat, where researchers working on societal impact leave the field altogether. The third path is compromise, where papers gradually lose their societal impact considerations on the way to publication. In rare cases, the papers may find relevance in sister communities, as observed in the findings where a sub-community of researchers interested in networking, HCI, and social science emerged. In the case of the Chameleon paper, we relied on a combination of persistence and compromise. This path of compromise is explored further below, as we consider the value of critical technical practice as a means for navigating these issues.

5.4 Navigating Compromised Agency in Critical Technical Practice

We do not claim critical technical practice to be a mainstream practice in CS research. Rather, we leverage it here to show how societal impacts are envisioned, worked toward, and rehashed within computer systems research. Building new infrastructure involves hours, months, or even years of charting technical goals and working in mundane ways across technical, institutional, and disciplinary seams. Critical technical practice offers us an important guide as we execute on and document life in the technical trenches. This journey in the trenches presents many opportunities for reflections on core values. Compared to values in design (VID) engagements that bring together teams of computer scientists and social scientists during the early stages of design, critical technical practice during infrastructure building presents a unique opportunity to consider alternatives as close to the moment of intervention as possible.

Moreover, critical technical practice offers more agency to pivot technical goals and modify the system as we navigate ‘deep’ infrastructures [37], an agency which social scientists in VID efforts typically lack. For instance, we faced a lack of electrical plugs for a gateway device in an SDF 1.0 greenhouse deployment; we pivoted to find new deployment locations in the greenhouse, but we also considered the role of unreliable power in SDF 2.0’s system design. In other words, critical technical practice allowed us to recognize and build in the need for more seamful infrastructure designs.

Through reflections produced through critical technical practice, we have demonstrated how systems research intended to produce positive societal impact is constrained and shaped by the need to navigate technical, disciplinary, and institutional seams. But what implications does this have for critical technical practice itself as a methodology to support considerations of societal impact?

In describing their own work producing critically reflective scientific research into ocean plastic pollution, Liboiron argues that all agency is fundamentally compromised by one’s position within social systems that are characterized by asymmetrical power relations [63]. In other words, there is no pure or innocent position from which to rework scientific practice and assure that it will be deployed in service of the good. Instead, recognition of and grappling with compromise must form the basis for critically oriented scientific work.

Similarly, we emphasize that critical technical practice as we have deployed it is not independent of the technical, disciplinary, and institutional pressures that we describe in this paper. Rather, it introduces *yet another seam*, one across critical and technical work, which has to be juggled in the process. It also adds new disciplinary pressures, embodied in this case through Phoebe’s interventions, to produce knowledge which is publishable to the HCI and CSCW audience; you have your own standards for novelty and empirical work to which this paper must hew.

What we would argue is that critical technical practice is of value to computational researchers who care about societal impact, not by enabling an escape from seam work, but by providing a means for intentionally navigating it. Critical technical practice enables us to anticipate structural factors that provide resistance to working towards envisioned societal impact goals and circumvent them, wherever possible. In other terms, critical technical practice entails reflectively navigating powerful institutions, mundane infrastructure issues, and compromises with disciplinary structures to make the best decision possible under the circumstances.

For example, our findings captured a research process that includes setting visions, building large sociotechnical infrastructures, writing for interdisciplinary audiences, and receiving feedback from the research community. An important aspect of the envisioning process is enrolling corporate / institutional stakeholders. Critical technical practice gave us a lens to continuously and critically probe corporate interests and power in these societal impact research engagements. We worked in our project to leverage their institutional power to work towards positive societal outcomes, while considering ways in which their perspectives might skew them. No doubt the results are, however, imperfect.

In summary, using critical technical practice to work toward societal impact produces a greater degree of agency to make technical change and to leverage institutional power towards our goals, but it also demands an awareness of how that agency is inevitably compromised. Critical technical practice helps us to look behind the simulacra and to recognize the illusion of seamlessness that our activities produce. It helps us to recognize our own seam work and to navigate the choices we make of how to bring seams into alignment.

6 IMPLICATIONS FOR THE HCI AND SYSTEMS COMMUNITIES

If we understand technical, institutional, and disciplinary factors shaping computer systems research in terms of structural seams, identify the role of seam work and simulacra in reworking societal impact, and consider critical technical practice as a means to see and reflect on seam work in the course of research development, then what does this imply about how we should pursue technical research for societal impact in HCI and systems research? Here, we provide three concrete suggestions. In particular, our suggestions are focused on academic CS research, where the combination of fundamental research, domain applications, and critical reflections is most prevalent.

6.1 Design New Structural Systems for Societal Impact

Prior work in HCI has argued that CS researchers do not actively anticipate nor engage with societal impact considerations [28, 91]. In contrast, our findings indicate that computer systems researchers and their collaborators do in fact consider social implications of their work. Granted, we observed some researchers care about broader implications in performative manners to secure and maintain funding. But even when a researcher cares about societal impacts in deeper ways, our work demonstrates that numerous technical, institutional and disciplinary factors make it difficult to orient consistently towards the envisioned societal impact.

This indicates that prioritizing societal implications is not only a matter of encouraging CS researchers to care about these impacts and design their technical systems differently, but also a matter of designing different institutional and disciplinary systems to structure that work. Jackson, Gillespie, and Payette argue for increased attention to policy in analyses of social computing, providing examples where design, practice, and policy of computing systems are fundamentally intertwined in ways that shape the societal outcomes of computing systems [55]. Here, too, we argue that societal impacts of computing systems research must not be treated as independent from the technical, institutional, and disciplinary structures that enable, shape, and constrain it. As a simple example, HCI researchers could help systems researchers develop new reviewing mechanisms for systems publications which better support novelty arguments rooted in domain-dependent research; conversely, systems researchers could help HCI researchers devise publication mechanisms that better support technically-inflected societal impact research (see also [100]). More generally, this suggests that the challenge of designing for societal impact should not be understood simply in terms of producing and testing specific systems with a positive societal outcome, but can be oriented to more broadly; we provide two examples in what follows.

6.2 Use Critical Technical Practice to Investigate Infrastructure

We recognize that because of the range of skills and institutional access critical technical practice requires, it is unlikely to become a mainstream practice in HCI. Nevertheless, there is one area where

the systematic deployment of critical technical practice would clearly provide great benefit: in the design of infrastructure.

We draw here on Edwards, Newman, and Poole’s identification of what they term the “infrastructure problem in HCI,” which they explain as the field’s limited engagement with technical (software) infrastructures. These infrastructures include applications and user interfaces (the typical playing field of HCI), middleware (e.g., software libraries and frameworks), and what they term ‘deep’ infrastructures (e.g., operating systems (OS), networking protocols, etc.). This limited engagement with lower-level infrastructures leads to problems at the end user level. Our findings revealed a concrete example of this phenomenon where middleware (e.g., Azure ML packages for Python) and ‘deep’ infrastructures (e.g., 32bit vs 64 OS on Raspberry Pi 4B) ultimately constrained which applications we could build in the SDF.

Our work suggests that the infrastructure problem in HCI is not limited to producing user challenges; it also produces societal impact issues. Each decision in the hardware/software stack, however small, molds societal impact by dictating what is possible in terms of infrastructure building, shaping what can happen at higher layers. We found in our work that the (software) layers between sensor hardware vendors and cloud providers provide important abstractions, but their implementations still pose constraints. Concretely, opting for the Raspberry Pi 4B as the small-board computer to operate in a farm has consequential impacts on which cloud providers we can use to collect and process farm data. Indeed, in line with the universal plug-and-play (uPnP) initiative [37], one of the SDF’s design goals is to decouple all the system parts so that we can plug-and-play any components of the system without regard to vendor. However, each layer of the decision stack placed its own constraints on this goal.

Note that, despite the influence of these mundane choices in infrastructure building, most of these choices do not appear in the higher-level findings of technical publications because they are considered irrelevant to the abstract intellectual contributions that the system is intended to make to systems research. This is why critical technical practice has such an important role to play in identifying and altering potential societal impact in academic research. We therefore argue that critical technical practice, with continuous documentation of and later reflection on the infrastructure building practice, is an essential practice for technical researchers designing for societal impact and a core interest for HCI. We acknowledge that this reflective process requires intense efforts and cross-disciplinary training not available to all. Nevertheless, we posit that critical technical practice is an important tool in highlighting how infrastructural factors shape societal impact considerations in CS research, in both explicit and subtle ways.

6.3 Support the Trilinguals of the Future

In the discussion section, we described how policy approaches to improving the broader impacts of computing focus on the development of bilinguals of the future. One outcome of being trained across domains is that it produces ways of knowing and doing that are not only innovative, but also inherently different. By virtue of working across fields, bilinguals can be marginalized in both communities. We can become, as Passi evocatively puts it, “remainders

of unresolved equations” [79] — applauded for venturing and doing something different, but also repeatedly questioned for not ‘getting it’ in one or both fields. In dealing with the powers of anonymous reviewers, then, bilinguals face double the imposter syndrome.

In this work, we have shown that there is a trade-off between the resulting longing for membership in multiple communities and remaining critically focused on desired societal impacts. We observed instances where critical reflection is at risk of being sidelined in favor of publication numbers intended, ironically, to gain the power to represent societal impact goals within the community. We therefore argue that improving societal impact of technical systems-building will require not just bilinguals but *trilinguals* - individuals experienced in computing technology, domain understanding, and critical reflection on how to bring them together. HCI and CSCW can play an important role here as a potentially hospitable environment for trilinguals to explore and get feedback on the critical dimensions of their work. To do so effectively requires care in judging unconventional work in HCI, too, to avoid reproducing the effects of the culture of rejection that characterizes CS reviewing.

7 CONCLUSION

In this work, we present a detailed, sociotechnical account of how societal implications are oriented and worked toward in computer systems and networking research, departing from the notion in prior work that computer scientists do not care for nor actively engage with societal implications of their work. To do so, we employ critical technical practice, understood as a unified practice of technical system development and critical (self)-reflection on technical decision-making. Our findings demonstrate how envisioned societal impacts are reshaped in the practices of setting technical visions, building technical infrastructures, and writing and receiving peer feedback within CS and other domains.

The key argument we put forward is that even when computer systems researchers envision and work toward positive societal impacts with various collaborators, their work in the different research stages entails what we term ‘seam work’ between organizations, disciplines, and within the academic hierarchy. This seam work reworks envisioned impacts and representations of the systems to produce shifting simulacra of impact at the different stages. The circus juggling act, then, is to maintain an illusion of seamlessness that (1) keeps our institutions/collaborators happy and their technical support flowing; (2) relies on simulacra of systems that serve as proof-of-concepts that are attractive to our agricultural and funding partners to keep them engaged, and (3) shows our reviewers that the systems are valuable by stripping the system ideas of contextual information until the findings masquerade as a linear, domain-independent story of us addressing important networking challenges. Equipped with a better understanding of structural factors that influence societal impacts, we provide concrete suggestions for pursuing technical research for societal impact in HCI and systems.

ACKNOWLEDGMENTS

This work was supported by the US National Science Foundation (NSF) under grants #1955125 and #DBI-2019674, the U.S. Department of Agriculture’s National Institute of Food and Agriculture

(USDA-NIFA) under grant #2023-77038-38865, a Microsoft Investigator Fellowship, and the Cornell Institute for Digital Agriculture (CIDA). Gloire Rubambiza was supported by an NSF Research Traineeship (NRT) in Digital Plant Science under grant #1922551, a Microsoft Cornell Summer Research Fellowship, and a SUNY Provost Diversity Fellowship. The authors would like to thank the anonymous reviewers for their constructive feedback during the manuscript’s preparation. The networking infrastructure for the software-defined farm (SDF) project would not have been possible without the help of numerous Information Technology professionals (especially Scott Yoest), building managers (especially Scott Albrecht), orchard managers, and carpenters at Cornell University. The SDF project also benefited from research, expertise and logistical support from Julio O. Giordano, Martin M. Perez, José F. Martínez, Shiang-Wan Chin, Braulio Dumba, Andrew Anderson, Fernando Romero Galvan and Katie Gold. The authors appreciate the support from Microsoft partners, including Ranveer Chandra, Tusher Chakraborty and Elizabeth Bruce. All opinions expressed are those of the researchers, not the NSF, USDA, Microsoft, or CIDA.

REFERENCES

- [1] Veronica Abebe, Gagik Amaryan, Marina Beshai, Ilene, Ali Ekin Gurgen, Wendy Ho, Naaji R. Hylton, Daniel Kim, Christy Lee, Carina Lewandowski, Katherine T. Miller, Lindsey A. Moore, Rachel Sylwester, Ethan Thai, Frelicia N. Tucker, Toussaint Webb, Dorothy Zhao, Haicheng Charles Zhao, and Janet Vertesi. 2022. Anti-Racist HCI: Notes on an Emerging Critical Technical Practice. In *Extended Abstracts of the 2022 CHI Conference on Human Factors in Computing Systems (CHI EA '22)*. Association for Computing Machinery, New York, NY, USA, 1–12. <https://doi.org/10.1145/3491101.3516382>
- [2] admin. 2023. About.
- [3] Philip E. Agre. 1997. *Computation and Human Experience*. Cambridge University Press, USA.
- [4] Philip E. Agre. 1997. Toward a Critical Technical Practice: Lessons Learned in Trying to Reform AI. In *Social Science, Technical Systems, and Cooperative Work: Beyond the Great Divide*. Erlbaum, Mahwah, NJ, 131–157.
- [5] Philip E. Agre and David Chapman. 1987. Pengi: An Implementation of a Theory of Activity. In *Proceedings of the Sixth National Conference on Artificial Intelligence - Volume 1 (AAAI'87)*. AAAI Press, Seattle, Washington, 268–272.
- [6] Philip E. Agre and David Chapman. 1990. What Are Plans For? *Robotics and Autonomous Systems* 6, 1-2 (June 1990), 17–34. [https://doi.org/10.1016/S0921-8890\(05\)80026-0](https://doi.org/10.1016/S0921-8890(05)80026-0)
- [7] LoRa Alliance. 2015. A Technical View of LoRa and LoRaWAN. *LoRa Alliance* (2015).
- [8] Morgan G. Ames. 2019. *The Charisma Machine: The Life, Death, and Legacy of One Laptop per Child*. MIT Press, Cambridge, MA, USA.
- [9] Leon Anderson. 2006. Analytic Autoethnography. *Journal of Contemporary Ethnography* 35, 4 (Aug. 2006), 373–395. <https://doi.org/10.1177/0891241605280449>
- [10] Paramvir Bahl, Ranveer Chandra, Thomas Moscibroda, Rohan Murty, and Matt Welsh. 2009. White Space Networking with Wi-Fi like Connectivity. In *Proceedings of the ACM SIGCOMM 2009 Conference on Data Communication (SIGCOMM '09)*. Association for Computing Machinery, New York, NY, USA, 27–38. <https://doi.org/10.1145/1592568.1592573>
- [11] Natã M. Barbosa and Monchu Chen. 2019. Rehumanized Crowdsourcing: A Labeling Framework Addressing Bias and Ethics in Machine Learning. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*. Association for Computing Machinery, New York, NY, USA, 1–12. <https://doi.org/10.1145/3290605.3300773>
- [12] Hannah Barrett and David Christian Rose. 2020. Perceptions of the Fourth Agricultural Revolution: What’s in, What’s out, and What Consequences Are Anticipated? *Sociologia Ruralis* (2020).
- [13] Oliver Bates, Kathy New, Samantha Mitchell-Finnigan, Matthew Louis Mauriello, Christian Remy, Roy Bendor, Samuel Mann, Simran Chopra, Adrian K. Clear, and Chris Preist. 2019. Towards a Responsible Innovation Agenda for HCI. In *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems (CHI EA '19)*. Association for Computing Machinery, New York, NY, USA, 1–8. <https://doi.org/10.1145/3290607.3299017>
- [14] Eric P.S. Baumer, Timothy Berrill, Sarah C. Botwinick, Jonathan L. Gonzales, Kevin Ho, Allison Kundrik, Luke Kwon, Tim LaRowe, Chanh P. Nguyen, Fredy Ramirez, Peter Schaedler, William Ulrich, Amber Wallace, Yuchen Wan, and

- Benjamin Weinfeld. 2018. What Would You Do? Design Fiction and Ethics. In *Proceedings of the 2018 ACM International Conference on Supporting Group Work (GROUP '18)*. Association for Computing Machinery, New York, NY, USA, 244–256. <https://doi.org/10.1145/3148330.3149405>
- [15] Nicola J Bidwell. 2019. Women and the Spatial Politics of Community Networks: Invisible in the Sociotechnical Imaginary of Wireless Connectivity. In *Proceedings of the 31st Australian Conference on Human-Computer-Interaction (OZCHI '19)*. Association for Computing Machinery, New York, NY, USA, 197–208. <https://doi.org/10.1145/3369457.3369474>
- [16] Nicola J. Bidwell. 2020. Women and the Sustainability of Rural Community Networks in the Global South. In *Proceedings of the 2020 International Conference on Information and Communication Technologies and Development (ICTD2020)*. Association for Computing Machinery, New York, NY, USA, Article 21. <https://doi.org/10.1145/3392561.3394649>
- [17] Nicola J. Bidwell. 2021. Rural Uncommoning: Women, Community Networks and the Enclosure of Life. *ACM Transactions on Computer-Human Interaction* 28, 3, Article 19 (July 2021). <https://doi.org/10.1145/3445793>
- [18] Mehrab Bin Morshed, Michaelanne Dye, Syed Ishtiaque Ahmed, and Neha Kumar. 2017. When the Internet Goes down in {Bangladesh}. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW '17)*. Association for Computing Machinery, New York, NY, USA, 1591–1604. <https://doi.org/10.1145/2998181.2998237>
- [19] Jenna Burrell. 2018. Thinking Relationally about Digital Inequality in Rural Regions of the US. *First Monday* 23, 6 (2018).
- [20] Matthew Chalmers and Arete Galani. 2004. Seamful Interweaving: Heterogeneity in the Theory and Design of Interactive Systems. In *Proceedings of the 5th Conference on Designing Interactive Systems: Processes, Practices, Methods, and Techniques (DIS '04)*. Association for Computing Machinery, New York, NY, USA, 243–252. <https://doi.org/10.1145/1013115.1013149>
- [21] Ranveer Chandra and Stewart Collis. 2021. Digital Agriculture for Small-Scale Producers: Challenges and Opportunities. *Commun. ACM* 64, 12 (Dec. 2021), 75–84.
- [22] EunJeong Cheon and Norman Makoto Su. 2018. Futuristic Autobiographies: Weaving Participant Narratives to Elicit Values around Robots. In *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction (HRI '18)*. ACM, New York, NY, USA, 388–397. <https://doi.org/10.1145/3171221.3171244>
- [23] Marshini Chetty, Richard Banks, A.J. Brush, Jonathan Donner, and Rebecca Grinter. 2012. You're Capped: Understanding the Effects of Bandwidth Caps on Broadband Use in the Home. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '12)*. Association for Computing Machinery, New York, NY, USA, 3021–3030. <https://doi.org/10.1145/2207676.2208714>
- [24] Marshini Chetty, Richard Banks, Richard Harper, Tim Regan, Abigail Sellen, Christos Gkantsidis, Thomas Karagiannis, and Peter Key. 2010. Who's Hogging the Bandwidth: The Consequences of Revealing the Invisible in the Home. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '10)*. Association for Computing Machinery, New York, NY, USA, 659–668. <https://doi.org/10.1145/1753326.1753423>
- [25] Marshini Chetty, David Haslem, Andrew Baird, Ugochi Ofoha, Bethany Sumner, and Rebecca Grinter. 2011. Why Is My Internet Slow? Making Network Speeds Visible. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '11)*. Association for Computing Machinery, New York, NY, USA, 1889–1898. <https://doi.org/10.1145/1978942.1979217>
- [26] Marshini Chetty, Hyojoon Kim, Srikanth Sundaresan, Sam Burnett, Nick Feamster, and W. Keith Edwards. 2015. UCap: An Internet Data Management Tool for the Home. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. Association for Computing Machinery, New York, NY, USA, 3093–3102. <https://doi.org/10.1145/2702123.2702218>
- [27] Julie Dickerson, Patrick Schnable, Theodore Heindel, and Carolyn Lawrence-Dill. 2020. NRT-DESE: P3 – Predictive Phenomics of Plants. *NSF.gov* (March 2020).
- [28] Kimberly Do, Rock Yuren Pang, Jiachen Jiang, and Katharina Reinecke. 2023. "That's Important, but...": How Computer Science Researchers Anticipate Unintended Consequences of Their Research Innovations | Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems. <https://dl-acm-org.proxy.library.cornell.edu/doi/abs/10.1145/3544548.3581347>
- [29] Paul Dourish. 2015. Not the Internet, but This Internet: How Othernets Illuminate Our Feudal Internet. In *Proceedings of the Fifth Decennial Aarhus Conference on Critical Alternatives (CA '15)*. Aarhus University Press, Aarhus N, 157–168. <https://doi.org/10.7146/aahecc.v1i1.21200>
- [30] Marisa Elena Duarte. 2017. *Network Sovereignty: Building the Internet across Indian Country*. University of Washington Press, USA.
- [31] Marisa Elena Duarte, Morgan Vigil-Hayes, Ellen Zegura, Elizabeth Belding, Ivone Masara, and Jennifer Case Nevarez. 2021. As a Squash Plant Grows: Social Textures of Sparse Internet Connectivity in Rural and Tribal Communities. *ACM Transactions on Computer-Human Interaction* 28, 3 (July 2021), 16:1–16:16. <https://doi.org/10.1145/3453862>
- [32] Michaelanne Dye, David Nemer, Neha Kumar, and Amy S. Bruckman. 2019. If It Rains, Ask Grandma to Disconnect the Nano: Maintenance & Care in Havana's StreetNet. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (Nov. 2019), 187:1–187:27. <https://doi.org/10.1145/3359289>
- [33] Michaelanne Dye, David Nemer, Josiah Mangiameli, Amy S. Bruckman, and Neha Kumar. 2018. El Paquete Semanal: The Week's Internet in Havana. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*. Association for Computing Machinery, New York, NY, USA, 1–12. <https://doi.org/10.1145/3173574.3174213>
- [34] Michaelanne Dye, David Nemer, Laura R. Pina, Nithya Sambasivan, Amy S. Bruckman, and Neha Kumar. 2017. Locating the Internet in the Parks of Havana. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17)*. Association for Computing Machinery, New York, NY, USA, 3867–3878. <https://doi.org/10.1145/3025453.3025728>
- [35] David Edgerton. 2011. *The Shock of the Old: Technology and Global History since 1900*. Profile books.
- [36] W. Keith Edwards, Rebecca E. Grinter, Ratul Mahajan, and David Wetherall. 2011. Advancing the State of Home Networking. *Communications of The ACM* 54, 6 (June 2011), 62–71. <https://doi.org/10.1145/1953122.1953143>
- [37] W. Keith Edwards, Mark W. Newman, and Erika Shehan Poole. 2010. The Infrastructure Problem in HCI. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '10)*. Association for Computing Machinery, New York, NY, USA, 423–432. <https://doi.org/10.1145/1753326.1753390>
- [38] Nick Foster, Luke Compston, and Daniel Barkho. 2006. MAIL: A Framework for Critical Technical Practice. In *Proceedings of the 18th Australia Conference on Computer-Human Interaction: Design: Activities, Artefacts and Environments (OZCHI '06)*. Association for Computing Machinery, New York, NY, USA, 205–212. <https://doi.org/10.1145/1228175.1228212>
- [39] Batya Friedman and David G. Hendry. 2019. *Value Sensitive Design: Shaping Technology with Moral Imagination*. The MIT Press. <https://doi.org/10.7551/mitpress/7585.001.0001>
- [40] Laura Galloway, Jessica Connelly, and Edmund Brodie. 2023. NRT-ROL: Interdisciplinary Studies of the Phenotype: EXPANDING Training in Research and Careers. *NSF.gov* (March 2023).
- [41] Maaz Gardezi, Bhavna Joshi, Donna M. Rizzo, Mark Ryan, Edward Prutzy, Skye Brugler, and Ali Dadkhah. 2023. Artificial Intelligence in Farming: Challenges and Opportunities for Building Trust. *Agronomy Journal* (2023).
- [42] Philip Garrison, Esther Han Beol Jang, Michael A. Lithgow, and Nicolás Andrés Pace. 2021. "The Network Is an Excuse": Hardware Maintenance Supporting Community. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW2, Article 464 (Oct. 2021). <https://doi.org/10.1145/3479608>
- [43] Cally Gatehouse and David Chatting. 2020. Inarticulate Devices: Critical Encounters with Network Technologies in Research Through Design. In *Proceedings of the 2020 ACM Designing Interactive Systems Conference (DIS '20)*. Association for Computing Machinery, New York, NY, USA, 2119–2131. <https://doi.org/10.1145/3357236.3395426>
- [44] William Gaver and Frances Gaver. 2023. Living with Light Touch: An Autoethnography of a Simple Communication Device in Long-Term Use. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*. Association for Computing Machinery, New York, NY, USA, 1–14. <https://doi.org/10.1145/3544548.3580807>
- [45] Sanjay Ghemawat, Howard Gobioff, and Shun-Tak Leung. 2003. The Google File System. In *Proceedings of the Nineteenth ACM Symposium on Operating Systems Principles*. ACM, Bolton Landing NY USA, 29–43. <https://doi.org/10.1145/945445.945450>
- [46] Inc Google. 2022. Cloud Computing Services. <https://cloud.google.com/>.
- [47] Barbara Grimpe, Mark Hartwood, and Marina Jirotko. 2014. Towards a Closer Dialogue between Policy and Practice: Responsible Design in HCI. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '14)*. Association for Computing Machinery, New York, NY, USA, 2965–2974. <https://doi.org/10.1145/2556288.2557364>
- [48] Rebecca E. Grinter, W. Keith Edwards, Marshini Chetty, Erika S. Poole, Ja-Young Sung, Jeonghwa Yang, Andy Crabtree, Peter Tolmie, Tom Rodden, Chris Greenhalgh, and Steve Benford. 2009. The Ins and Outs of Home Networking: The Case for Useful and Usable Domestic Networking. *ACM Transactions on Computer-Human Interaction* 16, 2, Article 8 (June 2009). <https://doi.org/10.1145/1534903.1534905>
- [49] David Hankerson, Andrea R. Marshall, Jennifer Booker, Houda Elmimouni, Imani Walker, and Jennifer A. Rode. 2016. Does Technology Have Race?. In *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems (CHI EA '16)*. Association for Computing Machinery, New York, NY, USA, 473–486. <https://doi.org/10.1145/2851581.2892578>
- [50] Shaddi Hasan, Mary Claire Barela, Matthew Johnson, Eric Brewer, and Kurtis Heimerl. 2019. Scaling Community Cellular Networks with Community CellularManager. In *16th \$USENIX\$ Symposium on Networked Systems Design and Implementation (NSDI 19)*. 735–750.
- [51] Shaddi Hasan, Amar Padmanabhan, Bruce Davie, Jennifer Rexford, Ulas Kozat, Hunter Gatewood, Shruti Sanadhya, Nick Yurchenko, Tariq Al-Khasib, Oriol

- Batalla, Marie Bremner, Andrei Lee, Evgeniy Makeev, Scott Moeller, Alex Rodriguez, Pravin Shelar, Karthik Subraveti, Sudarshan Kandji, Alejandro Xoconostle, Praveen Kumar Ramakrishnan, Xiaochen Tian, and Anoop Tomar. 2023. Building Flexible, {Low-Cost} Wireless Access Networks With Magma. In *20th USENIX Symposium on Networked Systems Design and Implementation (NSDI 23)*. 1667–1681.
- [52] Stephen Hilgartner. 2015. Capturing the Imaginary: Vanguards, Visions and the Synthetic Biology Revolution. In *Science and Democracy*, Clark Miller and Rob Hagendijk (Eds.). Routledge, 51–73.
- [53] Sarah Inman and David Ribes. 2019. "Beautiful Seams": Strategic Revelations and Concealments. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1–14.
- [54] Steven J Jackson, Paul N Edwards, Geoffrey C Bowker, and Cory P Knobell. 2007. Understanding Infrastructure: History, Heuristics and Cyberinfrastructure Policy. *First Monday* 12, 6 (2007).
- [55] Steven J. Jackson, Tarleton Gillespie, and Sandy Payette. 2014. The Policy Knot: Re-Integrating Policy, Practice and Design in Cscw Studies of Social Computing. In *Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing (CSCW '14)*. Association for Computing Machinery, New York, NY, USA, 588–602. <https://doi.org/10.1145/2531602.2531674>
- [56] Esther Han Beol Jang, Philip Garrison, Ronel Vincent Vistal, Maria Theresa D. Cunanan, Maria Theresa Perez, Philip Martinez, Matthew William Johnson, John Andrew Evangelista, Syed Ishtiaque Ahmed, Josephine Dionisio, Mary Claire Aguilar Barela, and Kurtis Heimerl. 2019. Trust and Technology Repair Infrastructures in the Remote Rural Philippines: Navigating Urban-Rural Seams. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (Nov. 2019), 99:1–99:25. <https://doi.org/10.1145/3359201>
- [57] 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 (June 2009), 119. <https://doi.org/10.1007/s11024-009-9124-4>
- [58] Joseph 'Jofish' Kaye and Liz Goulding. 2004. Intimate Objects. In *Proceedings of the 5th Conference on Designing Interactive Systems: Processes, Practices, Methods, and Techniques (DIS '04)*. Association for Computing Machinery, New York, NY, USA, 341–344. <https://doi.org/10.1145/1013115.1013175>
- [59] Eamonn Keogh, Erin Wilson Rankin, Christian Shelton, Daniel Jeske, and Anupama Dahanukar. 2016. NRT-DESE: NRT in Integrated Computational Entomology (NICE). *NSF.gov* (Sept. 2016).
- [60] Vera Khovanskaya, Maria Bezaitis, and Phoebe Sengers. 2016. The Case of the Strangerationist: Re-interpreting Critical Technical Practice. In *Proceedings of the 2016 ACM Conference on Designing Interactive Systems (DIS '16)*. Association for Computing Machinery, New York, NY, USA, 134–145. <https://doi.org/10.1145/2901790.2901860>
- [61] Kornelia Konrad, Harro Van Lente, Christopher Groves, and Cynthia Selin. 2016. Performing and Governing the Future in Science and Technology. In *The Handbook of Science and Technology Studies*, Ulrike Felt, Rayvon Fouché, Clark A. Miller, and Laurel Smith-Doerr (Eds.). MIT Press, 465–493.
- [62] Edward Lee. 2022. The Toxic Culture of Rejection in Computer Science – ACM SIGBED.
- [63] Max Liboiron. 2017. Compromised Agency: The Case of BabyLegs. *Engaging Science, Technology, and Society* 3 (2017), 499–527.
- [64] Joseph Lindley, Paul Coulton, and Miriam Sturdee. 2017. Implications for Adoption. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17)*. Association for Computing Machinery, New York, NY, USA, 265–277. <https://doi.org/10.1145/3025453.3025742>
- [65] Kenneth Lipartito. 2003. Picturephone and the Information Age: The Social Meaning of Failure. *Technology and Culture* 44, 1 (2003), 50–81. [jstor:25148054](https://doi.org/10.1215/00141801-2003-005)
- [66] Jen Liu and Phoebe Sengers. 2021. Legibility and the Legacy of Racialized Dispossession in Digital Agriculture. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW2 (Oct. 2021). <https://doi.org/10.1145/3479867>
- [67] Steve Lohr. 2018. M.I.T. Plans College for Artificial Intelligence, Backed by \$1 Billion. *The New York Times* (Oct. 2018).
- [68] Donald MacKenzie and Judy Wajcman. 1999. *The Social Shaping of Technology*. Open University, Buckingham, UK.
- [69] Michael A. Madaio, Luke Stark, Jennifer Wortman Vaughan, and Hanna Wallach. 2020. Co-Designing Checklists to Understand Organizational Challenges and Opportunities around Fairness in AI. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (CHI '20)*. Association for Computing Machinery, New York, NY, USA, 1–14. <https://doi.org/10.1145/3313831.3376445>
- [70] Garance Marechal. 2010. Autoethnography. In *Encyclopedia of Case Study Research*, Albert J. Mills, Gabrielle Durepos, and Elden Weibe (Eds.). Sage, Thousand Oaks, CA, 43–45.
- [71] Roberta M. Melvin and Andrea Bunt. 2012. Designed for Work, but Not from Here: Rural and Remote Perspectives on Networked Technology. In *Proceedings of the Designing Interactive Systems Conference (DIS '12)*. ACM, New York, NY, USA, 176–185. <https://doi.org/10.1145/2317956.2317984>
- [72] Microsoft. 2021. What Is an Azure Machine Learning Workspace. <https://www.twilio.com/docs/sms/quickstart/pytho://docs.microsoft.com/en-us/azure/machine-learning/concept-workspace>.
- [73] Lisa P. Nathan, Batya Friedman, Predrag Klasnja, Shaun K. Kane, and Jessica K. Miller. 2008. Envisioning Systemic Effects on Persons and Society throughout Interactive System Design. In *Proceedings of the 7th ACM Conference on Designing Interactive Systems (DIS '08)*. Association for Computing Machinery, New York, NY, USA, 1–10. <https://doi.org/10.1145/1394445.1394446>
- [74] Carman Neustaedter and Phoebe Sengers. 2012. Autobiographical Design in HCI Research: Designing and Learning through Use-It-Yourself. In *Proceedings of the Designing Interactive Systems Conference (DIS '12)*. Association for Computing Machinery, New York, NY, USA, 514–523. <https://doi.org/10.1145/2317956.2318034>
- [75] Helen Nissenbaum. 1998. The Cutting Edge: Values in the Design of Computer Systems. *SIGCAS Comput. Soc.* 28, 1 (March 1998), 38–39. <https://doi.org/10.1145/277351.277359>
- [76] Helen Nissenbaum. 2001. How Computer Systems Embody Values. *Computer* 34, 3 (2001), 120–119.
- [77] Giovanna Nunes Vilaza, Kevin Doherty, Darragh McCashin, David Coyle, Jakob Bardram, and Marguerite Barry. 2022. A Scoping Review of Ethics Across SIGCHI. In *Proceedings of the 2022 ACM Designing Interactive Systems Conference (DIS '22)*. Association for Computing Machinery, New York, NY, USA, 137–154. <https://doi.org/10.1145/3532106.3533511>
- [78] Justice Opara-Martins, Reza Sahandi, and Feng Tian. 2016. Critical Analysis of Vendor Lock-in and Its Impact on Cloud Computing Migration: A Business Perspective. *Journal of Cloud Computing* 5, 1 (April 2016), 4. <https://doi.org/10.1186/s13677-016-0054-z>
- [79] Samir Passi. 2015. The Remainder of an Unbalanced Equation: Moving from Science & Technology Studies to Information Science. *Critical Alternatives 2015 Workshop on Shifting Borderlands of Technoscience* (Oct. 2015).
- [80] Matt Ratto. 2011. Critical Making: Conceptual and Material Studies in Technology and Social Life. *The Information Society* 27, 4 (July 2011), 252–260. <https://doi.org/10.1080/01972243.2011.583819>
- [81] David Ribes, Andrew S Hoffman, Steven C Slota, and Geoffrey C. Bowker. 2019. The Logic of Domains. <https://journals.sagepub.com/doi/epub/10.1177/0306312719849709>. <https://doi.org/10.1177/0306312719849709>
- [82] S. Roberts, P. Garnett, and R. Chandra. 2015. Connecting Africa Using the TV White Spaces: From Research to Real World Deployments. In *The 21st IEEE International Workshop on Local and Metropolitan Area Networks*. IEEE, 1–6. <https://doi.org/10.1109/LANMAN.2015.7114729>
- [83] Gloire Rubambiza, Phoebe Sengers, and Hakim Weatherspoon. 2022. Seamless Visions, Seamliness Realities: Anticipating Rural Infrastructural Fragility in Early Design of Digital Agriculture. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems (CHI '22)*. Association for Computing Machinery, New York, NY, USA, 1–15. <https://doi.org/10.1145/3491102.3517579>
- [84] Abusayed Saifullah, Mahbubur Rahman, Dali Ismail, Chenyang Lu, Jie Liu, and Ranveer Chandra. 2018. Low-Power Wide-Area Network over White Spaces. *IEEE/ACM Transactions on Networking* 26, 4 (2018), 1893–1906.
- [85] Christian Sandvig. 2013. The Internet as Infrastructure. In *The Oxford Handbook of Internet Studies*, William H. Dutton (Ed.). Oxford University Press, 0. <https://doi.org/10.1093/oxfordhb/9780199589074.013.0005>
- [86] Christine Satchell. 2006. Cultural Theory: From Armchair Critic to Star Performer. In *Proceedings of the 18th Australia Conference on Computer-Human Interaction: Design: Activities, Artefacts and Environments (OZCHI '06)*. Association for Computing Machinery, New York, NY, USA, 199–204. <https://doi.org/10.1145/1228175.1228211>
- [87] Phoebe Sengers, Kirsten Boehner, Shay David, and Joseph 'Jofish' Kaye. 2005. Reflective Design. In *Proceedings of the 4th Decennial Conference on Critical Computing: Between Sense and Sensibility (CC '05)*. Association for Computing Machinery, Aarhus, Denmark, 49–58. <https://doi.org/10.1145/1094562.1094569>
- [88] Scott Shenker. 2022. Rethinking SIGCOMM's Conferences: Making Form Follow Function. *ACM SIGCOMM Computer Communication Review* 52, 4 (Dec. 2022), 26–30. <https://doi.org/10.1145/3577929.3577933>
- [89] Mark Shepherd, James A Turner, Bruce Small, and David Wheeler. 2020. Priorities for Science to Overcome Hurdles Thwarting the Full Promise of the 'Digital Agriculture' Revolution. *Journal of the Science of Food and Agriculture* 100, 14 (2020), 5083–5092. <https://doi.org/10.1002/jsfa.9346>
- [90] Katie Shilton. 2010. Technology Development with an Agenda: Interventions to Emphasize Values in Design. In *Proceedings of the 73rd ASIS&T Annual Meeting on Navigating Streams in an Information Ecosystem - Volume 47 (ASIS&T '10)*. American Society for Information Science, Silver Springs, MD, USA, Article 4, 4:1–4:10 pages.
- [91] Katie Shilton. 2018. Engaging Values despite Neutrality: Challenges and Approaches to Values Reflection during the Design of Internet Infrastructure. *Science, Technology, & Human Values* 43, 2 (2018), 247–269.
- [92] Katie Shilton. 2018. Values and Ethics in Human-Computer Interaction. *Foundations and Trends in Human-Computer Interaction* 12, 2 (July 2018), 107–171.

- <https://doi.org/10.1561/11000000073>
- [93] Katie Shilton, Jes A. Koepfler, and Kenneth R. Fleischmann. 2014. How to See Values in Social Computing: Methods for Studying Values Dimensions. In *Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing (CSCW '14)*. ACM, New York, NY, USA, 426–435. <https://doi.org/10.1145/2531602.2531625>
- [94] Shin-Han Shiu, Tammy Long, Karen Cichy, Brian O'Shea, Erich Grotewold, and Carol Buell. 2022. NRT-HDR: Intersecting Computational and Data Science to Address Grand Challenges in Plant Biology. *NSF.gov* (Sept. 2022).
- [95] Dilruba Showkat, Angela D. R. Smith, Wang Lingqing, and Alexandra To. 2023. "Who Is the Right Homeless Client?": Values in Algorithmic Homelessness Service Provision and Machine Learning Research. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*. Association for Computing Machinery, New York, NY, USA, 1–21. <https://doi.org/10.1145/3544548.3581010>
- [96] Ranjit Singh and Steven J. Jackson. 2017. From Margins to Seams: Imbrication, Inclusion, and Torque in the Aadhaar Identification Project. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17)*. ACM, New York, NY, USA, 4776–4824. <https://doi.org/10.1145/3025453.3025910>
- [97] Susan Leigh Star. 1999. The Ethnography of Infrastructure. *American behavioral scientist* 43, 3 (1999), 377–391.
- [98] Rosemary Steup, Lynn Dombrowski, and Norman Makoto Su. 2019. Feeding the World with Data: Visions of Data-Driven Farming. In *Proceedings of the 2019 on Designing Interactive Systems Conference (DIS '19)*. ACM, New York, NY, USA, 1503–1515. <https://doi.org/10.1145/3322276.3322382>
- [99] Adam Streed, Michael Kantar, Bill Tomlinson, and Barath Raghavan. 2021. How Sustainable Is the Smart Farm. In *Workshop on Computing within Limits (June 2021)*. <https://doi.org/10.21428/Bf6fb269.F2d0adaf>.
- [100] Miriam Sturdee, Joseph Lindley, Conor Linehan, Chris Elsdon, Neha Kumar, Tawanna Dillahunt, Regan Mandryk, and John Vines. 2021. Consequences, Schmonsequences! Considering the Future as Part of Publication and Peer Review in Computing Research. In *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems (CHI EA '21)*. Association for Computing Machinery, New York, NY, USA, 1–4. <https://doi.org/10.1145/3411763.3441330>
- [101] USENIX. 2022. NSDI '23 Call for Papers. <https://www.usenix.org/conference/nsdi23/call-for-papers>.
- [102] Deepak Vasisht, Zerina Kapetanovic, Jongho Won, Xinxin Jin, Ranveer Chandra, Sudipta Sinha, Ashish Kapoor, Madhusudhan Sudarshan, and Sean Stratman. 2017. FarmBeats: An IoT Platform for Data-Driven Agriculture. In *14th USENIX Symposium on Networked Systems Design and Implementation (NSDI 17)*. USENIX Association, Boston, MA, 515–529.
- [103] Janet Vertesi. 2014. Seamful Spaces: Heterogeneous Infrastructures in Interaction. *Science, Technology, & Human Values* 39, 2 (March 2014), 264–284. <https://doi.org/10.1177/0162243913516012>
- [104] M. Weiser. 1993. Hot Topics-Ubiquitous Computing. *Computer* 26, 10 (Oct. 1993), 71–72. <https://doi.org/10.1109/2.237456>
- [105] Richmond Y. Wong. 2021. Using Design Fiction Memos to Analyze UX Professionals' Values Work Practices: A Case Study Bridging Ethnographic and Design Futuring Methods. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems (CHI '21)*. Association for Computing Machinery, New York, NY, USA, 1–18. <https://doi.org/10.1145/3411764.3445709>