Title
Craft Knowledge and the Advancement of Science: The Role of Scientific Support Occupations in Shared Research Facilities

Author
Bovenberg, Danielle

Publication Date
2023

UNIVERSITY OF CALIFORNIA

Santa Barbara

Craft Knowledge and the Advancement of Science: The Role of Scientific Support Occupations in Shared Research Facilities

A dissertation submitted in partial satisfaction of the
requirements for the degree Doctor of Philosophy
in Technology Management

by

Danielle Elaine Bovenberg

Committee in charge:
Professor Stephen Barley, Co-Chair
Professor Matthew Beane, Co-Chair
Professor Mary Tripsas
Professor Gary Hansen

June 2023

The dissertation of Danielle Elaine Bovenberg is approved.

Gary Hansen

Mary Tripsas

Matthew Beane, Committee Co-Chair

Stephen Barley, Committee Co-Chair

June 2023

Craft Knowledge and the Advancement of Science:

The Role of Scientific Support Occupations in Shared Research Facilities

by

Danielle Elaine Bovenberg

To Elena Mae

ACKNOWLEDGEMENTS

It is a joy to me that, at the end of my path as a Ph.D. student, I have ended up in a different place than I would have been able to find at the start. It is a good place. I have only reached it because many people have helped me along the way.

My first debt of gratitude is to the individuals who participated in this study. The core of this thesis is formed by observations gathered from the everyday work of dozens of people working as scientific support staff. I am grateful for their trust and for the openness with which they shared what they found challenging and rewarding about their work. I am also grateful to those lab users who helped me see what made cleanroom work uniquely challenging. I am also grateful to the operational directors of the Alpha, Beta, Gamma and Delta facilities, who not only provided initial access to their facilities, but took steps to clear my path. Other study participants likewise went beyond providing simple access to their work, offering hospitality. I am grateful to the graciously offered invitations to dinner, drinks, coffee, concerts, lectures, lunches, and the occasional ramble across a beautiful campus.

My next debt of gratitude is to members of my committee. Throughout my time as a doctoral student, no one has kept me caffeinated quite like Gary Hansen has. Besides taking such a concrete concern for my well-being, Gary served as a gracious conversation partner in a methodological tradition unfamiliar to him, sharpening my approach, findings and implications along the way. I am grateful to him for boldly venturing into an ethnographic dissertation and making it better.

Mary Tripsas served as broker extraordinaire between my unfolding research and scholarly conversations about technology, innovation and strategy. She created a platform for the presentation and development of my work at conferences and made many fruitful introductions. As I developed papers and talks, I benefited greatly from her direct and persistent feedback when she saw that something could be better expressed or more clearly put. My writing and speaking are all the better for it.

Matt Beane has served as a grounded, generative presence throughout my time at UCSB. More than once, talking with Matt created much-needed space for honest and relaxed reflection amid the pressures of fieldwork and writing. This space to think led to better decisions. Matt has also been a terrific source of learning about concrete skills and approaches to academic work. His dedication to learning from feedback and his laser-focused pursuit of excellence continue to be touchstones.

This research would not have made it this far without the dedicated mentorship provided by Steve Barley. Steve has been palpably *with me* from the start of my time at UCSB. Over years of weekly meetings, Steve paid attention to the most mundane matters of my trajectory. He read and provided feedback on interview protocols, fieldnotes, draft emails, memos, grant proposals and drafts of the thesis. Steve showed up to some of my most difficult moments with empathy, understanding and belief in a way forward, often supported by the warm hospitality he provided together with Debbi. The biggest reward from our collaboration is that we got to (begin to) document what members of these overlooked occupations do and why it matters.

In the TM Department, many others were sources of support, grounding and collegiality. I want to thank Giulia Brofferio for innumerable things, but chief among them for her assistance in preparing our grant proposals and administering the grants. I further thank Katie Trigg for academic support, and Amanda Higham, Cindy Colmenares, Janet Kim, Travis

Candieas, David Telleen-Lawton, Dave Adornetto, Sarah Hilliard-Rueff, Victor Dowbusz, Jamie Hu, Eli Diaz, Rosaura Aguirre, Carolyn Harrah and Kathrin Wagner for being great colleagues. In addition, I count myself fortunate to have had such a humorous and considerate group of fellow students join our department during my time at TM. My thanks go to them for their companionship and collegiality, and for their patience and good humor as they watched me give yet another practice talk: Virginia Leavell, Dylan Kreis, Roni Shen, Sienna Helena Parker, Brandon Lepine, Fares Ahmad and Sai Kalvapalle. I thank Ioana Cristea for establishing a path as project scientist before the first Ph.D. students were admitted to the program, and for being an adventurous office mate and friend. Finally, in the TM Department, I owe special thanks to the faculty members who presided over my early training: Renee Rottner, Kyle Lewis and Paul Leonardi. From elsewhere on campus, I want to thank Robby Nadler and Karen Lunsford for their expert instruction in matters of writing. Finally, I am grateful to John Mohr for sharing so generously of his time.

I also want to thank members of the broader scholarly community that I have become a part of for their input along the way. A special thanks goes to John Walsh, who read many drafts of my work and whose commitment to understanding the careers of supporting members of the scientific workforce has been a source of refreshment and focus. I further thank the following individuals for helpful comments on drafts and presentations of my work: Daniel Sholler, Andrew Nelson, Curtis Chan, Julia DiBenigno, Amy Wrzesniewski, Jim Baron, Balázs Kovács, Ivana Katic, Oriane Georgeac, Mark Zbaracki, Michel Anteby, Ruthanne Huising, Arvind Karunakaran, Pedro Monteiro, Tomi Koljonen, Rohin Borpujari, participants at the CCC, WOEPSR, TORS, Medici Summer School, Industry Studies Association annual conference, and the Open Innovation in Science Research Conference. I also want to thank my fellow board members at the ASQ Student Blog: Luke Hedden, Yun Ha Cho, Maggie Cascadden and Abraham Oshotse.

I further gratefully acknowledge support for this project by the National Science Foundation under Grant No. 2151698 and the Alfred P. Sloan Foundation under Grant No. G-2022-17174. At the Sloan Foundation, I thank Josh Greenberg, Liz Vu and David Michel for helping refine and administer our grant. I also acknowledge support from several campus fellowships at UC Santa Barbara, among them the Chancellor’s Fellowship. This work was performed in part at several member sites of the National Nanotechnology Coordinated Infrastructure (NNCI), which is supported by the National Science Foundation under grant NNCI-2025233.

In closing, I wish to thank a number of friends and family members who have played key roles. Almost from the start of my time in Santa Barbara, my friendship with Camille Endacott, my South Pasadena soul sister, has been the source of a rich and distinct form of *gezelligheid*. Esther Showalter has been a peerless companion in hiking, in exploring language, and in reflection on matters great and small. John Fernandez has kept me supplied with bike talk, trash talk related to *voetbal*, and that essential fuel for this research, pizza. Even before I traveled to the West Coast, Sarah Esselborn has been sharing in the highs and lows of life with me, and I am grateful for our time spent “surfing the waves” together. I thank Jeb Bailey for stalwart insistence that all that matters is that I “do good science,” and Michael Abraham for sharing his infectious enthusiasm about material science. Housemates have a special role to play, and I am so very grateful for the exceptional people I’ve gotten to know in the most mundane of environments: Daisy, Audrey, Bree, Jose, Sofia, Noa and Tara. Among them, I want to especially thank Audrey Lopez for showing me what it was like to complete the Ph.D. with grace and humanity, and Sofia Hamrin for biking all those miles with me. I thank Carolyn McCall for providing the stage for all these friendships. I thank Sarah Douglas for keeping me grounded, connected and aware. I thank Nikki Ramage for spending taco Tuesdays with me. Friends who supported me from across the Atlantic were: Zoe and John Bazuin, Carel Roessingh, Januschka Veldstra, Isabel van Oorschot and Fieke Stoop.

My family is fundamental to the work I have done at UCSB. Jeanne and Lans Bovenberg raised me to love the work of translation across cultures. Aboo's stories from the field fueled my curiosity in the worlds of work we inhabit. Arie and Gracia have grounded me during countless cross-Atlantic calls and nurtured me on trips home with board games and beautiful design. Little Elena's joy in learning, and delight in getting back up again after she falls just for the joy of another attempt, were an inspiration as I wrote this thesis.

VITA OF DANIELLE ELAINE BOVENBERG
June 2023

EDUCATION

2023 (expected)	Ph.D. in Technology Management. University of California, Santa Barbara. Committee: Stephen Barley, Matthew Beane (co-chairs), Gary Hansen, Mary Tripsas.
2015	MSc. in Culture, Organization and Management. Vrije Universiteit Amsterdam, The Netherlands. *Cum Laude.*
2013	BA in Liberal Arts and Sciences. Utrecht University, The Netherlands. *Magna Cum Laude.*

PROFESSIONAL EMPLOYMENT

2023 (expected) Postdoctoral Associate. Yale School of Management.

PUBLICATIONS

Rottner, R.M., Bovenberg, D. E., & Leonardi, P.M. (2019). Social Media in Open Strategy: A Five-Flows Model of Strategy Making and Enactment. In D. Seidl, R. Whittington, & G. von Krogh (Eds.), *Cambridge Handbook of Open Strategy* (pp. 186–204). Cambridge, UK: Cambridge University Press.

Roessingh, C., & Bovenberg, D.E. (2018). The Hoover Mennonites in Belize: A History of Expansion in the Shadow of Separation. *Journal of Amish and Plain Anabaptist Studies, 6*(1), 100–116.

Roessingh, C., & Bovenberg, D.E. (2016). "No Sunday Business": Navigating Religious Rules and Business Opportunities in the Shipyard Mennonite Settlement, Belize. *Journal of Amish and Plain Anabaptist Studies, 4*(2), 133–148.

WORKING PAPERS

Bovenberg, D.E. Safe Crossings: The Role of Scientific Support Occupations in Attenuating Secrecy in Competitive Scientific Communities.

- Best paper proceedings, *Academy of Management Annual Conference* (2021)
- Finalist, Louis Pondy Best Dissertation Paper Award (OMT Division of the Academy of Management, 2021)

Beane, M., Sholler, D., & Bovenberg, D.E. Solve and Be Seen: How Workers in Deskilled Jobs Advance within an Organization. *Manuscript in Preparations.*

Bovenberg, D.E. Craft Occupations and the Flow of Scientific Knowledge: Tool-oriented occupations as brokers between application-oriented fields. *Manuscript in Preparation.*

Beane, M., Sholler, D., & Bovenberg, D.E. Absorbing Opportunity: How Co-Development Efforts with Underproven Technology Curtail Organizational Learning. *Manuscript in preparation.*

INVITED TALKS AND CONFERENCE PRESENTATIONS

Bovenberg, D.E. Safe Crossings: The Role of Technical Support Occupations in Circulating Knowledge in Competitive Fields. (2022, November). Presented to the Strategy Area Group at Ivey Business School, University of Western Ontario, Canada.

Bovenberg, D.E. Safe Crossings: The Role of Technical Support Occupations in Circulating Knowledge in Competitive Fields. (2022, November). Presented to the Organizational Behavior Group at the Yale School of Management, New Haven, CT.

Bovenberg, D.E. Scientific Support Occupations and the Development and Dissemination of Scientific Knowledge. (2022, October). Presented to the Department of Management and Organizations, Questrom School of Business, Boston University, Boston, MA.

Co-organized the symposium, *A "Relating" Lens on Occupations and Professions: Collaboration, Coproduction, and Brokerage,* with Rohin Borpujari. (2022, August) Organization and Management Theory (OMT) division of the 82nd Academy of Management Meeting, Seattle, WA.

Bovenberg, D.E. (2022, August). Tool-Oriented Occupations as Brokers in Application-Oriented Fields: The case of scientific support staff in shared laboratories. Presented to the OMT division of the 82nd Academy of Management Meeting, Seattle, WA.

Bovenberg, D.E. (2022, June). Shared R&D Facilities as Brokers in Innovation Networks: How tool maintainers create and disseminate knowledge between firms. Presented to the Industry Association Annual Conference, Philadelphia, PA.

Bovenberg, D.E. (2022, June). Safe Crossings: How Laboratory Staff Circulates Knowledge Amid Scientific Secrecy in a Shared Research Facility. Presented at the Consortium for Competitiveness and Cooperation, Toronto, ON.

Bovenberg, D.E. (2022, May). Safe Crossings: How Laboratory Staff Circulates Knowledge Amid Scientific Secrecy in Core Facilities. Presented at the Open Innovation in Science Conference at CERN, Geneva, Switzerland.

Bovenberg, D.E. (2022, April). Safe Crossings: How Laboratory Staff Circulates Knowledge Amid Scientific Secrecy in Core Facilities. Presented at the Workshop on the Organisation, Economics and Policy of Scientific Research, KU Leuven, Belgium.

Bovenberg, D.E. (2021, August). Safe Crossings: The Role of Scientific Support Occupations in Attenuating Secrecy in Competitive Scientific Communities. Presented to the OMT division of the 81st Academy of Management Meeting.

Bovenberg, D.E. (2021, June). Attenuating User Secrecy in Shared R&D Facilities: The Role of Scientific Support Occupations. Presented at the Industry Studies Association Annual Conference, Boston, MA.

Bovenberg, D.E. (2020, June). Safe Crossings: How Laboratory Staff Circulates Knowledge Amid Scientific Secrecy in a Multi-User Research Facility. Presented at the 12th Medici Summer School.

Bovenberg, D.E. (2019, October). Keeping scientific equipment of different ages running in a nanotechnology laboratory. Presented to the General Maintenance track of Maintainers III: Practice, Policy and Care, Washington, D.C.

Bovenberg, D.E., Rottner, R.M., & Eberhart, R. (2019, August). Causal Claims from Observational Data: An endogenous treatment effects approach on a matched sample. Presented to the Research Methods (RM) division of the 79th Academy of Management Meeting, Boston, MA.

AWARDS, HONORS AND GRANTS

National Science Foundation, Research Grant to investigate "The Role of Scientific Support Staff in the Creation and Dissemination of Knowledge Within and Across Core Infrastructural Facilities," March 2022, $139,800. PI: Stephen Barley.

Alfred P. Sloan Foundation, Research Grant "to support research on the roles played by scientific support staff in core infrastructural facilities," February 2022, $163,566. PI: Stephen Barley.

Doctoral Student Travel Grant, UC Santa Barbara, February 2022, $1,500.

Coronavirus Aid, Relief, and Economic Security (CARES) Act Minority Serving Institution (MSI) Summer Research Grant. UC Santa Barbara. September 2021, $7,000.

Finalist, Louis Pondy Best Dissertation Paper Award, OMT Division of the Academy of Management. For "Safe Crossings: How laboratory staff circulates knowledge amid scientific secrecy." August 2021.

Best Paper Proceedings, OMT division of the Academy of Management. For "Safe Crossings: How laboratory staff circulates knowledge amid scientific secrecy." August 2021.

Conference Scholarship, TIM Division, Academy of Management 2021. August 2021, $125.

Multidisciplinary Research on the Coronavirus and its Impacts (MRCI) Summer Research Grant. UC Santa Barbara. June 2020, $2,000.

Doctoral Student Travel Grant, UC Santa Barbara. October 2019, $200.

Chancellor's Fellowship, UC Santa Barbara. Full tuition and fellowship support for 3 years of doctoral training at UCSB. 2017, 2019, and 2021.

Wilhelmina Drucker Student Talent Award, Vrije Universiteit Amsterdam. September 2016, €6,800.

Johannes van der Zouwen Master's thesis award in the Social Sciences, Vrije Universiteit Amsterdam. June 2016, €1,000.

SERVICE

2021 – 2023	Board Member. Administrative Science Quarterly Student Blog
2019 – 2021	Contributor. Administrative Science Quarterly Student Blog
2017 – on	Reviewer. Academy of Management Conference, OMT and TIM divisions

ABSTRACT

Craft Knowledge and the Advancement of Science:
The Role of Scientific Support Occupations in Shared Research Facilities

by

Danielle Elaine Bovenberg

Over the past fifty years, the U.S. government has steadily increased funding for shared research facilities known as "core facilities," with the aim of fostering discovery, innovation and national competitiveness. Due to their place at the intersection of disciplines and research programs, core facilities are often seen as places where the fruitful exchange of knowledge occurs for the benefit of science. Yet, when picturing how such exchanges of knowledge occur, most accounts overlook the occupations ideally positioned to perform them: the scientific support staff who guide researchers' use of the facility's tools. Despite the extensive knowledge that members of these occupations possess of the instrumentation and techniques at the base of diverse lines of research, little research has examined how they synthesize and diffuse knowledge essential to the progress of science.

This study examines the role of scientific support staff in the development and dissemination of knowledge in core facilities through a multi-year ethnographic study of four academic nanofabrication facilities, core facilities that make semiconductor fabrication equipment available to researchers from a range of disciplines and industries. Drawing on sociological theories of craft knowledge and structural complexity, the dissertation documents the ways in which scientific support staff routinely select, apply and combine enduring bodies of technical knowledge to match the diverse, shifting, and technically demanding objectives of different lines of research. The findings thereby show that scientific support staff do not simply help researchers operate intricate machines and use esoteric materials, but that in doing so, they enable users to attain a level of structural complexity in the products of their work that these users would otherwise struggle to achieve.

The findings contribute to research on the changing scientific workforce, the organization of scientific work, and knowledge recombination in science-based innovation. In addition, this study extends our understanding of the role that craft occupations play in their communities, by documenting how one such occupation develops, maintains and extends capabilities fundamental to the advanced achievements of a range of fields.

TABLE OF CONTENTS

“Nobody cares about the way the Ferrari was built. They care about the way it performs. But the way it was built is pretty important. Think about architectural plans for a building. The building is built, you have the architectural plans, but you don’t know how all the craftsmen and the carpenters put it together. But if you’re well-staffed like this place is, you have a much better chance of telling people how the house was built, how the car was built.”

– Experienced user and former staff member at a nanofabrication facility in this study

Introduction

Scientists increasingly rely on complex and expensive technological infrastructures to pursue their research. For example, multi-million-dollar technologies are critical for identifying novel drug compounds and for developing carbon-capturing materials (Smith & Smith, 2020; Fu et al., 2017; Stephan, 2012). The cost of such equipment routinely exceeds the budget of a single research group and is, therefore, often borne by shared facilities known as “shared research infrastructures” (U.S. Department of Energy, 2015) or “core facilities” (Meder et al., 2016; Sia & Owens, 2015; Lippens et al., 2019). An organizational form that originated in the fields of high energy physics (Traweek, 1988; Doing, 2009), astronomy (McCray, 2006) and materials science (Mody, 2017), core facilities make key instruments available to researchers, enabling them to use otherwise unaffordable technologies.

Due to their place at the intersection of disciplines and research programs, core facilities are often seen as places where the fruitful exchange of knowledge occurs for the benefit of science. These depictions typically feature researchers from different disciplines sharing knowledge with each other around common tools, thereby advancing their respective fields and sparking interdisciplinary collaborations (Lippens et al., 2019; Meder et al., 2016). However, despite the narrative appeal and impact of such chance interactions among researchers, they are unlikely to form the primary way in which core facilities foster the fruitful exchange of knowledge between the fields they support. A core facility’s unique position lies in its capacity to consolidate and disseminate technical knowledge: knowledge of what can be done with the materials, tools and techniques that different lines of research share. Circulating this kind of knowledge requires familiarity with the instruments and contact with the work of many researchers, in order to recognize where similarities and differences create an opportunity for the brokering of knowledge (Hargadon & Sutton, 1997; Burt, 2004). The researchers who use these facilities are in many respects ill-positioned to carry out such exchanges of knowledge. Their role in the facility is such that they are unlikely to develop sufficient technical comprehension of the machines they use or sufficiently broad exposure to the work of others. In addition, their careers are structured in such a way that they typically have little motivation to develop, retain and widely share knowledge about techniques or tools (Katz, 2014; Howison & Bullard, 2016; Walsh & Hong, 2003).

In this dissertation, I argue instead that the synthesis and circulation of knowledge at core facilities relies fundamentally on the knowledge that the *scientific support staff* employed at these facilities develop. Scientific support staff guide researchers in their use of a core facility's tools. As such, they work closely with the instruments and interact with many different users. Moreover, the structure of their careers predisposes them to advance the capabilities of a core facility in such a way that these benefit the research efforts of multiple users. While prior research has considered how researchers in academic and commercial settings navigate the landscape of contemporary science – with its large scale, technologically intricate, multi-stakeholder and disciplinarily diverse research projects (Star & Griesemer, 1989; Galison & Hevly, 1992; Knorr-Cetina, 1999; Teurtscher et al., 2014; Kaplan et al., 2017; Perkmann et al., 2018; Vertesi, 2020) – this body of work has largely overlooked the role that members of scientific support occupations play in this landscape, particularly as these individuals move between research projects over their career and serve multiple projects at any particular point in time. In so doing, our accounts of modern science overlook the syntheses that these occupations create from their experiences across projects, and the unique opportunities at their disposal to share what they know with researchers. My research questions therefore are:

1. What kinds of knowledge do scientific support staff in core facilities possess, and how is this knowledge distinct from that of their facilities' users?
2. How does the support staff's distinct knowledge support the dissemination of knowledge among users of their facilities?

This study examines the role of scientific support staff in the development and dissemination of knowledge in core facilities through an ethnographic study of four academic nanofabrication facilities. Such facilities provide cleanrooms and semiconductor fabrication equipment to a disciplinarily diverse, often regionally clustered community of industrial and academic researchers. These researchers use these facilities to perform primarily early-stage research and development. By providing a space where proof-of-concept research can be performed on what are typically prohibitively expensive and – within any single firm – tightly controlled instruments, nanofabrication facilities form important nodes in the United States' semiconductor R&D infrastructure and are funded as such by the National Science Foundation (National Nanotechnology Coordinated Infrastructure, 2017; National Science Foundation, 2020). These facilities operate at the crossroads of disciplines and their staff members support a variety of research endeavors, including those of startups, established firms, and academic research groups.

The study's overall aim is to identify and explore the underappreciated paths by which knowledge important to scientific advancement arises. These paths are not well understood because they involve occupations whose roles in the scientific enterprise have not been well documented. Most of what has been written about scientific support occupations goes into little detail, does not take into account their perspectives and, thereby, underestimates both the complexity of their work and its significance, especially for the advancement of science through recombination of existing knowledge (Fleming & Sorenson, 2004; Lippens et al., 2019; Meder et al., 2016). Findings from this study therefore contribute to what we know about the work and expertise of support occupations in science (Milojevic et al., 2018; Lee & Walsh, 2021). Earlier research studied technicians who worked with a single group of scientists or other professionals (Shapin, 1989; Barley & Bechky, 1994). This study extends our knowledge of the role that scientific staff fill by documenting the strategies that they employ to accumulate and synthesize knowledge across projects and then deploy that knowledge to serve multiple, diverse research ends.

Overview of chapters

In the chapters that follow, I argue that staff at nanofabrication facilities develop unique knowledge because they inhabit a distinct role in the division of labor at their facilities. Through long-term experience in this role, they repeatedly encounter similar problems and interact with similar kinds of people. In solving these problems, they develop distinct occupational knowledge (Barley, 1996). Chapter 1 serves to name this role and articulate the kind of knowledge that accrues to its occupants. To define this role, I rely on a sociological account of the work of a craftsman: Douglas Harper's account of a master metalworker named Willie in *Working Knowledge* (1987). Willie occupies a central role in a community that – due to its poverty and remoteness – cannot rely on standardized solutions for fixing structures like cars, homes, and farm equipment. As a consequence, Willie develops what Harper calls "working knowledge." I define working knowledge as a resourcefulness with a relatively stable set of materials, tools and techniques developed through long-term experience in a central role in a community that relies on these materials, tools and techniques in its diverse objects of use. The foundation of working knowledge is a fine, embodied understanding of and control over the properties of numerous materials, but extends beyond shaping these materials alone. Willie builds, fixes and improves structures that are comprised of interlocking parts, and he recognizes the specific structural elements that are needed for the whole structure to function. In so doing, he shows mastery not only over the materials common to a range of structures, but also of their integration into more complex wholes (Simon, 1962). Working knowledge, therefore, shows itself in the ability to form and fashion the elements and elemental structures distinctive to a domain of practice – be it cooking, writing or engineering – and integrate them into structures of higher integrative complexity.

I argue that support staff members at nanofabrication facilities occupy a similar role in their user communities and, therefore, develop working knowledge of building micro- and nanoscale structures. They are able to control the properties of materials at the base of nanoscale structures and manage their integration into more complex wholes. As such, they are capable of creating a wide array of specialized structures from a relatively limited set of materials, tools and techniques. The distinct organizational capability of these facilities lies in their ability to routinely fashion, recombine and integrate the elements and elemental structures involved in nanofabrication, and, thereby, support research efforts in a range of disciplines.

To document the support staff's working knowledge, the dissertation proceeds as follows. Chapter 2 provides an introduction to nanofabrication facilities, giving a brief historical account of their emergence and growth as well as a primer on the technical fundamentals of nanofabrication work. I will also introduce the specific facilities that provided the settings for this study and provide an account of how data for this study was collected and analyzed.

The empirical chapters begin by describing an aspect of nanofabrication that staff considered crucial to any user's project: the work of "process integration." Making structures at the nanoscale involved recognizing that any structure was comprised of many properties that needed to be managed. Therefore, successfully fabricating a device meant ensuring that the fabrication of any one region of the device did not adversely impact the functioning of the whole. Chapter 3 documents and discusses the breadth of the physical properties of any given device that staff helped users manage.

Chapter 4 continues the empirical exploration of staff's expertise in process integration by asking how staff members ensured the integration of users' processes in practice. In particular, the chapter seeks to describe and account for the precision of the solutions that staff devised to user problems. The precise matching of tool capabilities with the fabrication needs of a given user's structure occurred over the course of many, repeated interactions between staff and user, in which they together examined (and discovered empirically) the ramifications of different concrete approaches for the user's device objectives and the known capabilities of tools. The purpose of these interactions was to develop a fabrication approach that was both *suitable* to a given nanoscale structure and *feasible* given the capabilities of the cleanroom's set of materials, tools and techniques. I describe how such matches were ultimately found by providing accounts of several problem-solving interactions between staff and users. This chapter further helps distinguish different roles that staff played in these interactions.

Chapter 5 is about one of these roles: that of “tool owner.” Tool owners were staff members responsible for a given piece of equipment: its maintenance, basic functioning, and configuration for the precise research goals of users. Tool owners helped users fabricate features with the precise properties called for by their design. As occupants of this role, tool owners were fundamental to the facility’s capability to build complex structures: they presided over the knowledge to control the many properties of many materials from which structures were assembled. Chapter 5 documents the recurrent work tasks through which tool owners developed this expertise. Tool owners maintained, troubleshooted and repaired their tools, over years becoming familiar with their tools’ physical makeup, distinct repair histories and idiosyncrasies. Tool owners also worked with users to optimize tools for creating specific kinds of material properties. By doing this work over the course of years, tool owners gained an empirical sense of the capabilities of their tools, remembering what had been done in the past and, in the context of proposed user projects, anticipating what could be done. In addition, tool owners extended their tools’ capabilities in the context of user projects when they helped users make features that hadn’t been made before in their facility. These greater capabilities then became available to the other users that relied on a similar structural feature – often members of the same research domain, but also members of other disciplines. Chapter 5 argues that by extending their tools’ capabilities in this way, tool owners kept decades-old machines in conversation with the emerging research domains of the day. This formed an important way in which nanofabrication facilities remained responsive to emerging fields of research.

In addition, by finely managing the properties of materials through the tools they managed, tool owners synthesized and disseminated knowledge in ways typically unavailable to users. Chapter 6 draws from interview data to suggest that users struggled to successfully repurpose solutions developed in the context of one project to other projects, a struggle that likely had to do with the situated nature of the solutions that they developed together with staff. Tool owners, on the other hand, often were able to trace the root cause of their tools' performance and therefore were able to use situated problem-solving experiences to bolster their understanding of the bodies of knowledge fundamental to a given problem. They were therefore in a position to repurpose solutions developed in the context of specific projects to broader classes of problems.

Finally, Chapter 7 is a short chapter that highlights how experienced staff members thought of users' projects in terms of what they called "simple structures": concrete, buildable substructures that recurred across device types and disciplines and that were responsible for particular aspects of device functionality. Chapter 8 concludes the thesis.

Chapter 1. Literature Review

Over the past fifty years, the U.S. government has steadily increased funding for core facilities with the aim of fostering discovery, innovation and national competitiveness (Mody, 2017; Chang & Grieder, 2016; Basic Energy Sciences Advisory Committee, 2018; NSF Office of Advanced Cyberinfrastructure, 2019). Two main rationales for supporting core facilities are that they enable science to progress more rapidly by more efficiently distributing expensive resources and that they enable researchers to interact in ways that create and diffuse new knowledge. Core facilities have the potential to "foster a culture of collaborative research at the interface of disciplines," thereby "facilitating and supporting interdisciplinary strategies" (Hockberger et al., 2018, p. 91, 83) and "broad interdisciplinary integration" (Basic Energy Sciences Advisory Committee, 2019, p. 32). They may also serve as "catalyst*(s)* for new internal and external collaborations" between researchers in universities and firms (Bikovski et al., 2020, p. 2). By enabling researchers from different disciplines and organizations to share tools, core facilities "foster exchange and integration of expertise" (Meder et al., 2016, p. 1090) and become "hotspot*(s)* of expertise." (Lippens et al., 2019, p. 2). In short, their funders anticipate that core facilities become what Galison (1997) called "trading zones:" places where knowledge germinates and spreads within and between scientific communities.

Such claims rest on the idea that physical proximity triggers unanticipated and fertile interactions among scientists and engineers, a tenet long central to the management of science and technology (Allen, 1977). As Meder and her colleagues (2016, p. 1090) put it, promising ideas "often arise from chance encounters, and core facilities promote such encounters between researchers." In this account of the value of core facilities, interdisciplinary projects are born when researchers from different disciplines meet around common tools and spaces. Core facilities foster valuable exchanges within disciplines when researchers enjoy a “culture of sharing and cooperation among scientists working on similar problems” (Basic Energy Sciences Advisory Committee, 2018, p. 47).

Although such serendipitous encounters undoubtedly occur and yield benefits to the researchers involved, they are unlikely to form the predominant way that core facilities foster the exchange of knowledge among the projects that they host. Core facilities offer a diverse set of researchers access to the same materials, tools and techniques, and are therefore ideally situated to leverage similarities and differences in these technical domains for the purposes of learning and the brokering of knowledge between research projects. However, owing to their role in these facilities and the career incentives they face, for several reasons the researchers that use these facilities typically do not possess the requisite technical expertise, network structure, or motivation to consistently share knowledge about how they use tools.

First, academic researchers in contemporary universities rarely hold the technical knowledge necessary to carry out highly technology-contingent research. Instead, they rely on a growing set of specialist occupations employed at universities (Barley & Bechky, 1994; Barley, 1996; Collins, 1974; Traweek, 1988; Galison, 1997). Lack of technical familiarity with the machines they use limits the degree to which researchers may in fact recognize resemblances between technical aspects of their work and that of others, as well as opportunities for recombination between their work and that of others, a prerequisite for effective brokerage activity (Hargadon & Sutton, 1997; Burt, 2004). It is therefore probable that when knowledge sharing among users of a core facility does occur, it takes place primarily around those technical similarities that users can easily recognize – and therefore around a limited set of highly visible and relatively close technical similarities between projects. Second, because most users stay at a facility for limited periods of time, their exposure to the work of other users is likely to be circumscribed. Instruments can only be used by so many people at any one stretch of time, and users of the same tool may never meet each other. Moreover, even if they do meet, users' familiarity with the work of others remains limited to those present at the facility at roughly the same period of time. As a consequence, researchers are unlikely to develop networks necessary to broadly disseminate any knowledge they do develop about the machines among other users.

Third, researchers may also be less motivated to share what they have learned with the facility's user community. Researchers' careers incentivize them to garner rewards by authoring publications and patents based on relatively context-independent knowledge claims (Merton, 1942; Stephan, 2012); knowledge about tools and techniques is more strongly tied to a single location and does not lend itself easily to such authorship and decontextualization (Collins, 1974; Katz, 2014; Ahalt et al., 2015; Howison & Bullard, 2016). In addition, researchers often hesitate to share information about their ongoing work with peers out of fear of being "scooped" (Haas & Park, 2009; Walsh & Hong, 2003; Nelson, 2016), further reducing researchers' willingness to make technical aspects of their research accessible to others.

To the degree that scientists who use core facilities are unable to overcome these limitations of their role, proximity alone is unlikely to induce the sharing necessary to turn a core facility into a vibrant scientific and technological trading zone. It is therefore unrealistic to expect that the benefits of running a core facility for its field will result primarily from interactions among its users. We must look elsewhere.

Technical support occupations in science

Scientists who use such facilities are not the only people involved in the research that occurs there. We know from ethnographic accounts of modern science that researchers rely on a cadre of other occupations who are involved, sometimes extensively, in scientific projects (Shapin, 1989; Barley & Bechky, 1994; Collins, 1974; Traweek, 1988; Galison, 1997). For lack of a well-accepted title, I shall call members of these other occupations "scientific support staff." Depending on the setting, scientific support staff may include technicians, technologists, process scientists, facilities engineers and so on.[1] Furthermore, depending on the role the staff play, they may hold degrees ranging from high school degrees to doctorates in a variety of scientific and engineering fields.

Although support staff have maintained equipment and overseen experiments since the 17th century (Shapin, 1989), modern times have seen tremendous growth in the number of people who support scientific endeavors as well as the emergence of new scientific and technical roles (Barley, 1996; Milojevic et al., 2018). Whereas in the past scientific support staff often were generalists serving only one research group at a time (Barley & Bechky, 1994), today a growing number of support staff specialize in particular instruments or procedures and assist multiple research groups in using these tools and techniques to achieve their research aims (Gould, 2015; Lee & Walsh, 2021). This is particularly true in core facilities where operational expertise related to the equipment resides with the facilities' staff whose job it is to support and consult with researchers. In other words, in core facilities the

---

[1] The titles used by different facilities for support staff occupations vary widely.

allegiance of support staff is not aligned with a single program of research but rather to the success of the facility and the facility's many users.

Even though support staff are ubiquitous in scientific divisions of labor, scant research has focused on the roles they play or the contributions they make to the scientific enterprise. Most of what has been written about scientific support staff skims over and, thereby, underestimates both the complexity of their work and its significance, especially for the recombination of knowledge necessary for scientific advancement (Meder et al., 2016; Fleming & Sorenson, 2004). Nevertheless, three scattered bodies of evidence suggest that support staff may be integral to a core facility's ability to foster the development and dissemination of knowledge envisioned by those who fund and administer core facilities.

First, the handful of studies that have focused on the work of scientific support staff indicate that technicians, staff scientists and staff engineers not only possess knowledge that is crucial for successful research but that their knowledge is different from, yet complementary to, what the scientists they assist know. In the 1990s, Barley and his colleagues pursued a coordinated program of research on technicians in scientific, medical and engineering settings (Barley & Bechky, 1994; Barley, 1996; Zabusky & Barley, 1996; Whalley & Barley, 1997; Barley et al., 1997; Darr & Scarselleta, 2001). These researchers proposed an abstract notion of a technician as a general role that support staff play in scientific, engineering and other professional settings regardless of their training and education. Barley and Bechky (1994, p.88) found that technicians in science laboratories manage "empirical interfaces." For science technicians, managing an empirical interface involved buffering the researchers that employed them from the contingencies of contact with the aspects of the material world upon which their research rested (such as natural phenomena, instruments and organisms). In the service of this objective, technicians took care of these physical systems and translated the state of these systems into symbolic indicators that other occupations could use. As a consequence, they developed expertise that was a blend of domain-specific, abstract principles (e.g., principles of cell biology) and context-specific, situated knowledge and skills (e.g., how to recognize "happy" from "unhappy" cells or optimally align a flow cytometer). Because technicians drew from both formal, abstract knowledge and their experience interacting with materials and machines, their resulting competence was not merely a subset of the expertise of the scientists with whom they interacted. Rather the technicians' knowledge was unique and involved competencies that few scientists possessed (Barley, 1996).

Second, although the specific contributions of scientific support staff to a program of research often go unacknowledged in scientific papers, those who administer core facilities recognize that support staff are indispensable to researchers' success. Typically, such comments acknowledge the staff's importance but provide few details about what they contribute. For example, Hockberger and his colleagues (2018) write of the role staff play in core facilities:

> "Core personnel are an indispensable part of the value proposition. They provide unique skills, expertise, and experience that foster relationships and build trust and confidence in researchers as they explore new and innovative technologies and applications. … By ensuring that core directors and staff are leaders in their fields and effective communicators and partners in the research ecosystem, the institution can be confident that faculty are getting the expert advice they need." (Hockberger et al., 2018, p. 80)

However, when scientists tell retrospective, behind-the-scenes tales of what led to a scientific discovery, they are more likely to speak candidly about the specific role that scientific staff played. For instance, Brian Kobilka shared the 2012 Nobel Prize in chemistry for his work with Robert Lefkowitz on G-protein-coupled receptors. Kobilka completed part of the research at the Department of Energy’s Argonne National Laboratory using X-ray beamlines. At one point in the research, his team had difficulty growing crystals that were large and robust enough for X-ray analysis. Kobilka later recalled that “what made the difference is that the X-ray facility scientific and technical staff … dug in to help, suggesting approaches and making successive improvements to the beamline’s capabilities.” The Argonne staff modified the equipment to create “micro-beams” that used only small amounts of radiation with which the researchers could probe their delicate samples. Dr. Kobilka concluded, “We couldn’t have determined the structure of the receptor without this collaboration.” (Basic Energy Sciences Advisory Committee, 2018, p. 29).[2]

Finally, results from a pilot ethnographic study I conducted of the work of staff scientists at a nanofabrication facility suggest that staff members play a unique role in the user networks at these facilities (Bovenberg, 2021). Findings from this study suggest that even though users of the facility were frequently unable or unwilling to share knowledge regarding methods and procedures with other users, the staff were able to integrate and disseminate knowledge about tools and techniques that would otherwise have been unavailable to the facility’s larger community of users. In this way, the staff leveraged their unique role in the laboratory’s division of labor to create transfers of knowledge that would not occur if left to the users themselves.

---

[2] In this case, Kobilka and his co-authors included the Argonne team as co-authors on the scientific paper that announced the discovery.

If core facilities are indeed “scientists’ partners in achieving their research goals,” as one influential commentator put it (Meder et al., 2016, p. 1091), then our current understanding of this partnership is one-sided; that is, the story is generally told from the user’s point of view. The costs of such an approach include misrepresenting the range of expertise that modern-day science requires and failing to explain how core facilities foster the accumulation and recombination of knowledge. To develop a deeper understanding of how science benefits from shared resources in core facilities, I studied ethnographically the role of scientific support staff in the creation and dissemination of knowledge within core facilities.

Two bodies of literature help further focus an approach to the questions guiding this research. Both serve to name the role that scientific support staff at nanofabrication facilities play, and therefore the problems they routinely address. Together, this helps bring into focus the distinct knowledge they develop and possess, as well as helping us see how this role may be similar to those played by other occupations.

The first body of literature examines the role that craft knowledge plays within a community that relies on this knowledge (Harper, 1987). The second examines what it means to build successively complex structures from basic elements (Simon, 1962). Together, these literatures help us see that scientific support staff do not simply help users operate intricate machines and handle esoteric materials. In a much larger sense, in doing so they help researchers attain a structural complexity in their work that these researchers would otherwise struggle to achieve. Scientific support staff help researchers apply and combine enduring bodies of technical knowledge to match the particular objectives of a given research project. This matching takes place, I argue, through a kind of knowledge that sociologists of craft have called "working knowledge."

Working knowledge

Working knowledge can be thought of as a kind of resourcefulness with a set of materials, tools and techniques that a person develops by occupying a central role in a community that relies on these materials, tools and techniques to achieve diverse goals. The concept of "working knowledge" was developed by sociologist Douglas Harper in his ethnography of a master metalworker named Willie in the North Country of New York state in the 1980s. The author met Willie shortly after moving to the remote North Country after several neighbors recommended that he visit Willie's workshop for repairs to his ageing Saab station wagon. In the ten years that followed, Harper got to know Willie as a neighbor, and in the last five of those years began to study and document the nature of Willie's knowledge and his role in the surrounding community, by spending time with him in his workshop, by photographing him at work on projects and later talking with him about those photographs, and by observing him in interaction with the assortment of clients that would drive many miles on back roads to request his services.

Throughout the book, neighbors and customers come to Willie's shop with various requests: to have their cars fixed and to have a custom stove built, for example. Willie also travels to clients' properties to fix farm equipment. A question that runs throughout the book is how Willie can meet such a wide range of needs for his neighbors with a limited set of materials, tools and techniques. Harper addresses this question primarily by showing us what Willie's knowledgeable action, or "working knowledge" looks like in practice. We are given many beautifully situated and detailed accounts of Willie at work, describing in his own words the nature of the problems at hand and the actions required to solve them. The following section identifies a number of aspects of Willie's role that are implicit in Harper's account and that underlie the kind of knowledge Willie develops. Articulating Willie's role and attendant knowledge provides a useful set of concepts for describing the knowledge that scientific support staff at nanofabrication facilities possess and how it is distinct from that of their facilities' users.

As with most occupational knowledge, Willie's expertise takes shape through his long-term experience occupying a role in a social structure. Inhabiting this role involves aspects of the work that are "relational" – involving interactions with other people – and those that are "non-relational" – involving recurring tasks and problems (Barley, 1990b; Nadel, 1957). These two facets of a role together shape the predominant interactions, situations and problems in which a member of an occupation forms and refines their skill on the job.

The relational core of Willie's role is his central position in a community that has needs for fabrication and repair that are unstandardized and, therefore, cannot be met by prefabricated, modular solutions. The nonrelational core of his work involves employing a stable set of materials, tools and techniques to meet diverse and unstandardized fabrication needs that are situated within what, following Herbert Simon, I will call "hierarchically complex structures" (Simon, 1962). That is, these structures are composed of many simple elements that interact to form successively complex units of functionality, together forming the overall properties of the structure, be it a car, stove, or piece of farm equipment. I will therefore argue that Willie's working knowledge involves three interrelated components: (1) a fine, embodied understanding of and control over the properties of numerous materials; (2)

a central role in a community that cannot rely on standardized solutions for their fabrication and repair problems; and (3) an ability to recognize the situated fabrication needs of complex structures. After supporting each of these points, I will propose a more abstract definition of working knowledge and suggest how it relates to the role that staff at nanofabrication facilities play in the development of knowledge in science.

A fine understanding of and control over material properties

Harper describes how Willie's understanding of materials, tools, and techniques forms a common root for his diverse activities. Willie's knowledgeable action in a variety of contexts is grounded in his understanding of materials and how materials will respond to his actions and behave during their use. Harper writes,

> The basis of Willie's working knowledge is his deep understanding of many materials. It is knowing how metal, wood, plastic, or even paper and cardboard respond to attempts to alter their shape, density, or pliability. The knowledge is so detailed it leads to engineering: forming materials into machines or correcting design problems in the process of repair. Fixing and making are often very close together on the continuum of Willie's working knowledge, both grounded in a basic knowledge of materials." (p. 31)

Harper provides numerous accounts of Willie building, fixing and improving diverse physical structures. Throughout these accounts, we see how his deep knowledge of materials applies equally to categories of objects that, for those who use them, are quite distant from each other. For instance, when Willie talks about shaping car fenders out of sheets of metal, he refers to this work fundamentally in terms of "moving metal," which, as we see later in the book, makes "fender work" rely on the same set of skills used to repair a warped stove door. Both are grounded in the practice of moving metal, as Willie recounts in the excerpt below:

> "A blacksmith could take a sheet of metal and do anything with it he wants to," Willie says. "If he's a *blacksmith.* We used to take a sheet of metal and make fenders—during World War II, when you couldn't get parts. We made Cadillac fenders, Chevrolet fenders; front ones, back ones— it didn't matter. We'd just take a sheet of metal; mold it. Heat and the hammer. All you had for a mold was the look of the one fender that was sitting there. You shaped it to that." (p. 34, emphasis in original)

Harper continues to describe how Willie's "hand knowledge" of materials was gained through attempts to manipulate them in his early apprenticing as a blacksmith:

> Willie learned about metal in his father's blacksmith shop. At the forge a person comes to understand metal in a fine and detailed way, through heavy handwork, altering metal with heat and then reforming it with the hammer and the cold of water or ice. (p. 31)

These repeated experiences gave Willie a feel for different metals, their behaviors at different temperatures, and which metals will weld together, and which won't. He learned what it meant to cool a metal "too quickly", to temper it "too hard", and the different ways to adjust one's tempering if one is using a flame temper, an oil temper, or a water temper. Through repeated encounters with this set of materials, Willie came to recognize subtle shifts in the properties of metals and his own ability to influence them. At the basis of Willie's knowledge therefore lay an ability to finely influence the many properties of the materials at the center of his discipline.

A central role in a community without access to standardized solutions

Harper suggests that a community relies on working knowledge when its needs for construction and repair cannot be met by the modular solutions produced in mainstream consumer culture and its modes of mass production. In Harper's account, it is the extreme remoteness of the North Country and the poverty of its agricultural community that define its need for Willie's expertise. For instance, Willie's neighbors need to keep old cars running in the absence of easily obtainable replacement parts and the region's farmers cannot simply purchase modular components for their antiquated farm equipment. The fact that his interventions must therefore be fashioned with full regard of the particulars of the situation at hand underpins a major pillar of Willie's identity as a skilled craftsman: that he is *not* what he refers to as a "parts exchanger," a repair professional who, working within a rationalized system of production and repair, ascertains his customers' needs for modular components but lacks an understanding of the machine as a whole.

An ability to recognize the situated fabrication needs of complex structures

How Willie fashions solutions in the absence of standardized repair needs and components constitutes the third element of Willie's working knowledge. Willie's interventions are custom-made in the sense that they must be fashioned to fit into the specific structures he is called to fix. He must understand afresh what each structure needs, and he then builds the parts that it needs to function. We could say that Willie works to discover the *situated fabrication need* of every structure he encounters. He then fashions from his materials, using his tools and techniques, a *feature* that will meet that situated need. The fabrication need is situated in the sense that its properties are defined by its role in an integrated structure, whose parts interact in many subtle ways. Features built to meet situated fabrication needs must therefore satisfy multiple requirements, defined by the unique structure of the overall system in which they are to function.

Willie's ability to build such bespoke features draws on his fine control of materials. An early example in the book shows how these two forms of knowledge work together in the completing of a custom-built repair. On one occasion Willie is called to fix a piece of farm equipment known as a blower, which is a machine that uses a large, fixed propeller to move silage from a wagon into a silo. The arms of the propeller have been warped in the course of someone else's repair attempts. Willie is able to straighten the warped arms of the blower's propeller by applying heat and cold to different parts of each arm. He explains:

> "The heat is the most important thing when you're straightening. You've got to heat the right places so it will bend back where it is bent. The heat *relaxes* the metal so you can bend it. Here, I'm using an acetylene torch, but I'm not using it full power. I heat it to color, and the color I bring it to depends on the heft of the metal. There's a range of colors, from white to cherry red." (p. 37, emphasis in original)

Willie's fine control over the metal's properties is informed by his comprehension of the full set of properties the blower needs to have to be operational. He indicates one such property, the propeller's overall balance, and the work he's done on its arms ("bars", in the quote below) to achieve it:

> "Every one of those bars has to be bent exactly the same way. […] And when I finished I balanced the bars. Most people don't do that. I added weld to lighter bars until it was balanced perfectly. When it runs it sounds just like a fan – no vibration."

The fittingness of the objects that Willie creates therefore results from a combination of his structural understanding of the machine he is working on – his ability to recognize the physical features and properties that a given structure needs to function – and his ability to finely shape the materials at hand into features that exhibit these properties.

Toward a more abstract formulation of working knowledge

Abstracting from Harper's account, we can say that the person who develops working knowledge presides over a relatively unchanging set of materials, tools, and techniques that she brings to bear to build diverse structures, the particulars of which are defined largely by the recurring and shifting needs of her community. It is through these engagements with the same materials in shifting configurations that she develops her facility with the enduring properties, recurring problems, and fundamental classes of operations needed to solve problems using her tool set.

Another way to abstract from the particulars of Willie's work is to recognize the larger class of objects to which mechanical structures belong. They are systems composed of basic elements arranged in successively complex, integrated units. In the case of the blower that Willie repairs, the basic elements are physical components such as bars, welds and shafts. These elements are arranged to form mechanical subassemblies such as the ailing propeller, which is itself part of a larger assembly for feeding material into the silo. Machines, therefore, belong to a class of objects that Herbert Simon (1962) refers to as "hierarchically complex systems." A defining characteristic of such systems is that they are "composed of interrelated subsystems, each of the latter being, in turn, hierarchic in structure until we reach some lowest level of elementary subsystem." (p. 468) An emergent property that such systems have is that their numerous parts "interact in a nonsimple way." (p. 438) That is, given knowledge of the properties of individual parts and the laws governing their interactions, it is not a straightforward task to determine the properties of the whole.

Building, fixing and improving upon such systems, therefore, requires appreciating their richly integrated character. In order to properly straighten the blower's propeller arms, Willie must recognize the propeller's nested and integrated location in the machine: how the propeller sits inside its casing, which sits inside a larger assembly that includes other components and is meant to function as a unit to move silage into the silo. As Willie's attempts to balance the propeller indicate, the propeller is integrated with this larger structure in several ways: not simply as a component in propelling silage into the silo, but also in terms of the vibration it produces. Its vibration must be kept within limits so as to not inhibit the functioning of other components. Fashioning an apt part, therefore, involves recognizing the integrated nature of the larger machine and the many conditions for functionality that one part of the machine must satisfy.

Many other things that people create exhibit this kind of complex structure. Other examples of lines of work that involve working with hierarchical complex systems include cooking and baking (Fine, 1996) and producing art (Becker, 1982; Long Lingo & O'Mahony, 2010). Cooks recombine a basic set of ingredients to create many different dishes. Composers arrange notes to create many different genres of music. In both of these examples, when we enjoy a culinary dish or a piece of music, we are enjoying something that has been created from multiple discrete elements such that their interaction amounts to something more than a simple sum of parts. We typically do not recognize these as being hierarchically complex systems because they are often presented to us in ways that obscure their composite nature. We enjoy them as wholes.

The fact that individuals develop working knowledge in the context of their involvement in making, repairing and improving such systems is important for understanding the quality of the knowledge they develop. It means that individuals in these roles develop not just fine control over the elemental materials of their craft, but that they also develop knowledge of the different rudimentary substructures that can be built using these elements, and which function within successively more integrated wholes. In the example of an accomplished cook, her knowledge not only consists of the most basic elements such as "salt, fat, acid, and heat" (Nosrat, 2017) and their role in a range of dishes, but also includes familiarity with basic units of more complex functionality that reoccur across dishes, such as stock or broth, fundamental types of sauces, and different kinds of dough. These elemental structures, in turn, form the basis of more elaborate systems of flavor and texture. Working knowledge involves both control over parts and over their integration into successively complex units of functionality, and the arrangement of these units into larger units in order to achieve specific overall properties. Accordingly, working knowledge results in being able to create novel, intricate and effective recombinations of known elements and substructures.

From this vantage point, what individuals with working knowledge preside over is an ability to routinely and reliably assemble unique hierarchical systems that can exhibit highly specific emergent properties. As an extension, we might say that organizations that occupy this role in their fields – such as nanofabrication facilities – specialize in the development and maintenance of the knowledge necessary to create a range of material properties, to build commonly-used units of functionality, and to assemble them in shifting combinations for novel, precise ends. The ability to routinely match well-understood materials, tools and bodies of knowledge to such a diversity of precise objectives is working knowledge.

Contributions to scholarship

In the empirical chapters that follow, I will argue that scientific support staff at nanofabrication facilities operated in a similar role to that of Willie. As a consequence, they developed working knowledge of the materials, tools and techniques involved in building micro- and nanoscale structures. These structures were hierarchical complex structures in that they consisted of simpler elements that interacted in nonsimple ways. Building them therefore required discovery of situated fabrication needs and fine control over material properties aimed at meeting these needs. Users typically came to these facilities with defined fabrication objectives in mind, but without the knowledge necessary to achieve them, such as familiarity with the practice of nanofabrication processing and the specific capabilities of discrete tools. The main task of the staff was to propose some combination of the capabilities of the facility's tools that matched the discrete fabrication objectives of a succession of users. I will argue that working in this role, staff members developed an understanding of how to finely influence the properties of materials, of how to build common elements of functionality they referred to as "simple structures" and of how to integrate these substructures into larger systems without destroying device functionality along the way, a discipline they referred to as "process integration."

I will further argue that what we see at nanofabrication facilities is not a single person with working knowledge, but multiple members of the technical staff who form a community of practice that stewards this knowledge (Lave & Wenger, 1991; Orr, 1996; Nicolini et al., 2022). Accordingly, examining the work of staff in these facilities enables us to see working knowledge in a contemporary setting, and examining its distribution across a community of practice allows us to make explicit some of the components of working knowledge that are not made explicit when they are contained in a single person, as they are in Harper's account.

Besides serving as an occasion to bring to the surface the nature of working knowledge, examining the work of scientific support staff at core facilities also provides an occasion to examine how working knowledge facilitates the circulation and recombination of knowledge within the community that it serves. By examining the work of engineers who staff these facilities, we are positioned to study not only how working knowledge serves the immediate ends of those who make an appeal to it (the users of the facility), but also how it serves as a distinct basis upon which knowledge about what can be done with a set of materials, tools and techniques circulates among a community that relies on working knowledge.

Such circulation is interesting in the context of studies of science and innovation, where innovation has long been defined as the recombination of existing knowledge (Katila & Ahuja, 2002; Fleming & Sorenson, 2004; Basalla, 1988; Gilfillan, 1935). As Harper's account of Willie's work suggests, the classification of problems based on working knowledge differs noticeably from classifications of objects based on their use. Few if any studies of innovation consider how such occupational classification systems may shape the way knowledge is recombined (or "search" is conducted) by different occupations involved in scientific and innovative projects (Levinthal & March, 1993; Levinthal, 1997; Gavetti & Levinthal, 2000). As mentioned earlier in this discussion, this oversight is probably due to a lack of detailed studies of the work of technical occupations that support multiple lines of research. Examining the working knowledge of scientific support staff at nanofabrication facilities, thus, also offers a way to extend what we know about knowledge recombination and more fully account for the ways that novelty arises in contemporary contexts of science-based innovation.

Before proceeding to an account of the work of scientific support staff in nanofabrication facilities, the next chapter provides an account of the settings and methods used in this study.

Chapter 2. Setting and Methods

Introduction

To examine the role of scientific support occupations in the development and dissemination of knowledge in core facilities, I studied four nanofabrication facilities. Before proceeding to the analysis of the empirical material that I collected, it is important to understand some basic facts about what nanofabrication facilities are, what role they have come to play in the U.S. science and engineering landscape, and the technical fundamentals of the work that is typically performed in them. In addition, it is necessary to report how data for this project was collected and analyzed. Accordingly, this chapter has four sections. First, I will briefly situate academic nanofabrication facilities historically, detailing the circumstances of their establishment and growth. Second, I will introduce the specific facilities that formed the setting for this study. Third, I will provide a primer in the technical fundamentals of nanofabrication necessary for understanding the empirical descriptions in later chapters. Finally, I will provide an account of how data collection and analysis proceeded for this study.

Historical background: Academic nanofabrication facilities in the United States

Academic nanofabrication facilities as we know them today are large, professionally staffed "core facilities," also known as "user facilities," which consolidate semiconductor fabrication equipment and make it available to users from academia and industry in a wide range of fields. The precursors to these facilities were founded in the late 1970s and 1980s. Historian of science Cyrus Mody (2017) traces the roots of present-day facilities to the late 1970s, when the first national "center" for engineering was established.

Founding a national center for engineering research

In the 1960s and 1970s, the National Science Foundation (NSF) had begun to define a new role for itself, shifting its focus from funding only basic research in the physical and life sciences to include engineering – and from funding only individual investigators to funding research centers as well (Mody, 2017).[3] In line with this shift, the NSF supported the development of a proposal for a national center that fostered research in engineering. Program officers in the Engineering Division of the NSF developed such a proposal in 1974, and this proposal was further refined in a series of workshops in 1975. The proposal ultimately described a center for research at the smallest scale attainable at the time, a "National Research and Resource Facility for Submicron Structures" (NRRFSS – which allowed for a tongue-in-cheek pronunciation, "nerfuss").[4] In a competition between UC Berkeley, MIT/Lincoln Lab and Cornell, the contract for administrating NRRFSS was won in 1977 by Cornell University.

At Cornell, the organizational form evolved to reach what Mody calls "good-enough alignment" with the university's organizational structure and mission (p. 119 – 120). The main challenge in reaching alignment, according to the facility's founding director, involved

---

[3] Mody shows how NSF's openness to the "center" model for microfabrication was aided by its acquisition of several centers from the Department of Defense (DoD). It had acquired these centers when an amendment was passed barring the DoD from funding basic research. The centers the NSF acquired in this way in the 1970s were a dozen Materials Research Laboratories (MRLs), two astronomy observatories, and the National Magnet Laboratory at MIT. In this period, the NSF also turned the Stanford Synchrotron into a "national user facility" to make beamtime available to researchers in fields outside of high-energy physics. The NSF also funded a Biological Research Resources Program, and in chemistry, the NSF funded a number of Regional Instrumentation Facilities that rented out experimental equipment.

[4] I am barely touching upon a great deal of history in recounting the events leading up to the NSF proposal. See Mody (2017), Wolf (2007) and Mody & Choi (2013) for a more fine-grained discussion of the events that led up to the formation of such a proposal, the major meso-level players (administrators and program officers) involved, and for speculation as to the unfolding of the numerous competitions for funding that mark the further evolution of the field of university micro- and nanoscale facilities.

balancing the demands of local faculty with the center's mandate to be a "resource" to users from across the United States (Wolf, 2007). The latter mission was eventually given primary importance, and in 1987 the NSF grant was renewed, this time under the name "National Nanofabrication Facility" (NNF).

The establishment of shared cleanrooms at other U.S. universities

By the 1980s, shared cleanrooms were being established on university campuses elsewhere in the United States as well, but without direct support from the NSF for a named, national center. Notably, in 1980 a cleanroom was established at Stanford University, under the aegis of the Center for Integrated Systems (CIS), which was supported by a consortium of semiconductor firms with financial support from DARPA and NSF equipment grants. In addition, in the course of the 1980s, the University of California at Santa Barbara had attracted and fostered a diaspora of former Bell Labs engineers as faculty, whose common background and research interest in a class of materials known as compound semiconductors led these professors to centralize several pieces of key equipment into a shared cleanroom. Likewise, shared cleanrooms were formed at Pennsylvania State University and Howard University during this era. These five universities then became the joint recipients of NSF support in 1993 when the NSF, rather than exclusively renewing funding for Cornell's facility, put out a call for a "National Nanotechnology Users Network" (NNUN), initially to be funded for ten years.

Resilience to shifts in science funding and relative prominence of research fields

The adaptability of nanofabrication facilities as an organizational form within the U.S. science landscape became clear over the course of the late 1980s and the early 1990s. In this period, funding for engineering research from the Department of Defense and the NSF flatlined, while funding for biotechnology research increased sharply. Mody shows how in that period, first the NNF and then its successor, the NNUN, were successful in orienting themselves more broadly toward the increasingly well-funded and prestigious field of biotechnology. The prefix “nano,” added to the NNF’s name and kept with the formation of the NNUN, lent itself well to this shift from a predominantly microelectronic user base to a broader range of applications. This successful reorientation and the subsequent endurance of nanofabrication facilities is notable, given that other institutional forms born in that era of institutional experimentation did not survive. For instance, consortia between industry and academia that were formed during that time in the semiconductor industry, such as SEMATECH, proved less resilient, at least in their initial form (Mody, 2017).

The network grows and its membership changes

Since funding the initial network of NNUN, the NSF has twice renewed funding for a national nanofabrication infrastructure. In 2004, it funded the National Nanotechnology Infrastructure Network (NNIN) and in 2016, it funded the latest network, the National Nanofabrication Coordinated Infrastructure (NNCI). With each renewal, universities have joined and left the network. Table 1 shows the changing membership of the NSF-funded network over time.

\-------------------------
Table 1 about here
\-------------------------

Alongside the network members, academic nanofabrication facilities of differing scale and scope have continued to emerge across the U.S. university landscape, and are still being founded with regularity. Outside of the current NNCI, nanofabrication facilities with comparable numbers of users to major NNCI nodes include facilities located at the

University of Michigan, Massachusetts Institute of Technology, Carnegie Mellon University, University of California, Santa Barbara, and University of California, Berkeley.

Historical accounts overlook technical staff members of nanofabrication facilities

In the description so far, I have drawn heavily on Mody's historical account of the emergence of microfabrication facilities, which is very helpful in understanding how these facilities emerged in a period of U.S. scientific history. But Mody himself is clear on the fact that his account is aimed at understanding the role of "meso-scale actors" (p. 17) in the emergence, proliferation and persistence of nanofabrication facilities. By meso-level actors, he means individuals such as NSF program officers, administrators of micro- and nanofabrication facilities, and professors with a stake in the success of these facilities. A focus on such actors is well-suited to accounting for the emergence of these organizations by showing how actors with the power to "propose, enact, oversee and judge among institutional experiments" (Mody, 2017, p. 17) did indeed do so.

This focus also means that the role of other occupations in these developments remains less well documented and understood. In particular, there is little in these early accounts about how the technical staff employed at these facilities enabled the emergence, equilibration and persistence of the facilities. Yet it is hard to imagine, for instance, how the NRRFSS struck its novel and workable equilibrium of offering a "resource" for a widening and diverse array of visiting researchers if it were not for a set of staff members primarily tasked with building, improving, modifying, and applying a technically complex infrastructure to a wide and shifting array of research endeavors. Moreover, presumably, these facilities' robustness to change has something to do with how well its expensive – and therefore durable and ageing – tools were applied, modified and rearranged as they were brought to bear on applications at quickly shifting and advancing research frontiers. In Chapter 5, in particular, I will be discussing how the character of staff's knowledge enabled them to keep in step with changing research frontiers.

The fact that nanotechnology is often referred to as a quintessential meeting place for different disciplines (Mody, 2017) is important. Indeed, Mody refers to nanofabrication facilities as "hosts" of "ties between microfabrication and the life sciences." But his account and those of others speak primarily about ties made between researchers pursuing projects together at the micro- or nanoscale. The importance of such collaborations notwithstanding, the multidisciplinarity that these facilities have fostered must to a significant degree be rooted in the broadening and deepening of its staff's knowledge – in their pragmatic attempts to help users from diverse fields use the same tools with precision.

In short, much of the success with which these facilities are credited must have taken place to a certain extent in and through the work of its technical staff. By closely examining the work of these members of the scientific workforce, this dissertation endeavors to enable more articulate acknowledgement of the distinct knowledge that members in these occupations developed, and how their expertise enabled them to play distinct roles in their facilities. I turn now to describing the facilities that served as field sites for this study.

Research sites

I conducted the research at four academic nanofabrication facilities, which I will refer to by pseudonyms in order to protect the anonymity of participants. I will refer to the facilities in the order with which I studied them: Alpha, Beta, Gamma and Delta. Each facility provided a nexus for a regional community of researchers, serving between 600 and 1,200 users per year and between 200 and 500 users per month. They each employed between 11 and 20 technical staff members. Users were academic and corporate researchers, primarily from the fields of Electrical Engineering (microelectronics, photonics), the Life Sciences (microfluidics), and Mechanical Engineering (micro-electromechanical systems, known as MEMS). Several of the facilities were or had been members of the NNCI or its precursors.

These facilities were in many respects suitable sites for exploring the kinds of knowledge that scientific support staff at core facilities possess. Each employed a large support staff. Many staff members had had long careers in core facilities and, as a result, had accumulated extensive knowledge of nanofabrication equipment and processes. Because the staff interacted with many different users on any given day around the problems that the users experienced, I was able to observe staff creating solutions for diverse projects and disseminating knowledge during their day-to-day work.

I chose to concentrate solely on nanofabrication, rather than study multiple types of core facilities, because two decades of studying scientific and technical work has demonstrated that ethnographers must develop a sufficiently deep understanding of the knowledge and practices of an occupation in order to document adequately what its members do and say (Barley & Bailey, 2020; Bailey & Barley, 2011). Rather than dilute my efforts by learning principles, tools and language of several “instrumental communities” (Mody, 2011), my approach ensured that I could better capture what was both similar and different among support staff at similar facilities. To understand the kind of work that staff performed, a basic description of the aims, tools and procedure of nanofabrication is a prerequisite.

Nanofabrication: Aims, tools, and techniques

Nanoscale structure design

Nanofabrication leverages the relationship between a material's structure and properties to create novel and useful physical devices.[5] The functionality of these devices is defined by their physical structure and the properties of their constitutive materials. Therefore, users of nanofabrication facilities build *structures.* Like the buildings we see around us, these small structures have architectures, geometry, dimensions of height, breadth and depth, and are designed to exhibit a specific set of properties. Figure 1, a photograph of material designed by researchers to exhibit an iridescent blue color like that of a *Morpho* butterfly wing, provides an example of a nanoscale structure. Zhang and Chen (2015) observed the structure of a wing's surface at the microscale and subsequently created a structure that mimics it, composed of parallel tapered posts with lamellae. When hit with white light, this array absorbs wavelengths of light that correspond to certain colors, while reflecting other wavelengths, consequently approximating the appearance of the wing found in nature. This is one example of a structure-property relationship, namely between surface texture and the absorption of light, and hence, the appearance of color to the human eye.

Figure 1 about here

There are countless other structure-property relationships that engineers exploit in nanofabrication. A well-known example is the relationship between a material's atomic

---

[5] A nanometer is one billionth (10^{-9}) of a meter. Micrometers (or microns, the units of the microscale) are units of 1,000 nanometers. The width of an average human hair is approximately 80,000 – 100,000 nanometers (80 – 100 microns). The width of DNA is approximately 2.5 nanometers.

structure and its conductivity, which is exploited in integrated circuits and other semiconductor devices. Another example involves using nanofabrication equipment to develop tiny mechanical systems, such as gyroscopes and accelerometers, to function inside personal wearable devices. Finally, in the life sciences nanofabrication is used to construct networks of microscale channels to finely sort the genetic material in biological fluids. These examples represent physical "buildings" constructed at a small scale with the aim of producing a desired emergent effect, be it unique kinds of light diffraction, conductivity, sensing, or fluid flow. By choosing appropriate materials and controlling their structure at a very small scale, engineers can affect distinct qualitative differences in the emergent properties of a device. Table 2 lists and defines key terms used in nanofabrication.

Table 2 about here

Nanoscale device fabrication

Nanoscale structures are built using a set of operations known as "processing." Processing involves selectively applying and removing materials to or from a thin substrate, which is called a "wafer" or "sample" as it is being processed. Process steps typically comprise one of four basic operations: *patterning a wafer* using light and light-sensitive chemicals (a procedure called photolithography, or when an electron-beam is used instead of light, electron-beam lithography), *applying a thin film to the sample's surface* (called deposition), *removing selected material from the sample's surface* (called etching) or *inspecting the sample for relevant properties* (called characterization). Most of these steps take place inside table-mounted stainless- steel vacuum chambers designed to keep particle-count low and shield users from the toxic, corrosive or flammable gases, high-energy plasmas and strong currents involved in many process steps. These chambers connect (at minimum) to a user interface, wafer transfer mechanism and vacuum pumps, and so entire nanofabrication tools can range in size from a tabletop printer to a small car. Each piece of equipment typically performs one specialized step: to either pattern, deposit, etch or characterize wafers. Using an orderly sequence of these tools, cleanroom users can build nanoscale structures with specific combinations of properties.

If construction is the reigning metaphor for the design of nanoscale structures, *cooking* is the image most often used to capture the complexities of nanofabrication processing. Cleanroom staff and users compared nanofabrication to cooking in order to indicate the importance of sensorimotor competence when handling samples, a comfort with the layout of the cleanroom, and a learned familiarity with how particular tools and materials respond to combinations of parameter settings, which were referred to as “recipes” for that tool. In addition, like in cooking, developing a sequence of nanoscale processing steps often involved anticipating and managing relationships of complex causality between multiple elements involved – steps performed early on in a process could become problematic only later in the process, along the way being amplified, dampened or otherwise affected by the intervening steps. (Chapter 3 discusses the diversity of different complex causal relationships that can emerge in this way.) Developing an effective process was a hard-won achievement, particularly for building structures that had many layers. Like accomplished chefs, those who most skillfully used the cleanroom developed what some called “intuition” for appropriate material combinations as well as for discerning which parameters to tweak next in their recipes to obtain a desired result.

It was not unusual for a given “process” or “process flow” (as the full set of steps to build a structure is called) to involve hundreds of handling steps and tool recipes. While processing, users moved their sample from tool to tool, programming recipes into tools and waiting for their completion, interspersed with periods spent at wet chemistry benches working with acids, bases or solvents. Throughout, it was vital to keep the sample as clean as possible, which was accomplished by removing particles in different ways. Figure 2 shows an excerpt from a finished process, developed at one of my field sites. The semiconductor device it is designed to fabricate is constructed of five layers. Each layer involves a patterning step using photolithography (“Stepper” in the figure), a deposition step (“E-beam”), as well as multiple instances of inspection (e.g., “Reflectometer”), cleaning (“solvent clean”, “DI rinse”) and wet chemistry (“bake”, “develop”, “spin”). The total process contains about 150 steps (only about 30 are pictured). This process, once fully developed, can take about two weeks to execute.

Figure 2 about here

At any one time there were dozens of users working in the cleanrooms at my field sites. Alone or in small groups, they worked through their own process, carefully designing or refining one step or another, in order to develop a working process: a set of instructions that, if followed, would build the structures that would exhibit the properties intended in their design.

Main staff roles

The staff at the facilities I observed maintained the facility’s instruments and infrastructure, trained users to operate equipment, and assisted users in developing research strategies. The technical staff of the facilities typically divided along the lines of "equipment," "process" and "facilities," though the organizations differed in how strict these distinctions were. Equipment and process staff were specialists in the techniques of nanoscale design and fabrication and were the initial point of contact for users seeking to design a fabrication strategy, select appropriate equipment or troubleshoot problems. Facilities staff monitored and ensured the conditions of the cleanroom, maintaining and upgrading critical systems such as air handling, chilled water, and acid waste neutralization. My study focused primarily on the work of those staff members whose role led them to develop familiarity with users' processes, and so I will focus my description of staff roles on equipment and process staff.

Equipment staff

Those staff members whose role was to be a "tool owner" (Chapters 5 and 6) typically specialized in a subset of tools, such as a specific kind of deposition system, etch system, or patterning system. Depending on the facility, each tool owner was responsible for between seven and 15 large tools, and a handful of smaller, less often used tools. They also typically served as designated "backups" for the tools of other equipment staff members. Equipment staff at facilities with fewer staff members per tool had to support a larger set of tools than their counterparts at better-staffed facilities.

Process staff

Staff members who specialized in "process" typically interfaced with users regarding their device as a whole or stretches of their process that incorporated several tools. Process staff worked with users to translate the desired performance objectives of their device into a physical structure that could in fact be made in the lab (Chapters 3 and 4). Process engineers were able to span these domains, in part, because they were familiar with the working principles of different types of structures. They also knew from experience how the fabrication equipment functioned to build these structures. Bridging these two domains of design and fabrication, process engineers described their work as “helping graduate students figure out what it is they need to make,” articulating “the this-lab aspect” of a commercial project, and asking, “What does the physical device have to do? How does the physical device have to work to do that?”

Typical user trajectory

Users typically arrived at a facility with a sense of the objectives of their device, but without familiarity with the tools and techniques necessary to make them. Their work progressed in two broad stages. The first stage took place outside the cleanroom and was geared at assessing the feasibility of a structure and mapping users’ process needs. The second stage took place largely inside the cleanroom and was geared at developing a working process.

Assessing feasibility and mapping process needs

Users’ journeys through a nanofabrication facility started with one or more meetings with process and equipment staff to make an assessment of the feasibility of their project. For staff, making this assessment involved determining *the process needs of a user*; that is, translating the objectives of a device as conceptualized by a user into the physical properties of the device to be built in this cleanroom. Making this translation involved parsing the design of a user’s device in terms of the “simple structures” that it would be composed of – buildable substructures which staff had either built themselves or seen built before in this lab (Chapter 7) – and defining the material properties of the device and comparing these to the ranges of outputs their tools were capable of producing (Chapters 4 and 5). Furthermore, assessing feasibility involved anticipating challenges with *integrating* the simple structures and various material properties of the device (Chapters 3 and 4).

Based on the user's process needs, staff conveyed to users how difficult they expected building this device would be – often in terms of how much time or money they anticipated it would take to discover and specify workable fabrication steps. Staff would also identify steps in the fabrication process where they expected the most work would be required, or points at which the feasibility of the device was in question. They would also recommend alternative fabrication strategies to increase the feasibility of the device or reduce its cost, for instance by proposing experiments that would quickly resolve feasibility questions or by proposing an alternative ordering of process steps that would increase a user's control over crucial parts of the device. The outcomes of these early discussions included an understanding among staff of the physical properties of the structure to be built, and a rough approach for process development: a prioritized list of the fabrication challenges the user would need to address. After these conversations, some users chose not to proceed – for instance, because they considered the costs or odds of success unacceptable. Discussions at this first stage tended to be quite concrete and thorough, because working in any of the cleanrooms in this study was expensive and experience had taught staff members that it was better to flag early on any challenges to a project's feasibility.

Process development

After progressing past an initial feasibility assessment, users began the work of discovering, defining and refining a sequence of fabrication steps with which to build their device – an activity known as "process development." During process development, users had to suit up in coveralls known as “bunny suits” and enter the cleanroom for long stretches of time. There, their objective was to tune a series of complex tools to deliver highly precise outputs. The goal of process development was to specify a sequence of steps that, when executed, led to the construction of the intended device. This work occurred in close collaboration with tool owners (Chapter 5 and 6). During process development, as users encountered problems or sought assistance in interpreting the results of their processing, they contacted staff members while in the cleanroom, arranged one-on-one meetings, or attended the open “process meetings” that these facilities typically organized: weekly open meetings in which users could reserve 30 minutes to an hour to discuss process issues with a group of staff assembled in a conference room or office.

During process development, the work of mapping process needs and assessing feasibility continued, but now with empirical data to inform these determinations. The user continued to develop their process until they had built a device they could use, determined that building the device was not feasible, or stopped using the lab for other reasons, such as lack of funding or time. Staff typically had little involvement in the final stages of a user’s project, because as users progressed, they less frequently required staff assistance. Moreover, in contrast to their entry into the lab, exit from the lab did not require users to meet with staff. Once they left the cleanroom, users typically took their devices to their own labs for further testing or use.

Nanoscale devices as hierarchical complex systems

In Chapter 1 I argued that staff developed their distinct expertise because they occupied a central role in a community whose needs for construction, repair and improvement were nonstandard and concerned structures that exhibit what Herbert Simon called “hierarchical complexity.” As I described in the section of this chapter titled “Nanofabrication: aims, tools and techniques,” nanoscale devices were structures. At their base was fine control over materials at the nanoscale, obtained while operating tools and executing techniques in the cleanroom. During process development, these material properties were then arranged into successively higher levels of integration: first, into simple substructures responsible for aspects of device functionality (Chapter 7) and then into higher levels of integration (Chapters 3 and 4). For instance, in the example of the *Morpho* butterfly wing given earlier in this chapter (Figure 1, not fabricated in any of the cleanrooms in this study), the final structure is an array of nanoscale features that resembled little trees: a vertical “trunk” from which horizontal “branches” extended. To build this structure, the authors had to first establish control over the shape, size and optical properties of its constitutive materials, which enabled the successive layering of the horizontal “branches” with portions of what would ultimately become the “trunks” of the trees. An array of such tree structures produced the overall optical properties that the publication describes.

In addition, the frequent comparisons between nanofabrication and cooking among users and staff alike attest to the challenge of building nanoscale structures as structures with hierarchical complexity: given an understanding of the properties of materials and the laws governing their interactions, it was nontrivial to anticipate precisely how a given combination of materials arranged in substructures would perform.

Data collection

Overall approach: Ethnographic fieldwork

The core of the study’s empirical work was ethnographic observations of and interviews with scientific support staff at each of the field sites, supplemented by interviews with users of these facilities. Ethnography involves “shadowing” participants as they go about their day-to-day tasks, observing how they speak and act and with whom they interact as the work unfolds. I strove to document the patterns and routines of my participants, as well as the meaning that participants attached to events and activities. This method enables researchers to capture mundane but important behaviors that participants take for granted and to create inductive maps of how insiders interpret their world (Geertz, 1973; Spradley, 1979; Agar, 1996).

Observation was my preferred method of data collection, for three reasons. First, earlier studies of technical work have shown that key aspects of technical competence are sensorimotor skills which are largely tacitly understood and only become evident in the performance of work (Orr, 1996; Barley, 1986, 1990a; Collins, 1974). Second, previous studies of technical and scientific work demonstrate that advances in knowledge are situated: they take place at particular times in particular places in the context of some specific problem that needs solving (Suchman, 1987; Barley & Bechky, 1994; Orr, 1996). Specifically, researchers have found that being able to document the steps and interactions that occur in the course of problem solving greatly enables them to track the sources and outcomes of knowledge creation (Lynch, 1985; Jordan & Lynch, 1992). Therefore, I expected that observations would be key to documenting the ways in which and the points at which the contributions of staff facilitated the progress of users’ research projects at their facility.

Finally, observation allows researchers to document events that participants might not remember. Because traditional interviews take place outside the flow of action, their adequacy for depicting that flow depends on participants’ recollections. I learned during my pilot study that staff contributions tended to occur sporadically over the unfolding of a user’s research project. Hence, it was unlikely that users and staff would remember all that they said and did when asked about it later. By observing scientific work in context, I aimed to capture the flow of events and then to use these observations to help participants' recall more accurately and in greater detail what they did and their reasons for doing it. Observation also allowed me to informally interview participants in the course of their work, with the goal of capturing their perspectives and reasons for doing things at the time they did them.

How the study unfolded

The data collection for this study took place from September 2018 through January 2023, in roughly six phases. There were four phases of fieldwork interspersed with two phases of analysis and writing.[6] A timeline of the data collection is provided in Figure 3.

The first phase involved what I am calling a "pilot study" at Alpha Nanofabrication Facility. This phase took place between September 2018 and March 2020, and culminated in the second phase, which involved the writing of a manuscript and several research proposals to granting agencies. When funding was granted, I proceeded into the third phase, which was a study of Beta Nanofabrication Facility. This phase took place between July 2021 and March 2022. After conducting fieldwork at Beta, I began analyzing the data I collected there between March and June 2022, while preparing for fieldwork at the next site. This constituted the fourth phase of the research. The fifth phase entailed conducting fieldwork at Gamma Nanofabrication Facility, which took place between June and December 2022. In June 2022, I also attended an industry conference. Finally, in the sixth phase, I spent a little

---

[6] I also conducted analyses during the phases of fieldwork, but these did not define a distinct phase of the project.

over a week in January 2023 collecting data at Delta, the fourth facility in this study. After this, I began to analyze the entire corpus of data. Below, I will describe in greater detail what I did in each phase of the study. Each phase included dedicated study of technical terms, building relationships with participants, and honing approaches to data collection.

Phase 1: Pilot study at the Alpha Nanofabrication Facility

The data collection for this phase took place in roughly three stages from September 2018 through March 2020. The first stage was aimed at gaining a basic understanding of staff members' work at the Alpha Nanofabrication Facility (ANF) by learning local technology, terminology, and roles. In this initial stage, I rotated among the ANF staff, typically interviewing each person once to build rapport and learn about his or her broader work context, followed by at least two occasions of observation. I professed willingness to observe any activity that, as I told my participants, would help me understand "the different kinds of work that keep the lab running." Hence, I accompanied staff members in a variety of contexts. I accompanied facilities staff as they performed inventory checks, replaced gas cylinders, completed small construction projects and prepared for the installation of a new piece of equipment. Equipment engineers allowed me to join them on periodic maintenance and gave detailed explanations of the functioning of several commonly used pieces of fabrication equipment. I accompanied process engineers as they did fabrication work in the cleanroom and met with users. The engineer who oversaw the cleanroom's original completion explained the facility's building plans. Working in this way, I interviewed and observed all staff members with the exception of the most recently hired person and another who declined to be interviewed. I participated to a modest degree, occasionally doing such jobs (under supervision) as cutting materials meant for the remodeling of part of the cleanroom, applying sealant to newly installed exhaust ducts, and handing tools to staff members carrying out repairs and maintenance, where possible.

While I observed, I took photographs, recorded conversations, and took notes by hand in a small cleanroom notebook. Because of the challenge of learning such a technical and, for me, esoteric body of information, I limited my observations to three to four hours a day, as others in technical contexts have done (Bailey and Barley 2011). I followed guidelines for ethnographic interviewing (Spradley 1979), aiming to understand how participants categorized their social world and the relations of the categories to each other. When I left the field, I elaborated the notes taken in the field into a running narrative. Then photographs which I took during the course of the observation were added to the account, followed by transcriptions of conversations. Including these different ways of capturing action helped to construct a record of the flow of events, maintaining as much as possible the referents for the speech of my participants and capturing the technical richness of the cleanroom environment. Completing a full account of a day of fieldwork often took two to three times as long as the initial observation.

During fieldwork, I encountered well-documented challenges to studying those in technical occupations. Engineers are notoriously difficult to study ethnographically, given their specialized vernaculars, reliance on equations, and use of numerous tools whose functioning may at first glance be opaque to the researcher (Barley & Bailey, 2020; Bailey & Barley, 2011; Leonardi, 2009). I therefore needed to develop scientific and technical knowledge, if I was to “hear what is said and notice what is done” (Bailey & Barley, 2011, p. 267). To meet these challenges, I completed an online course about nanofabrication technologies (Jokerst et al., 2018) and engaged in independent study of the novel terms that I encountered in the field, developing a glossary of terms. A user with a Ph.D. in materials science also agreed to meet regularly to discuss technical topics. Discussions with this participant significantly helped me to improve my fluency, which enhanced the fruitfulness of my interactions at my field site as time went on.

Because each day of fieldwork required two to three days to transcribe and annotate fully, and because I also needed to review the relevant technical literature, I visited the field site once or twice a week, on average. However, a great deal was learned in studying transcribed interviews and observations, and these initial observations served as a helpful basis for future, more focused inquiry.

The second stage began after a year of fieldwork, when I left the field in the fall of 2019 to take stock of my data and develop memos about the significant themes I had noted while in the field. Based on this round of analysis, I developed a more structured interview protocol that aimed to document the ways in which staff's role in the flow of knowledge in the facility was distinct from that of users. The third stage involved returning to the field in January through March 2020. This period was marked by productive and focused conversation and observation of staff members around structural questions earlier analysis had brought into focus (Spradley, 1979). I also interviewed twenty-eight users, to get a sense of their motivations for help-seeking and of the limits they encountered to sharing information with other users of the facility. Finally, I attended two lectures in which researchers presented work based on fabrication at the lab. I spent a total of 124 hours in the field, conducting 40 observations and 53 interviews and producing about 1,600 single-spaced pages of transcriptions and fieldnotes.

Phase 2: Analysis and writing manuscripts and grant proposals

Between March 2020 and July 2021, I analyzed data from fieldwork at the ANF with the aim of producing several written documents. In April 2020, I completed my qualifying paper (Bovenberg, 2021). That paper served as a basis for a journal submission, for conference talks and for a number of grant proposals. These proposals outlined a program of study at two other nanofabrication facilities to examine the questions that are the focus of this dissertation. In July 2021, a preliminary source of funding was found and the fieldwork began.

Phase 3: Data collection at the Beta Nanofabrication Facility

Between July 2021 and March 2022, I made several data collection visits to the Beta Nanofabrication Facility (BNF), spending a total of three months in the field. Because this was a time in which management of the Covid-19 pandemic at universities was still relatively variable and at times stringent, my fieldwork at the BNF was grounded in in-person observations but interspersed with a substantial number of virtual observations. My fieldwork at the BNF was marked by further learning about technical aspects of nanofabrication as well as the further sharpening of my approaches to data collection.

I started collecting data at the BNF with a round of interviews of all staff members of the facility, geared toward constructing each staff member’s career history, role in the facility, and everyday work tasks. After interviews, I set up times to shadow several staff members. I shadowed seven members of the technical staff, experimenting with different approaches to delineating the period of time during which I shadowed them. With three staff members, I arranged times to shadow them by asking them when they were planning to meet with users about their projects, and then joining them for those meetings. In this way, I observed a number of equipment trainings and qualifications, as well as early user onboarding meetings. I shadowed three other staff members for continuous stretches of time, typically three half-days within a single week. This too involved attending trainings and qualifications, but also gave me the opportunity to observe moments of downtime between meetings, staff members' posture of "being available" to users and responding to technical breakdowns that inevitably, though unevenly, occurred within their week, as well as the work that staff performed that did not involve face-to-face interaction with users, such as maintenance on tools and responding to emails. Finally, I shadowed one staff member who elected to conduct meetings with users almost exclusively virtually for that period of time. The first approach to shadowing allowed for observation of a greater number of staff-user interactions per unit time spent shadowing, while the second approach allowed for observation of downtime, a broader range of tasks, and email correspondence with users. The third approach allowed for easier comparison of meetings that took place in similar, virtual, settings.

I shadowed staff members in much the same way as I had at the ANF: recording audio and taking photos when given permission and taking notes in a cleanroom notebook. When observing a virtual interaction, I took notes synchronously on my laptop. At the end of each day, I wrote out the fieldnotes. On days that I was out of the field, I further completed the record of fieldwork by adding in photographs, inserting edited transcripts, adding metadata and cataloging the fieldnote in question.

In addition, when shadowing staff members, I attempted to follow observations of interactions between staff and users by interviewing both parties about what had occurred. My goal was to capture differences in goals, perspectives and expertise between staff and users. After I had observed an interaction, I typically arranged a time to interview the user as soon as possible. In conversations that ranged from 10 to 30 minutes, I asked users to describe the interaction with staff from their perspective: what had been their goals for the meeting, what were they hoping to learn from staff, what things had the staff member said or done that were especially useful, helpful or novel, and how had this interaction altered their fabrication approach, if at all. This served to explore how users viewed the staff's contributions to their research. Similarly, I asked staff members to describe the interaction from their perspective as soon as possible after the interaction had taken place. Here, I focused on asking the staff member what their goals for the meeting had been, what aspects of the user's statements, questions and process were especially important for the problem-solving in question, and how the interaction had served to advance their understanding of a given material, tool or technique, if at all.

I further supplemented my observations of staff with interviews. These interviews focused on adding clarity and detail to the events, practices, statements and ideas encountered during periods of observation. The goal in these interviews was to deepen my understanding of what staff did and knew, and to follow lines of inquiry that were inconvenient to pursue during the run of action captured in the fieldnotes. In addition, I observed weekly "process meetings," which were open meetings that the staff organized for users to ask questions about their processes. I attended 11 process meetings, ranging from 30 to 90 minutes in length. These took place virtually and I took notes synchronously on my laptop.

Concurrently, I interviewed users of the facility in order to understand what performing nanofabrication work was like from their perspective. I initially recruited users by email and in person. The director of the facility sent out an email to the facility's user base, introducing my study, inviting users to contact me by email if they wanted to participate. In addition, I recruited users by spending three afternoons in the facility's shared office space to which I brought free donuts. In this way, I recruited a first wave of participants. As the study proceeded, I also recruited users by way of snowball sampling, asking users to refer me to other users. In addition, as mentioned above, I interviewed users that I encountered while shadowing staff. Finally, when I had not yet interviewed users whose names came up repeatedly in conversation – as examples, for instance, of highly experienced users – I reached out to them directly via email to request an interview.

My most common interviewing approach was to ask users to describe their research in broad terms, and to tell me how their work at the BNF fit into that overall project or set of projects. I then asked them to describe one or more of the central fabrication challenges involved in building their device. Then I would ask a series of questions to ascertain how the user had gone about addressing this challenge, for instance asking them to describe how staff had been involved in helping them address the challenge. I also asked questions about what working in the cleanroom was like, the learning curve they had experienced in their early days of fabrication, the division of labor within their groups, and how their group maintained information and expertise regarding fabrication. Finally, a set of questions was aimed at documenting the different ways that users interacted with each other inside and outside the cleanroom.

I also experimented with a number of interview approaches to capture the development and dissemination of knowledge in specific technical domains of cleanroom work. In one interview approach, I chose three pieces of equipment in the cleanroom and interviewed the staff member responsible for the tools and as many users as possible about their use of those tools, to document how insights on how to use the tool were both developed and shared in interactions involving users and staff. Another approach I tried was to interview a subset of users whose processes had been disrupted by a supply chain issue relating to a core material in their process, in order to document how they were going about adapting their process (and the role that interactions with staff and other users played as they made these adaptations). In these different ways, I interviewed 32 users at this facility.

I spent an average of three days per week in the field. During this period, I alternated between two weeks of observation and a week of preliminary analysis. In my weeks away from the field, I edited and organized my notes, took note of salient topics, wrote memos, reflected on the outcomes I was getting from different data collection approaches, and adjusted my data collection approach accordingly. These moments of preliminary analysis were crucial because they allowed me to modify the research strategy to better investigate emerging concepts and hypotheses. At Beta, I spent three months in the field, performed 78 observations, 63 interviews, and spent 155 hours on site, which produced 1,700 single-spaced pages of transcriptions and fieldnotes.

Phase 4: Intermediate analysis and writing

After completing data collection at the BNF, I spent March to June 2022 analyzing the data I had collected there and preparing for fieldwork at the Gamma Nanofabrication Facility (GNF). This period of analysis cued me to comparisons that I wanted to explore at the GNF. This period also helped me recognize which interview and observation approaches worked well for me and my participants, and to adapt my approach to shadowing and interviewing at the GNF accordingly.

Phase 5: Data collection at the Gamma Nanofabrication Facility

Between June and December 2022, I conducted fieldwork at the GNF. The aims and approaches were largely the same as at those I had employed at the BNF, with a few differences that were informed by the preliminary analysis I had conducted after fieldwork at BNF. At the GNF, I chose to conduct fieldwork with a greater emphasis on shadowing staff members for continuous portions of their day, rather than simply joining them for scheduled meetings with users. I shadowed four staff members, each for three days within a single week. Compared to my time at Beta, I therefore conducted fewer discrete instances of observations, but clocked a greater number of hours on site.

At this field site, I also weighted my data collection more toward staff than toward users, due to two features of the facility. First, it became evident that staff members at this facility interacted frequently with each other about the problems they helped users solve. I had not observed much of this kind of interaction at the other facilities, and therefore chose to pay attention to how staff members interacted with each other around users' projects – capturing, for instance, how they queried each other about the capabilities of their tools and about their experience with a particular material or structure. In addition, average staff tenure at this facility was longer than the average in the field, and I chose to spend more time inquiring into the experiences that staff members had had over their careers, and the perspectives they had developed as a result of those experiences. I interviewed each staff member at least once and began to shadow individual staff members soon thereafter. I interviewed 14 users, in a mix of general interviews and interviews after observations, using the interview protocols I had developed while at the BNF.

I also attended eight of this facility's weekly open process meetings, as well as several meetings where new users to the facility described their processing goals to staff. Because Covid-19 restrictions at the GNF had eased at this point, most of these interactions took place in person. I also attended informal get-togethers with staff, such as concerts and picnics. At Gamma, I spent a total of three months in the field, performed 35 observations, 42 interviews, and spent 174 hours on site, producing about 1,000 single-spaced pages of transcriptions and fieldnotes.

Phase 6: Data collection at the Delta Nanofabrication Facility

In the final phase of the research, I collected data at the Delta Nanofabrication Facility in January of 2023. This was a period of productive and focused interviews with staff members. The interviews allowed me to confirm and refine many of my emerging findings across the other facilities, making use of the technical fluency I had developed over the prior periods of fieldwork. At Delta, I spent a total of eight days in the field, performed 7 observations, 27 interviews, and spent 37 hours on site, producing about 275 single-spaced pages of transcriptions and fieldnotes.

Throughout: Attending UGIM and joining the *labnetwork* mailing list

Two additional sources of data complemented the data collection described thus far. In June 2022, I attended the biannual meeting of nanofabrication facility managers and technical staff, known as the University, Government, Industry, Micro-Nano (UGIM) symposium. Attending UGIM provided a sense of the roles, work and problems shared among facilities, as well as differences in the ways these facilities were situated within their host universities and within their unique geographical settings. In addition, observing and taking part in formal and informal interactions between staff members of different facilities bettered my understanding of the occupational communities in which support staff participated. I also subscribed to *labnetwork*, the mailing list with which staff of nanofabrication facilities within and outside the United States query each other on topics of common interest. Labnetwork is a vibrant platform for discussions spanning the range of expertise it takes to run an academic cleanroom facility. For instance, I observed discussions on labnetwork that concerned requests for process parameters for specific tools and requests for manuals for antiquated equipment. There were also discussions that ranged from comparing vendors of much-used raw materials to ways of managing users who were behind on their payments. Finally, labnetwork served as a forum for posting job openings. Paying attention to labnetwork postings helped me get a sense for how widely shared the dynamics I had been observing at my field sites were, as well as providing topics to discuss with my informants that might not otherwise have emerged during shadowing.

Data included in the analysis for the dissertation

Having such a large corpus of data presented a difficult decision as to which data to analyze for the dissertation, and which data to set aside for later analysis. For this dissertation, I have chosen to focus on analyzing data from Alpha and Gamma facilities, for a number of reasons. First, research at the ANF and GNF took place at periods least affected by managerial restrictions put in place during the Covid-19 pandemic. Data collection at the ANF took place largely before the pandemic, and data collection at the GNF took place as restrictions were being relaxed. This meant that most of the staff's work took place in-person, and as a consequence, spending time at these facilities provided me with a greater sampling of the range of staff's activities. Data collection at Beta took place at a more restrictive phase of the Covid-19 pandemic. In addition, tenure of staff at the Alpha and Gamma facilities was, on average, longer than at the Beta facility. This made these facilities attractive places to examine the skills that staff developed over time by inhabiting their role. In addition, data collection at the BNF took place in a period of significant staff turnover at that facility, making it difficult to distinguish staff learning that took place as a result of a staff member being new to the job from learning that took place as a result of a user doing something new relative to other users. Therefore, for this dissertation, I have primarily focused my analysis on my time at the ANF and GNF, using data from my time at the BNF and DNF as it supports the patterns derived from that analysis.

This has a number of implications for the findings presented in this dissertation. First, the analysis presented in the empirics is aimed at uncovering similarities in the roles performed by staff members at different facilities, rather than differences among facilities. The dissertation is not comparative at the level of the facilities. Later papers may attempt to identify and trace differences between facilities. Second, the staff perspectives represented here will be weighted toward those of more experienced staff members, and staff members who work at facilities that, on average, retain their staff members for longer periods of time. Judging from conversations at UGIM and on labnetwork, turnover was common at many other nanofabrication facilities. It may therefore be the case that the empirics provide a picture of those facilities where working conditions were especially conducive to staff retention compared to other facilities. What underlies the differences among facilities in staff retention remains an empirical question that I cannot answer without further analyses and data collection.

Data analysis

A central challenge in analyzing data taken from the lived experience of individuals in a given setting is to accurately capture the meaning that words, actions and situations have to the people in the setting, while also communicating this significance in terms that an audience unfamiliar with the setting can understand (Geertz, 1973). The challenge of translating and conveying situated meaning is a feature of interpretive approaches to ethnographic research. The approach that anthropologists and other interpretive researchers have found to resolve this tension is to start very close to the particulars of a setting, capture the meaning of those particulars and work outwards toward consecutively greater abstraction (Barley, 1989; Gioia et al., 2013). In other words, theoretical assertions "can only emerge from, and must therefore remain grounded in, an understanding of the particulars of a variety of settings." (Barley, 1989, p. 4) Having understood what specific situations mean to the individuals in them is what grounds any generalization that later occurs (Geertz, 1973). In other words, the goal is to account for an experience that people in an organization have by developing concepts that are "adequate at the level of meaning of the people living that experience *and* adequate at the level of scientific theorizing about that experience." (Gioia et al., 2013, p. 16, emphasis in original) Moreover, this approach aims to make the link between the two levels transparent.

Interpretive data analysis can therefore be thought of as having three stages (Gioia et al., 2013). The first stage involves examining occurrences and statements made by participants with the goal of developing an adequate account of their experience. This approach presumes that participants know what they are trying to do and can articulate their thoughts and intentions. In this phase, I read my fieldnotes while asking, "What's going on here?" and "What are these people trying to achieve?" I read through my fieldnotes repeatedly with the goal of developing descriptions of instances of situated action that captured the character of the problems the individuals involved were trying to address. Initially, this involved adhering faithfully to the terms that my participants used to refer to their experiences. Concepts that are developed at this level of abstraction are typically called "first-order codes" (Gioia et al., 2013).

As this stage progressed, I increasingly worked to make these accounts accessible to audiences with little familiarity with nanofabrication. This resulted in a set of slightly more abstract concepts with very strong ties to the particulars I had encountered in the field. For instance, I sought to describe the recurring problems staff faced by speaking more in terms of "machines" they maintained than in terms of specific pieces of equipment for etching, deposition or lithography. This slightly more abstract set of codes is typically referred to as "second-order codes" (Gioia et al., 2013, p. 20). In developing second-order codes, I also increasingly narrowed my analysis according to my theoretical interest in the development and dissemination of knowledge in these facilities. In addition, I wrote memos about patterns that struck me as important or puzzling. For instance, at this stage, I began to notice that staff viewed every user's project as novel – for instance, saying, "no two groups are doing the same thing" – yet also made frequent reference to these projects as a recombination of recurring, known elements.

Throughout the first and second stages, I used *Atlas.ti,* a software tool specifically designed for coding and analyzing textual data, to classify the recurrent activities, strategies, and techniques that comprised the work of support staff and their interactions with the facilities' users. For example, I developed typologies of the problems that scientific support staff encountered, the tools they used, and the steps they took to solve problems. I distilled the observational and interview data into accounts of how staff moved from specific solutions to particular problems to general knowledge and techniques that could be widely distributed. These accounts raised themes and concepts that I then explored further, thereby triggering additional rounds of coding and analysis. This iterative process of analysis and conceptualization is central to the generation of grounded theory (Strauss & Corbin, 1990).

Rigorous grounding of emergent concepts in the particulars of data makes possible further discussion about these concepts at a more abstract level: in relation to each other and in conversation with the concepts and understandings described in related bodies of literature (Gioia et al., 2013). The third stage of analysis therefore involves thinking about how the defining features of action in a given setting relate to each other in light of a given outcome of interest, often in a more dynamic model. In the present case, I was interested in the distinct knowledge that staff possessed relative to the users with whom they worked. In the course of fieldwork, I had begun to more closely read Harper's book *Working Knowledge* (1987) and recalled a paper I had read by Herbert Simon (1962) about the properties of hierarchical complex systems. I formulated the "mini-hypothesis" that staff were helping users build structures that had a hierarchical complexity to them, and that therefore the distinct knowledge that staff possessed related to exerting control over the basic elements of these structures, the elemental substructures they could be arranged into, and the integration of the whole into a larger structure. This hypothesis would account for the pattern that I had observed earlier in my analysis: that staff saw user's projects both as in-principal a recombination of previously existing elements or bits of knowledge, but in practice as a novel combination the integration of which needed to be worked out for each structure anew.

With this mini-hypothesis in mind, I then proceeded to test the fit with the data, one instance at a time. At this time, I also further read Harper's account in light of the concepts that Simon developed, to test its fit with my developing account. The results of this hypothesis-testing and refinement became the body of this thesis. I wrote chapters on how staff helped users gain exact control over the properties of materials (Chapters 5 and 6) and how they helped users integrate the many steps of their processes (Chapters 3 and 4). I sent drafts of the chapters to the study participants featured in the chapters. They provided further technical detail and corrections, while supporting the broad argument made in the text.

Chapter 3. Process Integration: Significance, definitions and examples

Introduction

Building nanoscale devices involves layering different kinds of functionality within a single structure. For example, a classic transistor, whose purpose is to function as a conditional switch within an electrical circuit, must incorporate regions of material that conduct electricity, regions of material that serve as an insulator, and regions of material that can be influenced to perform either function. In nanofabrication, these regions are typically created one layer at a time: depositing material in the form of thin films, patterning them, and removing selected areas. The work of *process integration* involves ensuring that each region is fabricated in such a way that its fabrication does not adversely affect the ultimate functioning of other parts. In other words, someone who pays attention to process integration when building a device is trying to layer the functional elements of their device in such a way that they obtain only the combined properties and interactions that they want, and not the adverse properties and interactions that may just as well arise when one tries to apply diverse and potent fabrication processes to a sample that contains many materials and delicate structures. One senior staff member defined the achievement of building an integrated device as follows, during a presentation to a group of prospective users:

> So when you look at a Pentium chip, and you consider it the norm, I would look at a Pentium chip and I would say it's the exception: it's the one combination of 55 or 80 or 270 or 390 steps that *don't* interact [*adversely*] with one another.[7]

---

[7] Throughout the dissertation, I will follow the following conventions for presenting data: offset quotes or body text within quotation marks indicate verbatim quotes by participants. Within any given quote, italic text in square brackets indicates words that I have inserted to either anonymize a quote or clarify the meaning of a given statement. Unbracketed text in italics indicates the speaker's emphasis. Regular text in square brackets indicates actions taken by the speaker.

Therefore, the task of process integration can be defined as a search for a set of fabrication procedures by which the conditions for device functionality are met and the conditions for numerous fabrication failure modes are prevented. One fabrication failure mode has already been named: the fabrication of one region hurts the functionality of another. In addition to avoiding device-defeating failure, a fabrication process must also support the continued sound operation of the cleanroom's shared tools. Primarily, this means that processes must abide by the material restrictions that are set in place around different tools to avoid cross-contamination of samples that contain incompatible materials. That is to say, breaking the materials segregation rules of the cleanroom also counts as failure when integrating a process.

Integrating a nanofabrication process is difficult for several reasons. One reason is that, as we will see in Chapter 5 ("Being a Tool Owner"), the machines of the cleanroom are in principle capable of affecting samples in many different ways. The tools do not foster only those material properties that users consider to be focal to their process, but also several properties they may consider peripheral or detrimental to their process. Just because a user wants to make an *electronic* device, for instance, doesn't mean that they won't have to manage the *mechanical* stress that accumulates in their sample because they have layered materials that have different thermal expansion properties, resulting in small stresses at material interfaces that ramify throughout the structure as a whole. That is, the elements of a given a structure may interact in many different ways with each other and with fabrication procedures applied throughout a fabrication process, through mechanisms that lie outside the user's area of expertise. Therefore, successfully integrating a process requires familiarity with the working principles of a range of commonly-used tools and the diverse physical mechanisms by which fabrication failure (and success) can occur.

I begin the empirical section of this dissertation with a discussion of process integration because across my field sites, staff members said that process integration was very important for the successful fabrication of users' structures. Staff members saw their work as helping users *develop a working process*: selecting and sequencing tools and techniques such that a micro- or nanoscale structure with precise properties would result. Considering a user's overall process and ensuring its feasibility was therefore at the core of staff's work. To be sure, developing a process also involved other work and other kinds of knowledge, such those involved in skillfully operating the tools of the cleanroom – a topic Chapter 5 picks up – but the initial selection of appropriate tools and the broad approach to using these tools was often established at an earlier time outside the cleanroom. Process integration was in an important sense a design philosophy by which fabrication steps were devised and evaluated. Therefore, to understand the rationale that informed many decisions made within a user's process and the role that other forms of staff expertise played, it is important to understand the nature of process integration problems.

Occupational knowledge related to process integration

The importance that staff members across field sites accorded to process integration for the success of fabrication projects directs our attention to two important features of their role. First, it underscores the fact that staff at nanofabrication facilities preside not simply over a single machine whose interface they manage, but rather over an ensemble of machines and a repertoire of techniques, whose broad capabilities they arrange and adapt in shifting configurations to meet the specific construction needs of numerous projects. This recurring task of matching what users want to build with what the tools can do lies at the center of what nanofab staff do daily, and the unique role they play within their facilities.

Second, in attending to process integration, this occupation shows its kinship to other occupations that are organized around a discipline of construction and repair. John van Maanen (2011) writes,

> "In any hard discipline, be it engineering, gardening, auto repair, or ethnography, the learner must submit to things that have their own intractable ways, an authoritative structure that commands respect." (p. 219)

For occupations that conduct construction and repair, that authoritative structure is the imperative to build or fix integrated structures composed of many interlocking elements (Simon, 1962) such that these structures function as wholes when the job is done. Examining the work of nanofabrication staff in the domain of process integration allows us to consider the kinds of knowledge that these occupations develop around this central challenge.

To name this knowledge, it is worth considering resemblances between the work of staff and the work of Willie, the master metalworker who is the focal actor in Douglas Harper's (1987) account of craft skill and its role in a community. A characteristic of Willie's approach to problems is that his attention is distributed across the numerous factors that are at play in a given problem, and the balance he achieves of all these factors in a single, working solution. As Van Maanen (1990) writes,

> "To make a good, wise, sensible choice in Willie's workshop is to accept the interconnection of all possible factors involved in a particular problem and avoid the mistake of focusing on any one consideration – cost, time, appearance, efficiency, cleverness, strength, endurance and so forth." (p. 279)

The need to pay attention to a wide array of structural properties that together determine the success of a structure is most likely a common feature of the work of occupations whose focal discipline involves construction and repair. For the intervention to be considered successful, the structure must function once the job is done. For the structure to function, many properties must be considered. While users of the cleanroom may be able to segment into scientific disciplines with specialized knowledge about a single facet of the natural world, to effectively help users build successful devices, staff members must be conversant across a range of phenomena that must be managed for a functioning device to result. Staff's practice of process integration exemplifies this breadth.

In this chapter, my goals are to illustrate the central challenge that process integration poses for staff and users, and to show the breadth of physical properties to be managed while building a micro- or nanoscale structure. To do this, I will first briefly list the definitions that staff members at the Gamma facility gave of process integration. To exemplify the breadth of process integration challenges, I then describe several classes of integration challenges that staff spoke of during my time in the field. The findings presented in this chapter also serve to foster the requisite background knowledge for understanding the examples given in the next chapter, which describes how staff addressed process integration challenges in their interactions with users.

Staff definitions of process integration

Table 3 lists definitions and descriptions that staff members at the Gamma facility provided of the work of process integration. In their own words, these different staff members describe the necessity of process integration in nanofabrication and the character of the challenges involved. In so doing, they hint at the distinct skills that experience at process integration brings. The definitions share an emphasis on seeing the overall approach (Victor calls it "the whole picture") of a process and moving back and forth between this holistic picture and the concrete steps that comprise it, ensuring that individual steps align with the objectives of the overall project.

Table 3 about here

Common integration challenges

As indicated above, integrating a process involves meeting the conditions of success while avoiding the conditions for failure. Failure in a nanofabrication process can take place at various scales, in different time frames and through different physical mechanisms. Some failures can take place in an instant, such as when the chemical solution used to remove the top layer of a device also damages one of the lower layers. Other failures can take an entire process to unfold, such as when successive layers of deposited materials, due to their mismatched thermal expansion properties, generate sufficient compressive stress on the wafer that the wafer bows and its layers peel off. Some, such as the case of thin film stress, operate by processes of mechanical force. Others, such as the unwanted removal of prior layers in an etch step, operate by chemical interactions. Therefore, some failures can be thought of as "budgets" to be managed throughout a process, as in the case of the "thermal budget" of a device: how long can it be kept under high temperatures? Others can be thought of as "ceilings" to remain below: What is the maximum temperature this device can reach during fabrication?

In other words, one's ability to anticipate integration problems depends on one's ability to look at a process through the lens of diverse physical mechanisms. Common integration challenges that staff described while I was in the field were: managing thin film stress, attending to thermal budgets and ceilings, maintaining surface electrical properties, ensuring chemical compatibility when removing material from constructed stacks, maintaining alignment of features across device layers, monitoring the sample non-invasively throughout, and abiding by the cleanroom's system of material restrictions. The discussion that follows is not meant to provide a comprehensive catalog of integration problems, but rather to illustrate the types and range of problems typically encountered when trying to integrate a process. Table 4 displays these classes of problems in tabular format.

Table 4 about here

Managing thin film stress

Managing integration of a process takes certain main forms. One form it takes is managing certain sample qualities over the course of a project. As mentioned earlier, thin film stress is one of these qualities. Different materials have different rates at which they expand when they are heated. Therefore, when different materials are layered on top of each other, a process which often occurs at high temperature, small stresses are created at their interfaces, which are exacerbated when the structure is cooled to room temperature for further processing. The sum of stresses thereby exerts compressive or tensile stress on the structure as a whole, potentially causing material to delaminate from the wafer. Like many common problems in nanofabrication, thin film stress is also viewed as a valuable property by some application fields and is tuned to desirable levels throughout a process.

Anticipating and addressing integration challenges involving thin film stress means monitoring and managing stress over the course of a project. Each facility I visited had tools for measuring film stress. In addition, at Gamma there existed data gathered over time by staff and users reporting on and plotting the relationships between recipes (tool parameter settings) used to deposit a given film and the resultant film stress. Staff have also been known to recommend depositing compensating films on the reverse of a wafer, to balance out the net stress. These approaches can all be used in the service of managing film stress over the course of a process.

During an observation of an early feasibility meeting with a user who had not yet begun fabrication of his device, I observed staff caution him that film stress was likely to be a central problem to be managed in his process, which involved depositing thick layers of two materials that I will refer to as Materials A and B to preserve the user's anonymity.[8] Adam tells the user,

> I want to season you appropriately that there's a lot you're asking for. There's a lot of things we don't know about, like [*depositing*] this much [*Material A*] and [*Material B*], what that's going to work like, how the stresses are going to evolve – [*the materials*] might just jump right out of these [*features*] on you.

Adam is suggesting that it is possible that when the materials are deposited in such thick layers as the user's design calls for, the resultant film stress might be so great as to result in the materials peeling off of ("jumping out" from) the wafer. During another observation made while shadowing Tim, an equipment engineer at Gamma, he showed me documentation he had collected regarding the stress properties of the films deposited in a tool for which he is responsible (a role known as being the "tool owner" of a given tool):

> Tim and I also talk about "stress balancing" in films. Tim pulls over a nearby binder that says "[Tool A]" on it and flips to a page that displays a graph of film thickness versus stress for different films. Tim says ideally, you want stress to be at zero, because stresses will cause your wafer to bow, for example like a potato chip. When a wafer is bowed to a certain degree, you can't put it into some of the etching tools,

---

[8] This is a convention I use throughout this dissertation. I will restart numbering and lettering in every example. Material A in one example should not be taken to be the same material as Material A in a different example.

> because you can't flow helium on the backside of your wafer to cool it. Tim explains different scenarios involving different layers and different resulting total stress. (Excerpt from fieldnotes)

Attending to thermal budgets and ceilings

Another form that managing process integration takes is determining and abiding by maximum levels of exposure to some environmental condition. An example of this is the "thermal ceiling" of a sample. If a user wants to create a structure that incorporates semiconducting materials that are created through a process called "doping", in which ions are implanted into a material, then this effectively caps the maximum temperature at which the sample can be held in subsequent processing. Exceeding this maximum temperature level will cause the ions in the doped region to diffuse to other regions, effectively undoing the doping. Given that several classes of tools in the cleanroom routinely exceed temperatures of 200 degrees Celsius (392 Fahrenheit), having a temperature ceiling is a considerable limitation to take into account when developing a process.

Maintaining surface electrical properties

Another set of integration considerations involve the electrical properties of a structure's surface, which can be undone if the wafer is later processed by a tool that operates by bombarding a sample with ions. This is the case in a tool called the reactive ion etcher (RIE), that Adam mentions to prospective users:

> RIE damage: if you have a very thin device layer, it can get totally beat up by those energetic ions and get scrambled. So, it loses its properties. That's kind of horrible.

Adam tells users that one way to manage this potential hazard to one's structure is to protect electrically sensitive regions of a device with sacrificial layer while an RIE process takes place. This is another form that managing integration challenges takes: building features that effectively shield vulnerable portions of the structure from harmful forces involved in later processing steps. However, such temporary features must be removed at later stages of one's process, and this creates the set of integration challenges discussed next.

Ensuring chemical compatibility when removing material from constructed stacks

Another set of integration considerations is chemical, and involves the chemicals used to remove layers of materials from constructed film stacks (as layers of different films are often called). As a stylized example, it may be that the chemical most suited for etching Material A on one's materials stack also etches the layer below it – Material B – and does so more rapidly. This means that once the etch chemical "punches through" the top layer of the stack, it will begin to etch the lower layer as well. Unless one wants or can tolerate both layers etching at once, one must find an etch chemical ("an etch chemistry," in the local jargon) that etches away Material A and not Material B.

These considerations are significant if one considers that many of the features created on a sample throughout the process are placed there for temporary purposes (e.g., as sacrificial layers, as mentioned in the prior section), and must be removed for the final device to function as intended. But at the end of one's process, there are many more functional structures on one's wafer. So, finding the chemistry with which to remove only the unwanted materials, without damaging existing structures, is a challenge. Adam refers to a range of possible techniques for meeting this challenge when he tells prospective users about this set of integration challenges:

> There's a lot of ways to strip materials and get things off. Again, talk to us, you don't need to use a full on [*Tool Type A*]. Maybe you can use a [*Tool Type B*] or [*Tool Type C*] like one of the [*Vendor Name*] systems in the lab to remove resists or other materials. And there's a lot of other chemistry tricks and things of the trade. Again, talk to people, talk to us, we'll help you out.

Maintaining alignment of features across device layers

Another set of process integration considerations involves being able to ensure accurate placement of features on the sample's surface across multiple layers of material. Maintaining alignment throughout a process is typically done through the application and reapplication of features known as "alignment marks" (high-precision features that are used as the reference when positioning subsequent patterns) on each new layer of material, as Adam tells prospective users:

> You'll do this once or twice in your life and then never do it again, because it's a painful thing. If you design bad marks, or you're using somebody else's marks, and then you get down the road and you can't see the mark anymore, because you've got other films lying on top of it, or it got damaged during an etch or something like that. Trying to keep one set of marks or a couple sets of marks going through a process can be really challenging. A good thing to think about if you're doing multi-level processing with [*Lithography Tool Type A*] or [*Lithography Tool Type B*], is having lots of marks. Lots of marks. Big fan of lots of marks. Have a pristine set of marks that gets put on, aligned from the last step. Every new mask layer, you're going to have beautiful new marks.

Monitoring the sample noninvasively throughout the process

Another set of integration considerations involves being able to monitor the unfolding of your process non-invasively throughout. Sometimes, measurement techniques will damage or destroy aspects of device functionality, such that process integration will involve devising alternative ways of inspecting a sample. For instance, scanning electron microscopy involves bombarding the sample with electrons. Some forms of profilometry involve dragging a needle across a sample. For users who cared about properties of their sample that could not be measured without damaging their focal structure, staff recommended incorporating "test structures" onto their wafer, which Adam describes as "structures that are somewhere else on the wafer, away from your active devices, but in reasonable locations, that will reflect the processing." He also recommends adding instances of one's pattern (known as "die") for the purpose of "destructive testing" below:

> So putting an extra die, those die are really useful for destructive testing: you can saw them through cross sections, throw it in an SEM, and see what the [*Property A*] look[*s*] like, things like that. So there's use to these things.

Abiding by the cleanroom's system of material restrictions

A final and crucial set of integration considerations stems from staff's efforts at maintaining the performance level of tools through the enforcement of strict materials restrictions. It is not uncommon for materials that are used in one class of devices to destroy the functionality of another class of devices. For instance, gold and copper create electrical shorts when they appear in unwanted places on electronic devices. Therefore, gold and copper belong to their own "tool class" and wafers containing these materials can only be processed within tools of that class or lower classes. These material groups create pathways through the cleanroom that ensure that mutually destructive materials to not cross-contaminate across samples. In addition, the cleanrooms maintained separate spaces for handling chemical classes known to interact adversely with each other. These materials restrictions limited which tools were at a user's disposal given the materials in their device. Once a user had used a tool of the "dirtiest" class, they were no longer allowed to put it into "cleaner" tools. Therefore, at the start of their process, users had to anticipate which classes of tools they would need and find a way to sequence their processing such that it progressed only from clean to dirty and never in the reverse direction. Staff spoke about this as "creating a path" for users.

For instance, in the excerpt below, a staff member at Gamma who specializes in plasma etching tools talks about creating a path for users with uncommon materials, particularly materials that put them into a more restrictive tool class:

> [*For*] people with weird materials, we need to figure out, (a), what chambers [*i.e., tools*] they can put their samples in, and (b), what that means for the rest of their process. Because let's say they have a material that's not allowed in most chambers. Okay, well, we probably have an etcher that you could etch that layer in. But once you put your wafer in a dirty chamber, you can't go and put it in a clean chamber, even if it's for etching of a layer that is allowed in that chamber, because you've contaminated your sample. So there's integration challenges with certain materials that have to be worked out.

The prevalence of multiple process integration considerations

The examples given above highlight a single integration challenge at a time. Now, what if the functionality requirements of a user's device meant that they had to manage its compressive stress and low thermal budget, while also using an esoteric material in an early step, which limited them to the dirtiest and most restrictive class to tools? How were they going to abide by all these restrictions on their process while building the structures that will create the functionality of their device? This is where staff's integration expertise came in. Finding possibilities through layers of limitations is why one experienced user referred to staff's ability to integrate processes as "threading the needle through the process."

As the complexity of a device grew – as it incorporated increasing numbers of layers of functionality, for instance – the area of overlap between the functionality and failure conditions grew, and the call for integration expertise increased. A staff member experienced in integration could be exceptionally targeted in the solutions they devised. Storied successes at integration were those in which the window for success appeared insurmountably small but was found through clever use of the tools at hand. Being able to find such solutions involved understanding both the conditions for success and failure within the context of each specific project, as well as the tools at one's disposal to intervene in them.

The next chapter considers how the team of staff at a facility interacted with users and worked together to develop targeted solutions that met the fabrication needs of a particular structure and were also feasible given the specific capabilities and process integration constraints of the cleanroom's tools, techniques and materials.

Chapter 4. Process Integration in Practice

Introduction

The previous chapter served to illustrate the breadth of knowledge that staff applied to process integration problems. The present chapter is about the second noteworthy aspect of process integration: the *precision* of the approaches that staff ultimately devised. Because each nanoscale structure brought together a novel set of objectives and materials, the precise process constraints and objectives of each device had to be determined anew, in order to select tools, order process steps, and tune tools appropriately. Every feature had to be fabricated in such a way as to fit within a multistep process. Therefore, an important aspect of staff's ability to match what users wanted with what tools could do was their ability to discover the specific fabrication needs of any given structure. I'll call this *discovering the situated fabrication needs* of a device. The needs were "situated" in the sense that they were unique to a particular fabrication process, not standardized and, therefore, not amenable to modular solutions.

As I argued in the Chapter 1, a distinctive feature of the role played by Willie – the master metalworker who is the focal actor in Harper's (1987) account of the role of craft skill in a community – is that he is a central actor in a community whose needs for fabrication and repair are not amenable to standardized solutions. He must build, repair and improve upon integrated, hierarchical structures (Simon, 1962) in the absence of standardized needs and modular parts. Willie is not a "parts exchanger", but rather builds features that fit within particular structures that operate in specific contexts of use. Because Willie is a single person, Harper's account does not provide much insight into the process by which Willie discovers situated fabrication needs and matches these with the requisite materials, tools and techniques of his workshop. At nanofabrication facilities, a team of people worked on problems through consultation with each other. This gives us an opportunity to distinguish different roles, tasks and skills involved in the process of discovering and meeting the situated fabrication needs of a broad range of structures.

Overview of argument

My goal in this chapter is to describe how staff discovered and met users' situated fabrication needs during the problem-solving interactions that I observed. In doing so, my goal is to identify the distinct staff roles, tasks and kinds of knowledge involved in doing this work. In this section I will first offer the account in condensed form, and in later sections I will present further supporting evidence in the form of observations and interviews.

The objective that staff members pursued in technical problem-solving interactions with users was to bring into alignment the user's stated fabrication need with the capabilities of the cleanroom's tools. These two elements had to align for a fabrication approach to be both *suitable* to a user's project goals and *feasible* given the facility's capabilities. During such problem-solving interactions, staff members brought these two elements into alignment by pursuing three interrelated courses of action: (1) defining and refining the user's stated fabrication need, (2) generating roughly equivalent approaches to meet this need, (3) ensuring that the capabilities of discrete tools could together perform the proposed approaches. Each of these elements (the fabrication need, the proposed approaches, and the relevant capabilities of tools) evolved throughout these meetings, and did so as staff members brought them into relationship to the others. Figure 4 puts these concepts and their relationships in graphical form.

Figure 4 about here

Typically, a user approached staff with a stated *fabrication need*, such as a feature that they wanted to fabricate, which would ultimately function within their larger structure. With an initial understanding of the user's fabrication need, staff *proposed one or more possible fabrication approaches*, verbally sketching a rough sequence of process steps, involving specific materials, tools and techniques available in the cleanroom. Second, staff members then used one of these approaches to *probe the fabrication constraints and objectives of a user's device*, asking questions to uncover how well the proposed approach matched the user's specific structure. For instance, in one of the examples given later in this chapter, a staff member asks a user how well their device would tolerate high temperatures, given that the approach under consideration involves bringing the sample to high temperature.

Users responded in different ways to these probes. Users sometimes stated the problems they saw with the proposed approach relative to their objectives. When they did this, they *refined* their stated fabrication need to an extent that they had not yet done in earlier definitions of the problem. In other instances, users agreed to the suitability of the proposed approach. When they did this, they *supported* the case for its suitability to solving the problem. A typical probe is illustrated in the excerpt below. A staff member, Laurence, asks a user, Todd, whether his device can tolerate a given deviation from the technical drawings that Todd is presenting. This leads Todd to state that this deviation is acceptable, and further inform the staff about the device properties he cares about most:

> Laurence [gestures to the screen where a technical drawing of Todd's design is projected]: How critical is it for you to have a gap between the [*Material A*] and the silicon wall? Because you could do [*Etch Process 1*], and that would work really well.
>
> Todd [gestures to a region of the technical drawing]: Like this?
>
> Laurence: Yeah, exactly, right there.
>
> Todd: [*That*] shouldn't be problematic. We're only concerned about [*Property A*] here. So it's not going to be as problematic because it's not going to create any sort of [*interference with Property A*].
>
> Laurence: There's [*fabrication*] options, then.

Moreover, interspersed with these questions to users, staff members also used the proposed approaches to *query tool owners regarding the capabilities of their tools to underwrite a given approach.* These queries invited equipment specialists to comment on the feasibility of the approach under consideration and raise further integration issues that use of their tool introduced. Depending on their response, tool owners either *bolstered* or *weakened* the feasibility of the approach under consideration. In this way, the proposed approach anchored the discussion, serving as a means to refine the user's fabrication needs and match these with the capabilities of tools. Across several iterations, staff formulated an increasingly suitable set of alternative approaches, with which they continued to probe and query. This led to a fine matching of fabrication needs with concrete fabrication approaches. Figure 5 puts these concepts and their relationships in graphical form as they might evolve over the course of a stylized project.

Figure 5 about here

Successive probing of users' fabrication needs led to a progressively fine definition of their situated fabrication need. Querying equipment staff for the capabilities of their tools to underwrite specific fabrication approaches confirmed that the feature in question could in fact be made. Staff worked iteratively, entertaining several possible fabrication approaches during the course of a conversation and for each of these approaches moving through several such cycles of probing fabrication needs and querying tool capabilities before they reached one or more fabrication approaches that were *both suitable and feasible*. It was inherent in nanofabrication work that solutions were not reached in conversations but had to be fashioned in the lab. But when aided by staff members, users' efforts in the lab were couched within larger strategies aimed at creating alignment between the objectives of users and the combined capabilities of the facility's tools.

Occupational knowledge related to the matching of fabrication needs to tool capabilities

This account of how staff matches users' objectives to tool capabilities suggests that the staff's distinct knowledge included their ability to (1) probe users' fabrication needs based on an anticipation of how discrete tools and techniques would affect the properties of a user's structure, (2) generate alternative approaches to meeting a specified fabrication need, and to (3) query each other on the capabilities of their tools to underwrite a given fabrication approach. Moreover, staff's distinct knowledge resided in their ability to iterate across these three elements, using each to improve the elements' overall alignment.

Importantly, users were not typically positioned to discover the precise fabrication demands and constraints of their process with respect to the cleanroom's many tools, because it took a familiarity with how these tools operated to ascertain the suitability of a given fabrication approach for a specific structure. As the examples in this chapter show, in the course of these conversations, users often ended up providing more information about what they wanted to build than they did in their original statements, because they were responding to staff's precise probes, which were informed by staff's experience with these tools' effect on other samples. Some users also altered their designs in response to the limitations of tools and techniques, in instances of "situated redesign" (Rahman & Barley, 2017). The precision of a solution, in this context, then, was its degree of alignment with the objectives and constraints imposed by the user's overall goals for the device.

In addition, as the examples below illustrate, users were often not aware of the alternative fabrication approaches available to them. In part, this was due to the fact that generating alternative approaches to fabricating a given feature required an understanding of what the tools were capable of, an understanding that was informed by long-term experience with these tools (a topic that Chapter 5 covers in more detail).

Throughout the examples, we also see that these interactions required two different staff roles: one that guarded the overall integration of the process – managing what we might call a "conversation" or "negotiation" between what the user wanted and what the cleanroom's tools were capable of – and one that supported the matching of the demands of the device with the known capabilities of tools, a role most often played by "tool owners" of the tool in question (Chapter 5). Both were unique roles in the flow of knowledge in these facilities. As the examples in this chapter illustrate, the former role could be performed by different staff members throughout a problem-solving interaction. Most routinely, the role was performed by staff whose expertise was known to be related to "process" more than to "equipment," though facilities differed in the degree to which these roles were entirely separate.

Equally, it is important to recognize that the quality of the matches these staff members negotiated between fabrication needs and tool capabilities depended on how well tool owners knew the capabilities of their tools. For this matching between situated fabrication needs and tool capabilities to be sufficiently well-fitting, equipment staff had to know their tools in a way that they could respond to the precise queries they receive during these kinds of interactions. Chapter 5 on “Being a Tool Owner” is about how equipment staff got to know their tools in such a way.

Examples of problem-solving involving process integration

Integrating a user’s process took place anew for each user’s project and occurred throughout their process development inside and outside the cleanroom. Through the examples below, I will argue that developing solutions to process problems that were both suitable and feasible involved the processes illustrated in Figure 4.

First, a fabrication need is defined in the course of such a conversation, and after that a set of roughly equifinal approaches is proposed. Then, staff members use one of these fabrication approaches to probe the user’s process objectives and constraints. The user’s response to these probes further refines the user’s fabrication need, because the user will often either raise problems with a specific approach or respond that the approach is acceptable. Staff also uses the proposed approaches to query tool owners of the pieces of equipment involved whether their tool can perform the steps they’re proposing. In this way, the staff members refine the proposed approach both in terms of its fit with a user’s goals for the device and in terms of its feasibility given the capabilities of the tools of the cleanroom. The outcome of such discussions is a proposed approach for which staff have a measured level of confidence regarding its feasibility.

Example 1: Film stress and temperature ceiling

To show how process integration considerations are weighed in the context of a user project, let's consider Fred's process problem. Fred is a user at the Gamma facility who has come to the weekly open problem-solving session that staff at Gamma organize. He asks staff for their help in developing a feature known as an *etch mask*. Etch masks are layers of material used to protect some parts of a structure during etching while exposing other parts to etch chemicals. The principle resembles that of a stencil used for painting a wall, except the open regions of the etch mask are used to selectively remove material rather than add it. Fred has already deposited an etch mask made of silicon dioxide (SiO_2) onto his sample, but the mask has proven too thin for the plasma etching that he's had to do. By the time the plasma etch process has removed the desired depth of exposed material, it has also etched through the mask. Fred phrases his request to staff as follows, indicating that he also has concerns regarding the strain that will be introduced when he deposits more material onto his sample:

> Fred: So I'm coming to you and saying, Okay, I've got these wafers. They have oxide [*layered*] on them, and I was wondering what I could do to make the mask harder so that I can do my etching and I don't want to increase the strain in the wafer too much.

In the course of the 40-minute conversation that ensues, problem-solving proceeds roughly as follows: staff define the needed feature (a more robust etch mask); then they generate three possible approaches to fabricating this feature: depositing a mask made out of aluminum oxide (Al_2O_3), adding more SiO_2 to the existing layer, or applying a thicker photoresist layer than the process currently calls for. Then, staff assesses each approach in turn by exploring the tolerance of Fred's device to the process conditions that these approaches introduce. They also make targeted inquiries of tool owners about what their tools can do, for instance to assess the levels of stress that their processes introduce. This raises the possibility of depositing a layer of SiO_2 on the reverse side of Fred's wafer, to compensate for any stress introduced by the thicker etch mask. Finally, Fred decides to proceed with adding additional SiO_2 to one side of his wafer.

At the start of the conversation, Fred asks about the possibility of depositing a layer of aluminum oxide:

> Fred: So my question [*is*] about aluminum oxide. I don't know anything about how it will strain the silicon.

Adam begins to argue that the strain on the wafer from the Al_2O_3 film will be low, because the layer of Al_2O_3 needed as a mask will be very thin, and thinner layers create lower stress. The needed layer will be thin because the resistance of Al_2O_3 to the relevant etch chemistry (known as its "selectivity") is high:

> Adam: [*The layer will be*] very thin. Your selectivity is going to be a couple 1000 to one. So you can get away with a couple 100 nanometers [*of* Al_2O_3].

Joe, an equipment engineer responsible for a different class of tools than applies to this problem, queries Jason, who is the "tool owner" of the deposition tool that they are currently considering for the deposition of the Al_2O_3, as to the capabilities of the tool. Specifically, he asks whether the strain of the Al_2O_3 film can be modulated during deposition:

> Joe (to Jason): Can you adjust the [*strain*] in the dep [*deposition*] at all?

Jason responds that the strain cannot be modulated with this tool, but that he does not expect the thin Al_2O_3 layer to result in high stress levels, reiterating Adam's earlier point:

> Jason (tool owner): Not really. But there's only so much strain; and so for 100 to 200 nanometers, it's not that much strain.

Next, Fred and the staff discuss the second possible approach to making the etch mask: adding more silicon dioxide to the existing layer. Earlier in the conversation Jason had suggested this be done by means of a deposition technique known as plasma-enhanced chemical vapor deposition (PECVD), in a tool for which he serves as tool owner. Adam reiterates this suggestion:

> Adam: I think Jason's first statement about adding more oxide through a very readily, easily done [*method*] of PECVD makes a lot of sense. You know, bulk [*the layer*] up to something north of two microns.

Fred responds by reiterating his concern that this one-sided deposition might introduce strain on the wafer:

> Fred: And the fact that it's only on one side of the wafer? [*i.e., will this induce problematic strain?*]

Joe proposes that Fred can put SiO_2 on the reverse side of the wafer as well, to balance out the strain:

> Joe: You can put it on both [*sides*]

Adam suggests that Fred find out empirically whether the strain is a problem:

> Adam: It should be okay. But you can just sit there and run, I mean, it's so fast. Literally a half hour of your life and you're gonna put 700 nanometers on both sides and clean the chamber and be gone.

Laurence, a staff member responsible for a patterning technique, questions whether balancing out the strain is needed, given the tolerance of Fred's later patterning steps to deformation of the wafer:

> Laurence: I don't think you need that. Because if [*the wafer*] bows a little bit when you go and do the patterning, I don't think you're going to notice, actually.

Joe proceeds to query Jason as to the stress coefficient of the PECVD SiO_2 film:

> Joe (to Jason): It’s less than a micron [*thick*]. And how stressy is that stuff?
>
> Jason (tool owner): 180 megapascals.

Repeating Adam’s earlier suggestion, Joe suggests that Fred find out empirically whether the strain is a problem for his sample. Adam adds to Joe’s suggestion by proposing a concrete course of action for doing so:

> Joe: It may not make a difference anyways. Let’s just try it. I mean, it’s three minutes to see what happens.
>
> Adam: Measure it in the [*Stress Measurement Tool A*] before, throw 3, 4 minutes’ worth of deposition on there, throw it through the [*Stress Measurement Tool A*] again, see what you got. If it looks crazy bowed…

After staff have determined the in-principle feasibility of depositing 700 nm of SiO_2 on one side of the wafer with the PECVD as well as a means of testing for its fit within Fred’s stress constraints, Jason proceeds to probe Fred’s process for a second integration consideration: the temperature ceiling of the process. PECVD is a high-temperature process, typically operating between 200 and 400 degrees Celsius. Its place in the process must therefore be considered in light of Fred’s temperature limit:

> Jason (tool owner): What’s the maximum temperature limit on your [*sample*]?

Fred responds that a step in his process sets his ceiling at 100 degrees Celsius. In response, Jason recommends that Fred sequences his process steps such that he performs the PECVD deposition before he performs the temperature-limiting step in his process.

> Jason (tool owner): So you gotta do this [*PECVD*] before that [*the step that sets your ceiling at 100C*]

Joe asks whether there is a lower-temperature process Jason could recommend instead. Jason responds in the affirmative, stating that it's a possibility to use a deposition process known as high-density plasma chemical vapor deposition (HDPCVD), which operates at lower temperatures:

> Joe (to Jason): What about, don't [*we*] have a low-temperature version?
>
> Jason (tool owner): You could do HDPCVD, yeah.

At the conclusion of the discussion, Fred says he will try depositing PECVD SiO_2 on one side of his samples. If this doesn't work, then he might try depositing Al_2O_3 or thickening the photoresist. He thanks the staff members and leaves the room.

Brief synthesis

The example above shows staff probing the various process constraints of Fred's process by means of specific questions. These questions arise from a concrete approach that is under discussion at the moment. For instance, Jason asks Fred what the temperature ceiling of his process is, because the approach they're considering involves processing at high temperatures. We also see that the approaches that staff propose are concrete. They typically specify discrete tools (e.g., PECVD; Stress Measurement Tool A), global process conditions (e.g., 3 to 4 minutes of deposition) and sequences thereof (e.g., measure stress, deposit material, then measure stress again). We also see staff querying tool owners as to the capabilities of their tools with respect to a given approach. For instance, Joe asks Jason about the level of stress introduced by an Al_2O_3 film deposited using one of his tools. The next example shows staff performing similar moves, but in the context of a user's question about choosing the appropriate etch chemistry for her material stack.

Example 2: Selecting etch chemistry for a material stack

In interviews, users report that a common reason for asking staff for help was when they were trying a set of process steps that they hadn't tried before. They phrased such queries as wanting to ascertain "whether something will work" or "what will happen if I…?" Kim, the user in this example, brings such a query to the weekly process meeting at the Gamma site. She is building her structure on a wafer made out of quartz, onto which she has deposited a layer of aluminum nitride (AlN) topped by a layer of silicon dioxide (SiO_2). She asks how to best etch the silicon dioxide layer. Jason responds by naming two most commonly used etch chemicals (hydrofluoric acid, or HF, and buffered oxide etch, or BOE), but then poses the question to other staff members whether either of these chemicals will etch the user's underlying AlN layer – an undesirable occurrence given Kim's process goals. Tim, who is responsible for a tool that deposits AlN, responds that these chemicals will not affect the AlN layer.

> Kim: Yeah, my next question is about the next steps. So for the silicon oxide etching, how can I do that?
>
> Adam [gestures to the screen where Kim has projected a diagram of her process flow]: So that's at step where?
>
> Jason: Twelve. I don't know what HF or BOE does to aluminum nitride
>
> Tim (tool owner): Nothing.

Having established the in-principle fit of the approach of using HF or BOE as regards Kim's material stack, Jason then raises the issue of the underlying quartz wafer being susceptible to HF or BOE. Tim acknowledges that this is the case:

> Jason: Nothing? It's a quartz wafer.

> Tim: Oh shh

Adam begins to probe the Kim's fabrication tolerance for the quartz being etched slightly by BOE or HF:

> Adam: It is a whole wafer, right?
>
> Kim: Yeah, whole wafer

Tim says he's used BOE to etch SiO_2, further bolstering the feasibility of using BOE for this step:

> Tim: I've done processes where we do [*BOE*] on the top oxide

Adam continues to probe the user's fabrication tolerance for the quartz being etched slightly by BOE or HF. He specifically asks whether the user will mind whether the optical properties of the wafer suffer:

> Adam: That quartz wafer, there's no optical property to it? If it gets frosted up or becomes not so transparent, it's okay, right? It's just a big [*electrical*] insulator for you, right?
>
> Kim: Yes. Yes.
>
> Adam: So it might not matter.

Adam's query has surfaced the fact that even if the quartz substrate is susceptible to being etched by HF and BOE, the slight amount of damage that will be done in the course of the etching of the SiO_2 layer will not harm the intended functionality of the device, because this functionality does not depend on the optical properties – the roughness – of the quartz. Rather, the functionality of the device depends on the electrical properties of the quartz, which are not materially affected by being etched slightly.

Jason then reiterates that using BOE or HF would be easiest. Adam then proposes a concrete course of action, emphasizing that this approach is feasible – he is confident it will remove the SiO_2 – as well as steering clear of any problematic ramifications for the properties that the quartz needs to have for the success of this user's process:

> Jason: The easiest thing is just dip it in HF or BOE.
>
> Adam: Yeah, straight HF, it's going to go away like that [Adam snaps fingers]. And so just have straight HF [*and*] two dishes of water – bang. 10 seconds. Water, water, spray it off. And it probably won't do much damage to the quartz.

Brief synthesis

In the example above, staff have probed the user's fabrication needs sufficiently to establish that although HF and BOE will slightly etch the quartz wafer, this etching is tolerable. Relying on their own experience with BOE and HF, several staff members confirm that these chemicals will remove the SiO_2 layer and not the underlying AlN layer. With this, they have matched a situated fabrication need with the capabilities of a cleanroom technique.

Example 3: Refining a user's fabrication need

Sometimes users respond to staff's proposed approach by raising problems with it, which provides staff with more detailed understanding of the user's objectives and constraints, and therefore the fabrication need to be met. In the following example, a user and staff have been discussing a project that will involve significant work on aligning layers on a sample in such a way that a pattern can be placed on the backside of a wafer in precise alignment with a feature on the frontside of the wafer. Up until this point, the user has referred to the pattern she wants to deposit as a homogenous pattern, which her technical drawings indicate must be curtailed to a precisely specified region on the wafer in order to be aligned to the frontside feature. Adam probes the user's fabrication need by asking whether the patterned area needs to be limited to the area that the user has delineated in her technical drawings:

> Adam: Why does the [*patterned area*] need to be so discretely sized? Could it not spill over that whole backside?

The user responds in the affirmative:

> User: Oh, it can. Yeah. It could spill out here. Is that going to make things easier for fabrication?

Jason responds by proposing that the implication is that this eliminates the need for alignment of the front and backside – because the pattern can be cut to size later in the process. The user responds to this proposed approach by saying it will not work, due to properties of the pattern she has not yet told the staff about. The user then provides more detail about the specific properties of the pattern he wants to create on the surface, explaining why the process will still require alignment, even though the pattern can spill over a larger area:

> Jason: Then you don't have to align it.
>
> User: No, you still do. So, the [*pattern*] [*has properties that make alignment necessary*]
>
> Jason: Ah
>
> User: [*Elaborates on the specific shape of the pattern*] I didn't think it was relevant for fabrication, but I should have included that. [*Elaborates*] So making it larger is not going to solve our problem. We still need to make these two [*patterns*] aligned to each other. Does that make sense?
>
> Adam: Yes.

Staff then proceeds to explore other approaches to aligning the two sides of the wafer. Adam probes the user's fabrication constraints further. Specifically, he asks whether there is space on the wafer to place alignment marks. The user responds by specifying the space available for alignment marks:

> Adam: So, could you, if you had, out on that perimeter, that little edge there, a permanently-placed, deposited metal cross?
>
> User: That's not a problem.
>
> Adam: That's okay?
>
> User: Yes. Anything in this area here can be used for markers.

Adam then proposes an approach for attempting front-to-back alignment, to which Jim (another staff member) adds detail before Adam invites other staff members to comment on the feasibility of the approach:

> Adam: So if you took this [*sample*], after [*step A*], and you came to a [*lithography tool*], and [*you make alignment marks*], and then took those marks, and then measured them on some other tool to know where they're placed relative to [*feature A*]
>
> Jim: on the [*Measurement tool A*], image that
>
> Adam: Gather image data there. And then make corrections for how you do the [*next step*] based on that. Can that work? I'm kinda tossing it into the room and everyone
>
> Jim: Yeah, I think that sounds good. Because you also have to [*do additional step*]
>
> Adam: Yeah. You can do that with two [*lithography*] steps and [*metrology*] in between. And with a lift-off of a [*metal*].

Brief synthesis

By proposing different approaches and using these to probe the user's project objectives and constraints, staff members have developed a clearer picture of this user's fabrication needs. In fact, they have surfaced information that the user did not include in her original specification of the structure to be built. By querying other staff on the feasibility of a concrete approach, this need is mapped onto a set of steps that staff deems to be feasible, in principle. The approach that Adam proposes is concrete, referring to at least five different tools and nine different steps. The result of the conversation is an in-principle feasible approach that suits the demands and constraints of the user's process.

Example 4: Matching fabrication need and tool capabilities

Todd, the user in this example, has come to Gamma's weekly open user meeting to discuss his device, which will involve thick layers of [*Material A, hereafter simply A*] and [*Material B, hereafter, B*]. Throughout the portion of the conversation featured here, staff probes various aspects of his device's fabrication needs. The conversation shows the interchange between proposed approaches, fabrication needs, and tool capabilities as the objectives and restraints of Todd's process become increasingly clear in response to the capabilities and limitations of the specific deposition tools that staff discuss. The main equipment limitation in question in the following portion of the conversation is whether the deposition tools under consideration can deposit layers of A and B at the thickness that Todd's design drawings call for. The portion of Todd's process objectives and constraints that are being explored are whether the desired thickness of A and B might be achieved by other means: for instance, by alternating thinner layers of A and B. Throughout the conversation, staff proposes concrete approaches to fabricating the layers of A and B and invites Todd to comment on their fit with his process objectives and constraints. Over the course of the conversation, Todd prioritizes the properties that the device will ultimately exhibit.

Discussion dwells on the thick layers of A and B that Todd is planning on depositing. Noting that Todd is asking for a deposition of 5 microns of A, which is considered a thick layer, Jason asks Tim, the tool owner of the deposition tool in question, about the tool's capabilities in this regard:

> Jason: The only question for Tim is how long does it take to put down five-micron A in the [*Deposition Tool 1*]?

In response to this question Todd indicates some room in the device's fabrication constraints: the A and B can be layered, so that the individual layers are not as thick as originally proposed, using less material in each deposition run and thereby staying within the capabilities of the tool in question:

> Todd: Well, I'm not opposed to this being layered with A and B. The more continuous material, the better. But I'm not opposed to seeing a two-layer: however much maximal thickness of A you can [*deposit*] and then maximal thickness of B.

Tim responds by stating a tool limitation regarding the maximum amount of material that can be loaded into the tool (in the form of metal discs known as "targets") for any one run:

> Tim (tool owner): So in terms of the [*Deposition Tool 1*], we're using two-inch diameter targets that are an eighth of an inch thick.
>
> Joe: So how many do we need?
>
> Jason: How many targets you got? [laughter]
>
> Tim: Yeah, honestly, I don't know that there's five microns of depositable material in a single target.

Todd responds by describing some additional space in his design objectives that could be used to accommodate the limitations of the tool:

> Todd: Okay, one possibility is I can reduce the amount that we etch the silicon here. And I can actually make this more like 2 microns, and we can layer A and B as long as it [*still meets condition C*].

Tim recalls that a different deposition tool could deposit both materials A and B in the same chamber, a statement that he follows by a proposed approach:

> Tim (tool owner): Oh, [*Deposition Tool 2*] I could do A and B in the same chamber.
>
> Todd: In the same chamber, but sequentially?
>
> Tim: Yeah, you can switch between the targets in the tool. I could put them both in the same tool, so you could deposit under the same vacuum.
>
> Todd: So would two microns be possible with that?
>
> Tim: Sure. Yeah.

To this, Jason raises the question of film stress, a comment that is not picked up by others in the room until later:

> Jason: They're both stressy films

Instead, Tim continues to articulate the feasibility of the deposition step using Deposition Tool 2:

> Tim (tool owner): You could probably do, with reasonable certainty, I think you can do 2 microns of A and 2 microns of B. I think with reasonable certainty, I could say that.
>
> Todd: Okay, however much we can get is, the more the better. And as long as [*the overall structure has Property A*].

Next, Adam proposes another approach to depositing thick layers of material: electroplating the structures using an electroplating bath, a tool of which he is the tool owner:

> Adam (tool owner): You could think about, the A can be plated, but we'd need a seed of some sort. So you could just [*deposit*] a seed with Tim's tool, and then we could plate up A to 10s of microns if you please. The B, though, at least as far as I can think of it, you're kind of locked into doing that with Tim and [*Deposition Tool 2*]. But there may be some varieties there.

Adam follows this with a query into Todd's process constraints:

> Adam: Have you modeled or simulated different [*thicknesses*]? You're saying the same thickness for the A and B? Could they be like more A, less B? Or do they have to be in these ratios?
>
> Todd: That's a good question. I haven't looked into a layered material like this. However, I think because we're going to primarily rely on testing for [*Property A*] experimentally. So we could do variations of this and see what works best. I'm not – you know, it's possible to use
>
> Adam: Multiple layers? Thinner layers?
>
> Todd: What we could do is we could start with doing some reasonably equal thickness combination of A and B and [*for*] later fabrications, we can try just one or the other, and then see what [*Property A*] is. I'm not particularly eager to model for discontinuity of these, because there isn't that much information on this particular thickness will interact with [*Property A*].

Tim asks Adam, who is responsible for the electroplating baths, about the properties of films deposited in this way:

> Tim: Adam, what's the grain development like with electroplating? If you do a seed layer of A and then you deposit, electroplate two microns on top, [*could that*] get really big?

Adam responds by discussing the aspects of the plating bath that affect grain size, and then sketches a possible approach Todd could take:

> Adam (tool owner): They can. It's a pulse plating setup. So it's entirely dependent on how long we have the forward current on, and the balance of forward to reverse pulses determines the final morphology of the film. [*Elaborates*] So you can mitigate that by just being on for shorter periods of time before you do a reverse pulse. There's a lot of things you can do. [*Elaborates*] There's other flavors out there. [*Names more options.*]

Adam then asks Todd more questions about the thickness of the A and B layers:

> Adam: Todd, do you ever think, and maybe this helps Tim out better, could you do thinner layers and more of them? Like go a couple 100 nanometers and just do like A,

> B, A, B, A, B, a whole bunch of different layers and then as you extinguish a target, throw another target in there, and it's not like using a target right till you're [*damaging the tool*]?

Todd responds by stating a refined fabrication need:

> Todd: It's a good question. From what I understand of this, it's not something that's frequently available in the literature, having a discontinuity in [*feature*]. But I think the more continuous the thickness, the better. So if we could minimize the number of thin sections that we have of this – it is better to have maybe two or four layers of AB, rather than several of them. Because I think it'll [*mess with*] my [*Property A*].
>
> Adam: Okay.

Mark, an equipment engineer, queries Tim about [*Deposition Tool 2*]:

> Mark: [*Can we put*] multiple B targets in the instrument?

Adam replies and asks Tim about a subcomponent of the tool related to Mark's question:

> Adam: Possibly. I mean, these are all DC, right Tim? [The current type of the power supply (AC or DC) limits the type of targets you can install]

Tim responds in the affirmative, but also raises an aspect of the tool's setup that constrains the thickness of the targets that can be installed in the tool. He then proposes a way forward, which involves installing fresh targets right before Todd arrives and monitoring their depletion during deposition:

> Tim (tool owner): Yes. We're limited by – you can't put a normal quarter-inch thick target in the sputter tool with magnetic materials because the magnetic field lines don't leave the surface of the target. So you've gotta go with a half-thickness target. If you let me know before you're coming to town, Todd, I'll put brand-new A and B targets in the tool, so that we'll be able to maximize your deposition capability in terms of thickness.
>
> Todd: Okay, for sure.
>
> Tim: I think we can do it. If we have to, if we deplete a target, or we're getting towards depletion of a target, I can swap it. We just want to see that coming. So we'll

> monitor the voltage [*during*] deposition and see if we're getting toward a place where it looks like things are gonna fail.

Supporting interview data

I have argued that the interactions between staff and users reported above support the models depicted in Figure 4 and 5. Interviews further support this interpretation in three ways. First, interviews suggest that there is a discrete role played in these conversations by staff members who negotiate alignment between a user's process objectives and the capabilities of tools by proposing concrete approaches that lead users and tool owners to clarify both. Second, interviews suggest that one reason that probing a user's fabrication objectives and constraints is necessary is because it is often initially unclear to staff how well a user's stated fabrication need is grounded in knowledge of fabrication techniques and the cleanroom's tools. Third, interviews suggest that refining the user's situated fabrication need by exploring their process objectives and limitations is crucial for staff to be able to generate alternative approaches. Because users often have only a general sense of what the tools in the facility can do, they are not aware of the options open to them. Therefore, a key element of staff's knowledge of nanofabrication is the ability to generate alternative approaches to a given fabrication need. Interviews suggest that this expertise is relevant even in interactions with experienced users.

Different staff roles

In a later interview, Mark, a staff scientist who regularly attends process meetings, describes what happens in these conversations as follows. He describes these interactions as a search for a fabrication approach that is both suitable to a given user's unique process objectives and constraints and that is feasible given the lab's tools. A key role in finding this match is that of a staff member who queries equipment staff about specific fabrication approaches (in Mark's example, this is Rick), and then uses their replies to make feasibility pronouncements regarding the proposed approach. He describes Rick's role as identifying the facets of a user's process to examine and formulating targeted queries to the tool owner. Mark describes the tool owner's role (Jason, in this example) as supplying information about a tool's capabilities (etch selectivities), which Rick uses to make a pronouncement on the feasibility of a given approach. In this quote, Mark compares a user's overall process to a forest, and its steps to individual trees, and Rick therefore as moving the discussion between the two.

> I think Rick is a good forest-view person. And then some of the other [*staff members*] in the room are good at trees. And [*Rick*] goes, "Okay, I know it's *that* part of the forest [*that we need to talk about*]. What's the selectivity of that specific etching process?" And Jason goes, "Oh, it's 3 to 1." And Rick goes, "Okay, if it's 3 to 1, that should work."

Mark's description suggests that an important purpose that these meetings serve is making this match between user objectives and tool capabilities.

Why probing user's process objectives and constraints is often necessary

The reason why probing a user's fabrication objectives and constraints is necessary is because users tend to state their fabrication needs with equal confidence regardless of their foundation in experiential knowledge of how the cleanroom's tools work. As Mark describes:

> Oftentimes, it's hard to know when someone comes in and says, "Here's my work. Here's the process that I have written out, here's my 10-step process," it's hard sometimes to know whether that is coming from a place of knowledge, or a place of, "This is what my boss told me."

Probing process objectives and constraints enables a search for alternative approaches

Likewise, users often model their original fabrication plans on published papers, which are often based on fabrication conducted at different facilities that have different tool sets. This lack of rootedness in the practice of fabrication within the cleanroom of a given facility means that users often do not realize that they have many more options open to them than a single paper would suggest. This is something that Roger, an experienced equipment engineer indicates in the case of depositing a thin film, which can be done using a tool class known as furnaces and a tool class known as PECVD:

> We have multiple ways of putting thin films down on samples. Furnaces are just one of them. They have some advantages and disadvantages, because they're a higher temperature process. So maybe they don't need it. They're also slower in some cases, so maybe they can go with something a little faster, like a PECVD tool. So a lot of times, [*users will*] say, "Someone told me I needed to use a furnace film for this oxide." Or, "I read in a paper, there's a process flow that somebody had published and they used furnace films." And we could look at it and say, "Okay, well, for what you're doing, you probably don't really need the furnace film, and you could probably just use the PECVD film."
>
> That's something that's common. Because they're not familiar with the whole technology and the area. Maybe they're a physicist or a biologist, but they got a paper that said, "these are the steps I gotta run." And then it's sort of our job to say, "Okay, we could rework this and we could save you some time, or you get a better result if you use this step instead of that step." Just because somebody out at [*University A*] used a furnace, they may not have had a PECVD as an option when they did it, so they just used the tool that they had available. We got a little bigger set of tools, [*so*] we can say, "Well, you get better results if you do it this way." So that's our job, is to basically help guide them, streamline their process flow and kind of get them on track in terms of a better result and a quicker result.

This indicates that an important aspect of staff's knowledge is a knowledge of alternatives given a specified fabrication need. This probing of a user's process objectives and the subsequent generation of alternatives is something Mark emphasizes as he continues his description of what happens in process meetings:

> And I think when Rick is there [*at process meetings*], it happens more often that the question is, "Why is this [*proposed process*] like this?" And that's when [*the user*] says, "Oh, well –" if they have an answer, then that's great. If they don't have an answer, then sometimes [*staff*] start to go, "Well, how else could we do this thing? The only thing you care about is that [*your structure*] has holes in it at the end. How

> best could we make holes that are that size, this dimension, this, that and the other thing?"

In this statement, Mark describes how in the course of these conversations, Rick asks questions that lead to the fabrication needs of a user becoming increasingly defined. As the situated fabrication need gains definition – in the stylized example Mark gives, "the only thing you care about is that it has holes in it at the end" – then staff can generate alternative approaches given what their tool set can do: in this case, how best to make holes that have the dimensions and properties that this structure requires.

Jason provides a series of cooking and baking examples that illustrate the search for alternative approaches to reaching similar outcomes:

> I think it's like baking or cooking. Where if somebody comes in and [*says*], "I want to make cookies out of Crisco!" and then the cook's like, "You could use butter. Would that work, if you could just use butter?" And then the cookies are delicious.

In the quote above, Jason describes alternative oils used for baking cookies. In the next quote, he describes different techniques for cooking bacon, as well as a response to a user's specific objectives:

> Cooking bacon in a pan is a really good idea. Can you do in the oven? Yeah, there's some pros and cons to that. Can you do it on the grill? That's not the best way to make bacon. Like you probably can... [*Someone might say,*] "I tried it once and the bacon was really good, but I started a little grease fire," so I mean like, ehh, you know, just put it in a pan.
>
> "– But it was a hot day, I wanted to do it outside…"
>
> "Put your pan on your grill, that might work. Do that next time."

Jason also gives the example of boiling eggs to a desired level of softness. In some cases, using an alternative tool makes sense, such as using a sous vide:

> It's hard to get it timed right. And your altitude matters. Or you use a sous vide, "have you ever heard of this tool?" There's always another way.

Now referring to examples from the cleanroom, Jason makes a similar point: that given a specified outcome, there are almost always alternative ways to get there. His examples refer to the goal of depositing a layer of material in such a way as to avoid problematic film stress (see Chapter 3 for a definition of film stress). Users will often come in requesting a thick film to obtain a given device property, and staff will explore whether the user can tolerate thinner films instead to generate the desired property and avoid problematic stress:

> A lot of people are just like, "I want three microns of metal." And we're like, "Uhm, could you use 100 nanometers? Would that work?" And that's not as much of a problem [*in terms of film stress*]. It's rare that it *has* to be three microns thick, at this one thing and it's the only way to make it. […] And there certainly are those problems, but most of them there's a different, there's an easier way. You just have to think about it and ask more questions and come up with an easier way.

Even experienced users need this kind of interaction with staff, where staff probes various constraints on the user's stated process objectives:

> They'll tell you, "I need this and this and this." But certainly, like in [*our process meetings*], we'd say, "Well, does this matter? Can you get away with, does it have to be [*a*] really good conductor?" Or with metals it's questions like, "How much voltage are you going to apply?" Or "how much current do you need to pass through this device?" We'd ask those sorts of questions. And then we'd say, "Oh, if it's 10 volts, you only need like 30 nanometers of oxide. It's gonna be fine. Yeah, just just go thinner."

But sometimes a thicker layer is necessary, given the user's objectives:

> But if they said, "I need a 1000 volts," we'd be like, "Uhm, you need microns, all right, okay. All right. Let's figure out how to get you microns." So they might not know all of the implications. They know they need to measure this small of a current or create this much of an [*effect*]. They know those kinds of end-use parameters that determine the material requirements [*or*] the geometries or whatever. And so, higher voltages, higher currents, those obviously are harder to make, and so they're gonna have more restrictions and they have to consider more things to make it work properly.

The difference made by the quality of users' participation in problem-solving

A core feature of process development meetings is that staff will query users' process tolerance given the properties of a potential approach. Users' openness to these probes into their process shapes the kind of help they receive, because it affects the discovery and refinement of their situated fabrication need and the matching of that need with precise tool capabilities. Users who do not share their processing objectives with staff must make this match largely on their own, deciding for themselves which features they need to build and which tools and techniques they will use to do so. An example of reticence in this domain, and, in the eyes of staff, a missed opportunity for the user, occurred during a process meeting I observed at the Beta facility. Staff discussed among themselves a meeting they had just had with a user:

> Arnold: That first user, [*Name*], he's come to us a couple times. I don't think I fully understand his process, what he's trying to do.
>
> Kate: You know, I'm not exactly sure what they're [*analyzing*]. Their lab is pretty secretive about a lot of stuff. I think the challenge in helping them is that I'm not sure it's entirely clear that their process needs are mapped out. Like, their device needs. So, they ask very specific questions about this [*technique*]. And then if that doesn't work, they go on to the next thing. Whereas I think what they really need is some help with a more holistic approach. I'm not sure [*that*] if they get the [*precise feature they asked for*], they're going to get what they want.

Coming to grips with an integration challenge often necessitated users giving staff the insight they were requesting into a process. Some users were not willing to provide this kind of access without assurances of confidentiality in place. As Kate's comment suggests, users who chose to develop their processes without staff's advice on integrating their whole process often formulated their own fabrication needs and then asked staff very specific questions. It is likely that users' success at developing a complex, integrated device in this way varied strongly with their own level of experience at fabrication.

Similarly, Roger at Gamma remarks, after describing how he can help users identify better alternatives when he understands their process needs:

> Sometimes people are, I guess, a little more protective and closed off. So they'll just say, "I just need a furnace film." That's all they're gonna tell you. Then, okay, I'll train you on the furnaces, right? And then you may find out weeks later, it's like, well, you could probably get a better result if you used the PECVD. So, we do need people to share that information, so that we can help them. The more they can share with us, the more we can help them.

Conclusion

These two chapters have served to document the ways in which staff ensured that users' process steps were integrated: that no process step hampered the functionality of the entire device. The observations documented in this chapter illustrate how process integration was achieved over multiple interactions with users and among staff, often over the course of a project.

This chapter has shown a few forms that process integration takes in action. In so doing, it has shown that staff worked together to help users build well-integrated structures. The tools of the cleanroom were so specialized that no single staff member knew them all well enough to ensure that a given approach was well-supported by known tool capabilities. Instead, knowledge of tools was distributed across tool owners. In addition, these capabilities of tools were brought into the conversation by staff members who negotiated alignment between users' goals and what the tools could do. The staff members in this role proposed approaches – sequences and combinations of specific materials, tools and techniques – and then tested for their fit with users' objectives and tool capabilities. By documenting the work that staff members in this role performed, we begin to get a sense of how the staff collectively identified and met *situated fabrication needs,* one of the elements this dissertation argues is a part of working knowledge.

It could, of course, also occur that no match was found between the objectives of a user and the capabilities of tools. For instance, it could happen that the user's objectives were such that no path through the cleanroom's material restrictions could be found. Any of the integration challenges identified in Chapter 3 could likewise be prohibitive. In addition, it was also possible for the user to withhold sufficient detail about their fabrication objectives, which also hampered the precision of the match that could be found in these conversations.

As this chapter has illustrated, a key role involved in the work of identifying and meeting situated fabrication need was the role of the tool owner. Matching the objectives of users with the capabilities of the facility required distinct expertise regarding what the tools could do: knowledge of specific tools that was responsive to the queries posed in process integration discussions. How staff developed this kind of expertise is the topic of the next chapter, about being a tool owner.

Chapter 5. Being a Tool Owner

> "[*Staff*] will give you the basic training, get you started and get you going. And they'll unstick your wafers when they're stuck in the tool. But if you need to get deeper into a process, and get it to higher levels of precision, they will sit and work with you, and teach you how to really make the tool sing."
> – Experienced industry user, Gamma site

> "[*Users*] do all kinds of oddball things. They do different materials, they do things the machines aren't designed to do. So when they do that, we learn something from them. We learn: here's a capability we do or don't have. Or: here's what happens when you do this."
> – Experienced equipment engineer, Gamma site

Introduction

As we have seen in the previous chapter, tool owners played an important role in process integration discussions because they responded to targeted queries about their tools' capabilities to support the fabrication approach that was being discussed. In their responses, tool owners often also raised additional process integration challenges that the use of their tools introduced and formulated further probes into the user's fabrication objectives and limitations. In so doing, tool owners bolstered, weakened or qualified the feasibility of the fabrication approach under consideration, and thereby aided in the discovery of a user's situated fabrication needs and the fine matching thereof with the facility's tool set.

Equipment staff's ability to make these responses, pronouncements and queries rested on their long experience in the role of "tool owner." My goal in this chapter is to document the core features of this role and to suggest how experience with its recurrent tasks and problems fostered the development of distinct occupational knowledge. Specifically, I'll argue that tool owners developed knowledge of their tools' capabilities that was responsive to the situated fabrication needs of diverse projects. Recall that users' fabrication needs were unstandardized: each project typically represented a unique combination of device objectives and integration constraints. Pre-set combinations of tool parameters (“recipes,” as they were called) were therefore only of partial use. In the absence of standardized needs, tool owners had to gain a sense of the process possibilities that their tools represented and to manage the application of these possibilities to specialized user goals.

Tool owners learned what their tools were capable of by performing their role, which was composed of two main parts (Figure 6). First, tool owners engaged in a number of activities that kept their tools operational (such as maintenance, repair and troubleshooting) and through this work they learned the physical mechanisms and subcomponents that supported and circumscribed aspects of tool performance. Second, tool owners helped users configure a given tool for their specific processes, and through this work they learned to tune aspects of the tool for finely differentiated ends. In so doing, they witnessed the range of outcomes the tool was capable of producing. These two aspects of their role enabled tool owners to effectively map what their tools were capable of, and therefore to comment on the feasibility of a given approach even when they hadn’t seen the specific fabrication need before.

Moreover, tool owners’ knowledge of what their tools could do evolved as they encountered new use cases for their tools. Even tool owners who had managed their tools for decades noted that they were “still learning” about what their tools could do. Tool owners described learning about their tools when they were prompted to work with novel inputs (such as different substrates or wafer sizes), to fabricate features in the context of unprecedented process integration constraints and to develop finer control over a particular dimension of tool performance. Tool owners met these novel demands by developing new recipes with users, honing their familiarity with aspects of tool configuration that affected a given outcome, and sometimes by making modifications to tools. From these experiences, tool owners gained further experience with what their tools could and could not do.

Figure 6 serves to illustrate the relationships I make in this chapter and the next chapter. Specifically, this chapter will discuss the parts of the figure that are typed in bold font, while the next chapter will discuss the other parts of the figure. I will discuss the two main aspects of tool owners' role, documenting recurring tasks and problems. I argue that experience in this role laid the foundation for distinct occupational knowledge related to the capabilities of their tools. Finally, two implications follow from the role and occupational knowledge of tools owners. The first implication is that tool owners enabled the fine matching of users' objectives and the capabilities of tools. In so doing, the expertise of tool owners underwrote the collective working knowledge of staff documented in the previous chapter.

Figure 6 about here

The second implication is that the way tool owners knew their tools allowed for their facilities to keep step with emerging research frontiers. Over time, through the many interactions that tool owners had with users around their processes, their labs became better suited to the research fields represented among their users. For instance, in this chapter we see a tool owner's efforts at making his tools hospitable to a novel, but ascendant, sample type – samples that come in small pieces, rather than in full wafers. His efforts at reconfiguring his tool to create this capability align with a trend in nanoscale research: a growing number of users no longer process only silicon, but have shifted to exploring more esoteric materials for which full wafers are too expensive to procure. This tool owner's adaptation of his machine beyond its initial specifications, completed over the course of months in the context of a single user's process, rendered a core cleanroom tool better suited to the fabrication needs of several research fields. One could argue that the ability of these emerging research fields to pick up steam relies on these ostensibly mundane development efforts between staff and users, where staff extend what they know about their tools, and thereby what users can do with them in their R&D efforts.

Helping users extend the capabilities of tools was especially important given that these cleanrooms operated expensive and ageing machines to perform work at a continually shifting set of research frontiers. Many of the tools in these facilities were second-hand to begin with, and the average age of these machines lay in the double digits, with the oldest tools pushing thirty years. The final implication of tool owners' knowledge for the capabilities of their facilities was, therefore, that it kept old tools in conversation with novel applications. Consequently, staff's ability not only to keep these tools going, but to adapt and optimize them to serve the research frontier of the given moment was of more than a little importance.

What tool owners did

Tool ownership, across facilities, involved a fairly consistent set of tasks, geared toward two key domains: keeping a tool operational and guiding users on its tailored use for their process. In the words of an equipment engineer at the Gamma facility, his role encompassed "knowledge and maintenance."

Keeping tools in working order

Maintenance, repair and upgrading

In the first domain, keeping a tool operational, tool owners performed a number of tasks that could be considered maintenance of the technical core of a tool: servicing the tool's main subsystems, performing periodic rebuilds and cleanings of subassemblies, replacing worn parts, and repair. Periodic rebuilds of subassemblies gave tool owners the opportunity to replace temperamental or delicate components for more robust ones and to perform other "upgrades" to their tools. Senior tool owners also took the opportunity of these teardowns and rebuilds to instruct junior equipment staff on the anatomy of their tools. When repair involved extensive technical troubleshooting, staff gained familiarity with the role that parts and their interactions played in the tool's functioning, as Roger, an equipment engineer responsible for a major class of tools in the Gamma cleanroom indicates, describing how he has gotten to know his tools over the years:

> It's when something breaks on a tool. Whether it's a furnace, an etch tool, lithography tool, when it breaks, you've got to fix it, right? A lot of times, there's some troubleshooting you got to do and figure out why it's not giving you the result you need. So then you understand the mechanics behind it, and why it's running.

In cases when a staff member had been present for the original installation of a tool in the facility, the staff member could say that with the exception of the equipment vendor, they had been "the only one to have turned a wrench on this tool." As tools often remained in use for decades, experienced tool owners were therefore well-acquainted with the individual repair histories and idiosyncrasies of the machines they maintained. Maintenance contracts with vendors were typically reserved for the newest and most advanced tools in the cleanroom. For staff who owned these tools, being a tool owner involved interfacing with vendors when these came by for periodic maintenance or service calls and engaging in remote troubleshooting.

Restocking and reconfiguring

Tools often had a limited number of slots for materials and chemicals to be used in processing. Users' process needs spanned a larger set of materials. In their daily or weekly routines, tool owners also responded to user requests for the installation of specific materials into tools. For some of the most heavily-used tools, this was a daily or weekly occurrence. This meant ordering, storing, installing and de-installing the various common and more idiosyncratic materials and gases that users requested. In addition, some tools could operate in multiple distinct modes which required the tool owner to adjust aspects of the tool's internal configuration. For instance, one of the largest photolithography tools in the Gamma cleanroom could run multiple sizes of wafers, but changing from one size to the other required a day of work from its tool owner.

Preventing minor breakdowns and bringing tools back up

In addition to maintaining the tool's technical core, tool owners performed other caretaking tasks that were less technically demanding but nonetheless crucial to keeping a tool operational. Some of these tasks involved guarding against an array of known, but minor, breakdowns that did not mortally wound a tool but nonetheless removed it from operation – that "brought the tool down," in the local jargon. Tool owners worked to prevent such breakdowns by regularly inspecting and servicing parts of their tools known to be vulnerable. For instance, I observed Joe, an equipment engineer at Gamma facility make his "daily rounds" across his tools, which involved using a spray bottle of solvent to moisten the chemical dispensing nozzles on a tool, to prevent them from drying out and clogging. On another large, decades-old and mechanically complex tool, Joe’s rounds involved assessing the state of the tool’s wafer stage by carefully moving it back and forth in its different directions of motion, a task he referred to as “exercising the stage.” Not all minor breakdowns could be guarded against, however, and staff therefore spent time responding to users who brought a tool down and restoring a tool to functionality.

Across facilities, staff did their best to avoid these vulnerable spots being accessible to users, by asking vendors to make modifications to interfaces, by physically obstructing certain parts of screens with plexiglass, by posting signage, by emphasizing the correct use during training and telling stories meant to deter users’ erroneous use of the tool. There nevertheless remained numerous opportunities for users to bring a tool down, and therefore, tool owners regularly received calls to bring tools back up again.

Training users

Making and keeping a tool operational also involved training users on the tool’s operation. New batches of users came in at regular intervals, and for owners of the most popular tools, this often meant running at least one training a week. Training involved familiarizing users with the basic components of the tool, the order of operations for its basic use, key operating errors to avoid, and the materials restrictions in effect for the tool. During training, tool owners often asked users about aspects of their processes that were relevant to the tool’s operation, such as whether the user was running whole wafers or “piece parts” (which informed how the sample would be loaded into most tools) and what materials they were working with. Staff would run, or would supervise users while they ran, a basic recipe. While the recipe ran, the staff member narrated the basic operations that the tool was performing. Depending on the technical complexity, monetary value and delicacy of the tool, as well as the user's level of experience and chemical hazards involved, training could last anywhere from thirty minutes on a single day to several hours over the course of several days, involving several instances of supervised operation and a test of the user's declarative knowledge.

Writing documentation

Tool owners further instructed users in how to use tools by writing documentation with instructions. These documents were distinct from the manuals that tool vendors sometimes provided, which staff considered too technical or poorly written for a typical user to use competently. For instance, Joe describes the work he'd done to develop a user manual for another one of his tools that made programming the tool ("writing a job," in the local jargon), easier to do:

> For the [*Stepper*], writing a job for it is awful. So when we bought the tool, [*the vendor*] sent me, on a CD, their training program for writing jobs on the tool. Well at that point, I'd been writing jobs for steppers for [*20*] years. I went through the whole thing, I went to the tool, [*and*] I couldn't write a job. Because it's just that opaque. So, after I got good at it, I sat down and wrote – it's in the front [*of the binder*] there – several pages: my cheat sheet. If you follow these steps, you can write a basic job. And that's worked pretty well. People are able to read through that and make a simple job.

Gathering reference materials

In addition to writing operating instructions, staff also gathered and made available technical reference materials regarding basic chemicals and materials needed for navigating a given technical domain. For instance, photolithography involves applying photoreactive chemicals known as "photoresist" to the surface of a sample. Coating a sample with photoresist is most commonly done by adhering (usually by vacuum) the sample to a rotating chuck and spinning it at high speeds while dispensing the photoresist from above, allowing the centrifugal forces to ensure an even coating of the chemical across the sample. This technique is called “spinning resist.” There are hundreds of different types of photoresists, differentiated on several dimensions, available from different vendors, known by a plethora of overlapping trade names and chemical names. Each required a distinct set of spin speeds, bake times and temperatures to have its designed effect. Staff made available to users information necessary to select, spin, bake, expose and develop many different resists, in an effort to facilitate the comparison and selection of the appropriate resists for their specific application.

Developing baseline recipes and characterizations

Maintaining technical reference materials served another important domain of work staff did to make and keep a tool operational: developing what were known as “baseline recipes” on tools. These were publicly available parameter settings that had been developed with broad classes of applications in mind. Staff referred to these as “starting points” for users, and staff typically first referred users to a baseline recipe before engaging in more targeted tuning of a recipe for a user’s application (see later section on process development). For instance, Joe refers to this genre of documentation in his domain as follows, beginning with the technical reference materials previously described:

> If you look in the [*photolithography area*], there’s three black books. There’s one big one. And that has all the vendors’ datasheets for the photoresist. Which are, believe it or not, extremely difficult to get. Also in there is other information, like which resists you should use on which tools, baseline processes for dose and bake times and temperatures, developers and time, just our standard house processes for those things. And that gives [*users*] a guideline. They will always ask me, “What dose should I use?” I say, “It’s going to be unique to your situation.” I say, “Go look in the book. There’s a number in there and your stuff could be a factor of two in either direction. But you got to find it. So go do the test.” But that gives them a starting point.

Developing baseline recipes for a tool involved preparing dozens of samples and running them through different parameter combinations and then measuring the result on several dimensions known to be of interest to users. I observed Jonathan, an equipment engineer, while he was in the process of developing baseline recipes on a newly installed tool in the Gamma cleanroom. The tool had been purchased for its capability to deposit materials at an angle, a capability which had the potential to create features heretofore unachievable by other methods. Jonathan's task in the weeks ahead before opening the tool to wider use would be to recreate the features that a paper claimed the tool was capable of, a task he compared to "having a race car, but you don't know how to drive it yet."

> Jonathan: So to give you an idea of what I'm working on. [Jonathan pulls up a recently-published academic paper on one of the nearby computers and scrolls through its pages] This is why we're doing all this stuff, is we bought that tool to do this kind of deposition. So, what you can do, is you can make some really neat structures. Like this [Jonathan gestures to an SEM image in the paper of a fabricated feature], or that, but there's some really cool pictures in here. Like this, see? Which would be pretty cool, huh?
>
> Danielle: Yeah
>
> Jonathan: So I gotta read up on this [Jonathan laughs]. There's a lot to this. It's like having a race car, but you don't know how to drive it yet.

Developing baseline recipes is often done with an eye on matching the variation in features that staff typically encountered in users' structures. In the excerpt below, Craig, a process engineer and tool owner at the Alpha facility, describes the work involved in developing baseline recipes on an etch tool for specific features, which he refers to as "characterizing" a machine:

> [*Characterizing a machine means*] going in and running different processes with the machine where you vary the parameters. So, lots of knob-turning and seeing how that affects the results, so you can get an understanding of how the machine's working

> and what the knobs are doing. So if we have a better idea of how things work in a system, what things do, then we can help the users with their specific work, where they might want – say, if they're trying to shape something: do they want it to be shaped like *that*, do they want it to be shaped like *that*? Basically, to help them be able to get what they want.

Characterization involves varying parameters within a system and observing the effects of these changes on the net results a tool produces, with the goal of understanding the individual contributions that different parameters make to the net results. Here Craig indicates an important purpose that characterization served: to help users get the control they sought over the tools. Staff learned their tools in such a way as to be responsive to the variety typically found in users' fabrication objectives. In this way, tool owners developed their ability to engage in the second main class of work involved in being a tool owner: working with users to tailor tools to their specific fabrication needs.

Guiding users in the tailored use of tools

The tasks described up to this point were all ways that tool owners made and kept their tools operational, ensuring that the tools met a basic level of operation and users were inducted in their basic operation. In addition to doing these tasks, staff engaged in a set of tasks that were more closely tied to the individual processes of users. These were troubleshooting individual users' problems and a set of activities geared toward developing processes on their tools that met specific sample properties and performance objectives.

Troubleshooting problems

Being a tool owner involved troubleshooting problems that users reported on one's tools when these problems originated at the tool. As the chapter on "diagnosing process problems" illustrates, users' problems often originated from a more complex array of sources to be found throughout their process. In these cases, equipment staff typically engaged the expertise of a broader range of staff, including those whose knowledge was geared at diagnosing such “process problems.” But sometimes problems fell squarely within the operation of tools or the process module surrounding the tool. In the example Joe provides below, which he gave while he anticipated his day ahead, he describes the potential sources of a problem that a user has emailed him about, which involve one of his tools (a contact lithography tool) and a particular photoresist:

> I’ve got another guy, he’s been emailing me. I’ve gotta go see him and help him. He’s having an issue with doing contact lithography, which can be funky. And sent me pictures of his process, I have to go look at his wafers and try to decide, okay, what’s causing this? He’s got contamination on the wafer, it looks like. I got to figure out, is it his wafer’s contaminated or is the photoresist he’s using contaminated, which is possible. And he’s telling me he’s using this resist, which is, the date on the bottle is three years old. And for some resists, most resists, it doesn’t really matter that much. This particular resist, which is a negative resist, which there’s only like one of, ages very poorly. So if it’s three years old, it’s not going to be any good. So that could be part of the problem. So I just gotta go [*into the lab*] and hoe through all that and see what he’s doing.

Users reported relying heavily on equipment staff for the interpretation of the unexpected appearance that their samples sometimes took on. An experienced academic user recounts an instance when he relied on a staff member’s expertise on a tool when the wafer he had been running came out looking unexpectedly “dirty”, something that equipment staff Jason and Victor traced to the operation of the tool after they had inspected the user’s samples:

> The [*Etch Tool A*] is another etcher for silicon. And working with Jason and Victor, I had started to see particulates appear on my wafers as I was running processes. And I had interpreted that as something in the chamber being dirty and depositing; like there being redeposition of material that was in the chamber. But then I started to show images to staff and they were like, “Ah, you know what I think this actually is, is the tool, to passivate the sidewalls while you’re etching, deposits [*Material A*]. It’s basically this etch-proof material. And I think the tool is over-depositing that. And the way that you’d be able to tell this, is to do some sort of chemical analysis and look for [*Chemical Compound related to A*] in the particles that you’ve seen. And then also, if

> you HF dip it and it goes away, that's a surefire sign that it's that and not the other thing." And sure enough, [*I did an*] HF dip and it was gone.

Another user described his work with Tim, an equipment engineer in charge of several deposition tools in the Gamma cleanroom, around the goal of lowering the stress in a titanium film he was depositing on his sample. The user described Tim "guiding" him "through a set of experiments" in which the user varied the pressure in the recipe on a deposition tool. When that didn't work, Tim worked on the tool itself to eliminate possible sources of the problem: cleaning the chamber, pumping it down to a lower vacuum pressure, and installing fresh titanium material (referred to as "targets"):

> There was a while ago where we were doing depositions of titanium on one of Tim's tools. And [*the depositions*] were turning out less [*metallic*] than they should. So Tim and I talked, [*and*] it was him guiding me through a set of experiments, trying to troubleshoot that. […] I think he even did some things where he did a chamber clean and pumped out [*the chamber*] to see if we were getting any oxygen in the chamber, and that was basically forming Ti oxide as opposed to titanium. And then also, changing up the targets, to see if there was some sort of target contamination that was leading to part of the problem. And then doing some additional characterization techniques, depositing thinner film, seeing if the rate was different as a function of pressure. So we did a whole bunch of testing.

Developing processes with users

The second way that tool owners engaged with the specific processes of users was in the work of "process development," which meant helping users develop a recipe that matched their sample properties and fabrication objectives. The tools of the cleanroom were in principle capable of accepting a range of inputs and producing a range of outcomes. Process development meant configuring these machines to accept the user's specific sample properties and to optimize for the limited set of outcome measures the user was aiming for. Troubleshooting and process development often went hand in hand, as process development campaigns often raised unexpected results to diagnose. As one staff member put it, process development meant working on a tool with "people that want specific things." Another called it "honing tools to a sharper edge."

Basic process development: Optimizing an existing recipe

Staff distinguished different kinds of process development. The simplest kind of process development is the kind that every user had to do, because their structure was never exactly identical to that of a prior user. Even if users were using well-characterized tools to work with known materials to build substructures known to staff to have been built before, the precise materials layered on their sample (its "material stack") and the topography of its surface would still be slightly different than that of other users. In addition, the performance of the tools themselves changed subtly with use. Each distinctive combination of features required a user to find the optimal set of parameters for a given tool. But in cases where the materials were known and fabrication objectives were in principle routine, finding this optimum was a fairly straightforward task. It involved taking an existing recipe and running experiments to find the optimal parameters for one's particular sample properties and performance objectives.

In these cases, staff helped users by providing them with a previously successful recipe that approximated their process needs and sample properties and by providing guidance on the design of experimental procedure to find the optimal parameter settings. Depending on a user's experience with running such process development, staff may also provide guidance about experimental design, note taking, documentation and strategies for robust processing.

Such optimization took structured and less-structured forms. At the most structured end of such experimentation, users ran what was known across facilities as a "design of experiment" or "DOE." It involved a disciplined approach to gathering data about a tool's performance in a range of relevant conditions in order to find the optimum, as Peter at the Alpha facility explains in the case of etch tools:

> You'll have 10 different parameters in the machine: gas flows, pressures, powers, and you'll start varying them and recording, taking measurements for each one. You'll do this one, vary one parameter, then do this one, vary one parameter and you'll do that 20 times or 10 times and then tabulate them all and say: "OK, so when I vary the pressure, it makes it [*the profile*] go from slanted to undercut. OK, whenever [*I use*] this gas, the polymer gets more rough." It helps you develop an intuition for [*each of the things it does*]. And then if you start running [*recipes*] it actually does what you want it to do. You make a couple of plots: roughness goes up with this parameter; roughness goes down with this parameter; therefore if I want my material to act like this, then I pick the best set of parameters for that.

Peter indicates that the extent of one's DOE depends on the "intuition" one already has about how the various tool parameters affect one's outcome of interest:

> [*The extent of your DOE*] depends how scientific you want to get. Sometimes you just vary parameters until you get what you want. And you just go entirely on your intuition. And sometimes on the second try, you nail it, or the third try, you nail it. If that doesn't happen, then you take a step back and say, "OK, now I want to [*do*] it meticulously and design an experiment." Because my intuition didn't get me there. And now I need to develop my intuition – I need to redevelop my intuition on this particular machine. And then you design an experiment that gives you the plots that help you have an intuition for why it worked that way.

A user at the Alpha facility recounts receiving this kind of help from staff:

> Sometimes cleanroom staff has already established the recipe, and you can ask them to share that with you. Or sometimes, they just have a better feeling of which parameter you should tweak first or tweak more. For me, when I first developed a recipe – [*I felt like*] I'm blind: I'm like in a dark, dark room. But the cleanroom staff told me which parameter would be potentially better or which parameter would be potentially most important, [*saying,*] "You should try that first then try that." So the experience of the staff is really important.

Requests for starting recipes often coincided with tool trainings. For instance, when I shadowed Tim, an equipment engineer at the Gamma cleanroom, he trained two users on an etch tool. At the start of the training, he asked the users which materials they planned to etch. They each responded by stating the name of their material. After the training, one of the users lingered and spoke with Tim about etching his material, titanium oxide. Tim said, "I wonder if someone has a titanium oxide recipe in there already." He consulted the recipe list on the tool's graphical user interface, scrolling down the dozens of recipes already created on this decades-old tool, but didn't find a recipe in the list of recipes there. Tim recommended the user contact Tim's colleague Victor, who specialized in developing etch recipes, to see if Victor had a recipe in his files that would serve as a starting point.

Advanced process development

Other times, process development involved doing something that staff hadn't done or seen before in their tool: a new combination of process integration constraints, a new sample property, a new or demanding kind of structure to be built, superior control over a set of outcome parameters, or a fabrication need that pushed a tool to its limit in a particular area of operation. These were the collaborations with users that staff recounted when asked to describe how they learn about their tool in ways that extended what they had learned from developing baseline recipes and running basic characterizations.

Meeting a new combination of process integration constraints

Process development campaigns were sometimes informed by the unique process integration constraints of a particular structure, leading staff and user to explore how to meet the combination of objectives and limitations that defined a specific situated fabrication need. Jason, a plasma etch and deposition engineer at Gamma, gives an example of such a process development campaign below, where a group is seeking to deposit a film in one of his tools without damaging an underlying layer in their structure:

> We had a group that's making these really sensitive semiconductor surfaces, and they want to passivate them with an insulator. And all the deposition processes they were trying were destroying their devices. And so I've been trying to help them create that first little bit of material, to put it down in a way that doesn't damage the layer underneath.

Jason describes the interactions he has had with these users while exploring parameter values that would provide the requisite film:

> And so *they know* what the material is and whether or not it's damaged, and I don't know exactly the answer, but I know which knobs are the most likely to help. And so we would do changes, and they would take measurements. And, [*the outcome was*] orders of magnitude better just by working together. […] And a lot of that was email. A lot of that work was like, "Can we try 200 Watts instead of…?" and just changing it. And then they'd run it and it's like, "Oh yeah, that was a little better, but can we try 150 Watts?"

This pattern of interaction was common, where users would steward their fabrication objectives while staff recommended incremental changes to the values of select parameters, followed by several iterations in which users assessed how well the outcome matched their fabrication objectives and staff would recommend further changes, until a satisfactory process was found. In the excerpt below, a user describes how he worked with Victor, an etch engineer at the Gamma site, with the aim of etching features with slanted sidewalls, a departure from the straight vertical sidewalls most users endeavored to create:

> I was working [*with Victor*] on an etch process for etching out the circuits of [*my device*]. And going back and forth with him, basically trying a bunch of different recipes on the [*Etch Tool A*]. What I was trying to do was get slanted sidewalls, which is the opposite of what most tools are designed to do. You [*typically*] want them to go straight. This one, we were trying to slant it. I basically teamed up with [*Victor*] and did a bunch of recipes and trials. It was clear, he has a wicked knowledge of that tool. Like, really deep. There were all sorts of things that he was recommending that I just never would have known to try.

I asked the user how he noticed that Victor's knowledge was so extensive compared to his own. The user's reply indicates that Victor has knowledge of more possible approaches ("chemistries"), and has a sense of which parameters will alter specific dimensions of the tool's performance, and how the relative influence of different parameters develops in a nonlinear way throughout the tool's phase space:

> Basically, [*he recommended*] a bunch of different chemistries, and then [*gave*] me access to controlling the bias and the power of the recipe. Starting with one chemistry, but then changing, [*for example*] gas concentrations. It was several different knobs. And he's like, "Here's [*the variable*] I would expect to do this. And this one would probably help up until some point, and then it would [*no longer help*]." It's all those non-trivial understandings of the processing.

In the example below, Jason notes that novel combinations of process integration constraints often occur in the processes that more experienced users are attempting. These users have often already tuned baseline recipes for their purposes and have found them insufficiently precise or unsuited to the additional process integration demands incumbent on their structure. In such cases, Jason works with these users to tune the tool to a level of precision that goes beyond what basic characterizations and documentation on the tool prescribe. In this example, Jason discusses the process integration consideration of film stress (see chapter on process integration for further discussion of film stress), a property that users who deposit thick or high-stress films onto their sample must manage for their device to be structurally sound. On one of his tools, Jason has prepared two baseline recipes: one for a "low-stress" silicon nitride film and one for a "high-stress" film. One basic piece of guidance that staff give to users seeking to manage film stress is that film stress increases with the thickness of the deposited film:

> I think [*the questions that experienced users ask*] are just deeper integration questions. They usually know that there's some lower-stress nitride [*recipe*], and they

> say, “I get great results when I use this nitride, but I need it thicker.” And then they usually have already tried a lot of things. And so the problems are more interesting and more complex to solve, because they’ve already tried this and that other thing that we always tell them to try, because when they were a newbie, they were told, “Well, you can do this or this or this.” And so then at that point, it usually becomes trying to think more out of the box about their different options, or alternative techniques, or just trying new processes, and trying to find a way to decrease the stress of this or that *and* keep the microstructure this way.

Another basic piece of guidance staff typically give users seeking to manage film stress is that the stress of any particular film can be tuned by varying the pressure in the chamber where the film is deposited. Jason describes how an experienced user might interact with the process characterization documentation that staff has prepared for a tool:

> So [*Deposition Tool A’s*] data book has stress as a function of pressure, right? But an old-hat fabber would say, “Okay, well, I know I want to be around 10 millitorr [*of pressure*], but I’m still getting this issue.” Well, there’s other ways, right?

Jason proceeds to name some of the alternative approaches open to the user at this point, which involve adjusting parameters that most users do not become acquainted with:

> You could heat the substrate, and that’s going to change the stress. Or you could apply a bias and you bombard the wafer with argon ions, and that’s going to change the stress. It’s like a peening effect, like blacksmithing. So we might say, “Well, why don’t you try these other things that we don’t usually talk about, because it’s kind of more annoying and easier to get wrong, but you’re good, so go.” So there’s always another way to do things. But, either we haven’t taken the time to characterize them, or most of the time, the pressure is good enough [*for tuning film stress*].

Jason goes on to articulate how these additional adjustments depart from baseline recipes prepared on the tool:

> And usually, those sorts of parameters that we haven’t mapped, it’s when you’re trying to optimize for multiple things. And most of the recipes that we share are kind of good enough for most things. And so it’s not that they’re bad recipes, it’s that it’s not the best at everything. It’s a Swiss Army Knife sort of thing, where I’d rather have a Swiss Army Knife than just the sharpest razor blade ever. I’d rather have it be where it kind of does everything pretty well, and then you got to adjust it if you need better whatever [*performance measure*].

Jason describes two examples of processes requiring film stress to be managed conjointly with other performance measures:

> [*They'll say,*] "So I'm using it as a waveguide, and I have too much loss. So how could I improve loss without changing the strain?" Or, "When I run the low-stress molybdenum, my resistivity goes up. How could I keep the resistivity down while also keeping the stress down?"

Similarly, Tim – responsible for a class of deposition tools in the Gamma cleanroom – describes starting process development campaigns with a more generic recipe and then proceeding to less commonly used tool parameters. He gives a recent example when he helped a user create a superconducting niobium film, a basic element used in a subset of quantum computing devices. The goal for the user he was working with was to create a niobium film that would exhibit superconductivity at higher temperatures. Their first attempt with the tool's baseline recipe resulted in a film that was superconducting at a given temperature, which made Tim optimistic about future attempts to improve the film:

> We were able to get that value just doing a room temperature, plain Jane niobium deposition, which is really good. That means, now we start tuning and playing more. So: if you do heated depositions, and you play with the chamber pressure, you can change the crystal structure a little bit, get different grain sizes in the metal and so on, and hopefully start seeing improvements.

Supporting this view, an experienced industry user describes wanting staff to help him get tighter tolerances on one of the photolithography steppers, a collaboration that he anticipates will lead to him becoming acquainted with some of the tool's deeper settings:

> User: If you need to get deeper into a process, and get it to higher levels of precision, [*staff*] will sit and work with you and teach you how to really make the tool sing.
>
> Danielle: Is there a recent example you can think of there?

> User: A recent one, not so much, but I'm going to be [User chuckles], they don't know it yet. But I'm going to be bugging Joe and Jim on how to get a little more out of the [*Stepper A*].
>
> Danielle: Okay. Like smaller feature sizes?
>
> User: Just more reproducibility. It's how to get tighter tolerances on my feature sizes, as opposed to smaller, necessarily.
>
> Danielle: Okay. What would be the kinds of information you'd share with them to help them help you?
>
> User: I would probably talk to them about my current process. And my recipe. I would show them some SEM images. I would ask their advice on like, "What parameters should I change?" We'll start with the simple dose-focus changes, and then we'll move on to like, "Okay, what are some of these menus that are buried a little bit deeper in the software?"

The basic structure of a process development campaign with a user, then, was to define a specific tool performance outcome, or set of outcomes, to optimize for. Then, staff provided users with a baseline recipe and guided them in the adjustment of parameters most likely to affect the desired outcome, progressing from most commonly-used parameters to more rarely-used ones. Users run the experiments and collect data. The user returns to the staff member to discuss and interpret the results and plan the next set of experiments until a fit was found or the approach was abandoned.

Learning from process development campaigns

It is implicit in staff members' accounts of more ambitious process development campaigns they've embarked on with users that staff learns from these collaborations, because these involve mapping unmapped areas of tool functionality. But there were also times when staff explicitly talked about what they had learned from these campaigns, and how what they had learned had the potential to benefit future of that tool.

Learning about tool limits and capabilities

Different tool owners provided several examples of what they had learned from what users were doing with their tools. For instance, even after working with many of his tools for between 20 and 30 years, Joe said he still learned from working with users with novel use cases for his tools, which he called users "doing oddball things" in the quote at the very beginning of this chapter. While I shadowed him, Joe gave examples of the "bits" he learned from users doing oddball things. One kind of bit involved information about the limitations and capabilities of his tools. In one case, Joe described how one user successfully ran a wafer with dimensions outside the stated range of the tool. This taught Joe that there is more room inside the tool than he thought there was based on the tool's documentation.

> So, you know, there's always little bits. [*One user*] put in a wafer [*into the tool*] that was 1.2 millimeters thick. And the absolute limit, it's supposed to be 1.1 millimeters. And it [*the wafer*] went through. And that tells me: well, we've got a little more [*room*] there than we thought.

In a second example, one user of a tool that Joe had maintained for decades had attempted to program the tool with a uniquely lengthy set of wafer handling instructions. This user's attempt, and failure, taught Joe about software limitations that he did not know about:

> He's run into some real weird software stuff. He taught us that if you have more than 10 passes in a job, you can write them, but [*the tool*] won't run them. It stops at ten. So, software limitations. It's pretty interesting.

Learning how to handle a new class of input types

Sometimes Joe's engagement in a user's novel use of a tool was more extensive, prolonged, and challenging to his own understanding of the tool. Such experiences resulted in Joe better understanding how a whole subsystem of a tool operates and improved his ability to master a class of problems on the tool he had not mastered before. In the example below, Joe describes what he learned from accommodating a user's request to run a different input type through a tool than the tool had been designed to handle. Specifically, the user was requesting to run pieces of wafers rather than whole wafers on one of the photolithography steppers – a complex and integrated machine not designed to process pieces. Configuring the tool to nonetheless run pieces prompted Joe to deepen his knowledge of the alignment and leveling systems of the machine. He did this by working with a user for four months, getting the user to document his approach, and by obtaining documentation from the vendor on how the relevant subassemblies of the tool worked:

> I learned an awful lot about how the alignment system works and how the leveling system works, and how to work around it for unusual situations, pieces in particular. Which I never would have figured out unless I'd had to – I had a guy [*user*] working on it for four months. That's what he did. And I got him to document a big part of it. But I also got a lot of info from [*the tool vendor*]. They [*released some of the documentation to me*.] So now I know a lot more about the tool than I did before.

Joe says that what he's learned about the subsystems necessary for running piece parts on the Stepper now gives him the ability to instruct – and appropriately caution – future users who want to try something similar:

> And if somebody has that issue [*i.e., wants to run piece parts through the Stepper*], I've at least got something to point at and say, "Here's the kind of problems you're gonna run into." And there's actually a lot of them. Because [*the vendor*] designed that thing for whole wafers. It has no idea why you would put a piece [*in*] there. And nothing's set up to work right when you do that. So, it's a real nightmare. You can do it, but boy, be prepared. […] But at least now you don't have to start from scratch. And we know what to look out for.

From this experience, Joe gains knowledge that he will bring into early discussions with users, to assess the feasibility of their device and identify where in their process development they will have to invest more time and money. In this way, tool owners' everyday work in process development feeds back into the feasibility assessments they make early on in users' projects.

In addition, the user's dedication to a specific novel use case gave Joe the opportunity to learn to use a machine for a class of inputs unanticipated by its designers. He learned how several relevant systems within the tool work in such a way that he could create the conditions for successful processing in conditions that were outside of the tool's stated performance specifications.

Broader impacts of staff's learning about tools

That Joe has extended the Stepper's capacity to process piece parts as well as full wafers is emblematic of a class of work that staff at these facilities do: they adapt old tools to maintain step with emerging research frontiers. Joe's tools were originally built in the 1980s and 1990s to process whole SEMI-standard silicon wafers. In the intervening years, micro- and nanoscale researchers have become increasingly interested in patterning more exotic semiconductor materials, such as Gallium Nitride (GaN), which are more expensive to make (or "grow" in the local jargon). Users, especially academic ones, cannot afford full wafers of these materials. Therefore, there has been an uptick in processes that require patterning piece parts in recent years.

Joe's work on the Stepper is unique for the technical depth and complexity of the alterations that he has made, but it is emblematic of staff's numerous and continuing attempts across the cleanrooms' tool sets to keep ageing tools in step with the process demands of emerging research frontiers. As the vignettes presented thus far indicate, these adaptations take place in the context of specific user projects, as tool owners and users address the issues they encounter when trying to do something novel. When involved in these process development campaigns, staff learns about the limitations of tools, linkages between relevant subassemblies and tool outcomes, and overall, how to tune the material properties that matter for a given emerging research frontier.

Joe had also adapted another one of his Steppers for processing piece parts. This capability drew users to the facility who had this specific fabrication need, as one user told me:

> So I rely a lot on Joe's help with this tool, just because it is an advanced tool that is really capable for working on pieces. It has like 16 different chucks that you can use, depending on what your substrate is, which is really special, because there's really nowhere else where they have that kind of capability. A lot of other fabs will work on 4-inch wafers, or 6-inch wafers, but they won't work on cut pieces. It's too much of a headache.
>
> […]
>
> So that's what's so special about here, is that I can do work on pieces. Many other places, they won't let you do that; they only let you do whole wafers because it simplifies a lot of things and most of the tools are configured in a certain way to work with that. Joe has special chucks that work with pieces, so that's why [*I came here*].

Learning to develop superior control over a set of outcome parameters

Another type of novelty that tool owners encounter during process development campaigns is when working with users prompts them to gain control over a section of the tool's outcome space to an extent that staff has not had before. An example of a staff member learning from such a process development campaign with a single user in ways that they anticipated were valuable for a broader set of cleanroom users comes from Tim, a deposition engineer at Gamma. He indicates that what he's learned becomes relevant for future users interested in optimizing for that outcome measure for that material.

Tim describes a user working on micromechanical applications who wanted to deposit a film of aluminum nitride on his sample in a layer that was characterized by low compressive stress. This stress is a problem for structures that are meant to be entirely level to function, that use very thick films, or whose functioning depends on tightly calibrated force transfer, such as micromechanical devices. This user wanted to get the stress in his layer of aluminum nitride as close to zero as possible. The combination of the user's material, film thickness and stress requirements was something Tim hadn't seen before. In addition, they were using one of the newer tools in the cleanroom to develop this process, a tool on which relatively few recipes had been developed up to this point. Getting the results closer to where the user wanted them to be required Tim to talk to the tool vendor. Tim describes how he and this user worked together to hone a recipe:

> Tim: I worked with a professor doing a lot of this aluminum nitride stuff. And so, I learned by talking to his students, [*and responding to them about*] what's wrong with the tool? In turn, they give me all this information that I can pass to the next person.
>
> Danielle: What's the kind of information you get from interacting that you can pass to the next person?
>
> Tim: This tool is a great example of that. One of the students wants low-stress thin film aluminum nitride. He wants to do 100 or 200 nanometers, and have it be low stress. It's normally incredibly compressive stress, like, two gigapascals. So getting that to zero is a challenge. But we worked together on it. I talked to the guy at the company that invented the tool, he gave me some [*suggestions like*], "Move things this way." We got them working, and that kid [*the student*] went and ran wafer after wafer, and was getting me the data, telling me, "This is moving in the right direction. We're not there yet." And then the tool breaks [Tim chuckles]. So it's a tough tool. Takes a lot of work. But that's [*an*] example of learning together.

Tim configured the insides of the tool with guidance from the vendor, and the student performed the experiments and collected the data. Tim indicates that both he and the student were "learning together." As I left the field, Tim and this user were still working on this process. In this case, tuning the tool to optimize a specific performance outcome gave Tim a better understanding of how to control given aspect of the tool. He says he is now better able to help future users who want to create low-stress aluminum nitride films using his tool.

The difference made by the quality of users' participation in problem-solving

Alongside their statements about learning from users' processes, staff also indicated that the opportunities to learn from a given user's process could be sharply curtailed by the user's openness, specifically regarding the results of a given process development campaign. For instance, Joe indicates after making his statement about users "doing oddball things," that it is hard to get users to document what they've done:

> Joe: The problem is, we don't always document everything. We've got it up here [Joe taps the side of his head]. And it's nice to document stuff, we try to document some things. But some of it, I don't even know how you could document it.
>
> Danielle: Yeah, that's a good question. How you would document something like that
>
> Joe: Well, there's a user's wiki page. And they're supposed to write down things like [*that*], but they won't do it, either because they're too lazy or they don't want anybody to know what they did, because it's a secret. So, it's kind of hard to get that information out sometimes. It's always been a problem, where users have got a project, and someone [*else*] asks us if we can do this. I say, "Yeah, this guy did it." [*They say,*] "–How do you do it?" [*I say,*] "– I don't know. You can get ahold of him to see if he'll tell you. Because he wouldn't tell me."

The variability in users' willingness to report back how a given approach went is supported by statements by experienced users:

> Danielle: Do you think [*staff's value to you as a user*] is also because staff is able to pay attention enough to everyone's problem and also learn from that, and then help the next user with it?
>
> User: There's a certain amount of that. Some users, because they, say, are commercial, or are very secretive because they need to publish, are not open about their processes. Some are open about their processes. So the accumulated experience helps.

As described in Chapter 4, it was possible for users seeking to divulge as little as possible about their overall processing strategy to avoid process integration conversations and contact tool owners directly, with precise fabrication needs that they themselves formulated. Staff describe being able to recognize when this is the case, ascribing it to a desire to keep things secret:

> Jonathan: I had a guy that was working for a company [*and*] had an SBIR thing. He was doing different wattages on [*this tool*]. And he was really trying to get really specific things on the tool that no one's done before.
>
> Danielle: Did you learn from him things about how to use that tool?
>
> Jonathan: It's tough, I actually didn't. Because he kept a lot of that stuff under his vest, and he came here at night.

Jonathan showed me some of the email correspondence with this user, in which the user made very concrete requests of two different staff members. Jonathan surmises,

> There's a whole process here that he has, that Jason and I are getting little pieces of. And I don't really know what he's up to.

On another occasion, I am shadowing Jonathan and a user walks up to him and asks about using a particular tool, for which she needs the staff's assistance because only staff can make changes to recipes on this particular tool. Jonathan and this user interact, and afterwards Jonathan takes note of her hesitance to share what she is doing and how this may hamper her getting on the appropriate tool for her sample, which is a piece part and not a full wafer. The tool she was asking about processes large batches of full wafers and may pose a risk to her small sample.

> Danielle: What was her question?

> Jonathan: She wants to take the [*Etch Tool A*] and she wants – you saw that, she doesn't exactly tell me what she's doing. You could see that she's kind of hesitant in saying like, "This is what I'm up to." And ... you can take that for what it is [Jonathan chuckles]. But she wants to take a tool that has, the recipes are locked. And she wants to do it at a lower wattage. And, some tools are better for manipulating these things and some are not. So that's a barrel etcher. That's made to put like 25 wafers in them [*at*] one time. She has a chip. Chip's going to get sucked right out through the vacuum. That's a vertical tool, so the vacuum's here. So if she lays a chip, it might go right through. [Jonathan gestures to other nearby etch tools] These tools are different. They're made that the vacuum's underneath the plate, so it takes a lot to suck [*the sample*] out.

As indicated in the chapter about process integration, staff tended to probe the fabrication needs users initially expressed, because these expressed needs were sometimes not grounded in knowledge of fabrication practice its tools. This habit also carried over into the tool domain, where tool owners preferred to ask users questions about their sample properties and processing objectives in order to match them to the appropriate tool. In tool owners' eyes, users' reticence hampered this matching.

Conclusion

By examining the role of tool owners, we have seen an important part of how the facility as a whole endeavors to match tool capabilities to users' situated fabrication needs. What this chapter has specifically shown, is that many of these tool capabilities are not written in the tool manuals or vendors' tool specifications, but rather are discovered by staff as they endeavor to meet novel fabrication needs posed to them by users. Staff's sense of what their tools could do was rooted in an experiential understanding of the physical systems that supported and circumscribed tools' performance, which they developed by installing, maintaining and repairing the tool, and honed further through attempts to tune the tool during process development campaigns. Staff further built a picture of the capabilities of their tools through numerous, repeated interactions with users in the course of troubleshooting and process development. By presiding over knowledge of the technological possibilities that the facility's tools represented, tool owners underwrote the flexibility of the facility to foster new research frontiers.

The next chapter completes discussion of Figure 6 by considering the role that tool owners played in circulating knowledge within their facilities' user communities.

Chapter 6. Tool Owners' Role in the Flow of Knowledge

Introduction

The previous chapter explored the role that tool owners played in their facilities by considering their daily work, as it involved keeping their tools in working order and working with users to tailor a given tool to their specific structure. This chapter builds on that discussion to describe how the distinct occupational knowledge that tool owners developed allowed them to not only help individual users tune tools, but also positioned them to learn from these process development campaigns in such a way that these efforts had the potential to benefit the larger user community. I'll argue that staff were better positioned than users to trace the root causes of any given outcome that their tools produced, and that because of this they were able to use their experiences developing processes with users to bolster their knowledge of the materials, tools and techniques involved in the process. As tool owners improved their ability to use the tools and techniques at hand, they met future process development efforts with greater control over the properties of materials at the base of nanoscale structures.

At the heart of the argument that I make in this chapter is a difference between the knowledge that tool owners gained from process development efforts and what users learned about the tools they used. The main difference between staff and users' knowledge that I will discuss lies in their respective ability to account for the outcomes their tools produced. In their equipment maintenance activities and through their involvement in a range of process problems, staff's knowledge of their tools was substantially aimed at developing greater and finer control over their tools. These experiences acquainted them with a set of factors involved in tool performance typically invisible to users. Three such factors that will be discussed in this chapter are: the links between discrete machine components and process outcomes, the role played by sample properties that users typically kept constant within their own process, and the range of byproducts that tools produced in their operation. Due to their role, tool owners were better positioned than users to ascertain the role that these factors played in the outcomes they observed. In addition to considering a broader set of factors with which to account for tool performance, tool owners also developed a finer lexicon with which to name the performance of their tools.

The first implication of this aspect of tool owners' knowledge is that tool owners became repositories for discoveries that would otherwise not be available to users other than those who made the original discovery (Figure 6). Staff's grasp of the conditions for a given processing outcome allowed them to leverage their experiences within specific projects to develop knowledge in domains that were fundamental to the diverse pursuits of their facility's users. I use the term fundamental in the sense that this knowledge supported the fabrication of elements common to numerous structures being built in the cleanroom. The second implication is, therefore, that tool owners were in a position to repurpose solutions developed in the context of specific projects to advance the work of a broader set of users that relied on similar materials, tools or techniques. Users, on the other hand, described struggling to repurpose with precision what they had learned from a given process development campaign to structures that resembled theirs in one aspect, but differed in others (for instance, the structures being made by group mates). This chapter therefore provides a sense of the kind of development and dissemination of knowledge that tool owners were able to carry out by virtue of their role in the division of labor.

Challenges to users' attempts to repurpose solutions across structures

The fact that users' fabrication needs – and therefore the features fashioned to meet them – were situated within complex structures posed a challenge to users attempting to repurpose what they had learned from fabricating a given structure to the fabrication of other structures. For instance, an experienced academic user at Alpha facility, whose group makes electronic structures with relatively many layers said,

> [*When processing in the cleanroom*], there's always these little details, like if you encounter this very specific situation and the solution is this very specific thing you can do. And it actually was a bit of a theme through graduate school, trying to understand how, as a lab group, we could take the information of what problems other people have encountered in their very specific projects and maintain that knowledge, rather than having to solve the problem again, which I think, unfortunately, a lot of cleanroom processing is solving the same problem over and over again. So trying to find a way to write it down in a way that you can really use later, it's hard.

The user then talks about ways that his group has tried to retain and repurpose knowledge from prior projects, which included adding more detail to their process descriptions, but to little avail:

> There was an attempt at some point to use a wiki [*within our group*]. There's been attempts to write down in a lot of detail, a lot of processes, with varying levels of success. Though, ultimately, everybody that starts a new graduate program ends with a totally different process than anyone has ever done in that group because, otherwise, it wouldn't be graduate work. So yeah, it's hard even when people really wrote down details.

The user continues to describe that a challenge for groups' attempts at sharing useful problem-solving information among themselves is that solutions are situated within multi-stage, multi-material, and ultimately distinct, structures:

> When you're processing, you start with a clean wafer. Maybe if everybody was just comparing the first step of their process it would be really easy to have suggestions for another, but when you're in step seven it's like, "Well, we had this problem because the gallium nitride in step three of our process was etched 6 microns." And

> you can’t really, how do you put that in a wiki and say, “This is a problem that’s caused by—”?

He also indicates that the time and effort needed to develop this kind of knowledge – knowledge of the complete set of causes of a situated problem – does not match the goals of users motivated primarily to obtain an integrated outcome:

> And I think the answer to that question [*I raised above*] is you have to understand the problem well enough, but when your Ph.D. is not understanding the problem, but it’s to build a device, nobody really wants to dig into the individual issues that much. Actually, I think that’s generally an underlying point that’s worth thinking about, is that we [*the Ph.D. students*] are mostly electrical engineers, not processing engineers. So our goal for the Ph.D. is to build a thing and then characterize [*i.e., measure*] the thing, and you write the paper on the thing, and usually there’s a section in your paper that says how you fabricated it. Sometimes you might spend 80% of your time trying to figure out how to fabricate the thing, but that’s not really what your paper is about and not really what your Ph.D. is about. So, learning how ICP etch chemistry works really, really, really well so that you can optimize an etch is not really desirable. It’s kind of a giant time sink, which is why it’s [*of*] such fundamental importance that the cleanroom [*staff*] handles the equipment themselves.

In other words, even though this user’s group tried to write out their processes in greater detail and specificity, this did not aid the generation of transferrable knowledge, because the problems and solutions were so “specific” to each structure being fabricated, and users were not motivated to trace the full set of factors underlying the effectiveness of these situated solutions.

This user’s experience suggests a more general problem with retaining and repurposing knowledge gained from building complex structures. Because of the great number of conditions affecting structural success, each solution found is situated to a high degree. Repurposing a solution therefore requires understanding the full set of conditions that fostered the feature’s fabrication in order to adapt it to other contexts. In addition, as staff explained, some of the factors underlying processing outcomes were typically invisible to users. The next section considers what such conditions were, according to tool owners.

Differences between tool owners' and users' knowledge of tools

As we have seen in the previous chapter, a dominant concern for tool owners in maintaining their tools and assisting users was to develop and maintain precise control over their tools' operation. Consequently, an important kind of knowledge that tool owners developed in the course of their everyday work centered on accounting for outcomes that their tools produced. Herein lay a key difference between the knowledge of tool owners and that of users. *In accounting for outcomes, tool owners typically considered a broader set of factors than users did, as well as distinguishing more finely between outcome dimensions their tools could produce.*

Accounting for this broader set of conditions and outcomes allowed tool owners to relate what they learned from situated fabrication problems to more fundamental bodies of knowledge: knowledge about the elements common across nanofabricated devices and that underwrote the success of many different structures being built in the cleanroom. As indicated above, users struggled to learn generalizable lessons from their situated processing efforts, in particular in transferring what they had learned to different problems. Tool owners' deep familiarity with their tools allowed them to learn from experiences they had in the context of specific projects to develop knowledge that was useful to other cleanroom users.

In the next section, I will describe three factors that tool owners considered in accounting for their tools' performance, but that were typically hidden from users. These factors are unlikely to form a complete list, but they are the ones that tool owners mentioned during my time in the field. Tool owners traced the performance of their tools to (1) the tool's components and their interactions, (2) the properties of a user's sample, and (3) the byproducts and intermediate products that tools produced even as they created the focal outcomes.

Tracing outcomes to tool components and their interactions

Because tool owners endeavored, through maintenance and process development with users, to control and account for subtle aspects of tool performance, they learned to trace outcomes to mechanical and chemical operations occurring inside the tool, naming these as products of interactions between discrete components and subassemblies. Jason describes how, as a tool owner, his familiarity with his tools' entire physical and control infrastructure enables him to trace process outcomes to their mechanical cause:

> I fix the piping, and I figure out the software-hardware stuff, and I figure out the process and the process integration [*on my tools*]. […] I think it's really nice to know [*that*] if you saw this non-uniformity on the wafer, that it was caused by this specific hardware sort of failure. And if you're not somewhat [*responsible for the tool in this way*], and have seen tools break, you're not necessarily going to know why the process looks the way it does, or the devices perform the way they do.

Staff said that they had gained much of their ability to trace outcomes to respective tool operations in the course of troubleshooting. Tool owners described the way they developed what they knew about their tools as a "trial by fire," involving a series of uncomfortable experiences early on in tool ownership in which the tool failed, the sources of failure were unknown, and the tool owner bore sole responsibility for getting the tool operational again amid frequent and numerous inquiries by users regarding its status. When repair involved extensive technical troubleshooting, the staff member gained further knowledge of the mechanisms that governed the tool's functioning, as Jason indicates:

> I learn from the problem solving. A big teacher is just how [*the tool*] fails, and then how we account for it or fix it. Whether [*the tool is*] broken or not, how to account for some issue.

A staff member whose own expertise lies predominantly in process integration draws a distinction between his expertise and that of staff members who have "a tool approach." He describes tool owners as "relentless" in troubleshooting:

> Some of these guys [*i.e., the tool staff*] are awesome, because they go and they will be relentless, and they will just figure it out and fix it. They have this analytical mind for where the problem is for all these things.

The technical anatomy and operation of a given tool was therefore the first major set of factors that staff considered when accounting for a given outcome. According to staff, users were often only dimly aware of the influence of these factors on their process. When I was shadowing him, Tim said he had considered offering a course to advanced users to teach them more about how his tools – a set of deposition tools known as sputter tools – functioned. Because, as he puts it,

> …a lot of [*users of this* tool], they treat it like a microwave: "Put my wafer in, tell it what to do, and I get my wafer back." They don't really understand the processes, fully.

I asked him what topics he would cover in such a course. He replied by naming components and subsystems of the tool that were responsible for control, measurement, and power supply, specifically around a subassembly of the tool known as its "guns", into which bulk material was loaded in the form of flat discs known as "targets" to then be dislodged and propelled toward the user's sample through a series of electrical and chemical processes:

> There's a lot. Just knowing what any of these components even looks like, right? [*We'd*] talk about a capacitance monometer [*i.e., a type of pressure gauge*]; they probably never physically held one or seen one. Looking at some of the controls and the pneumatics. Like, [*Tim gestures to the guns of the sputter tool we are standing*

> *near, whose vacuum chamber he has opened for maintenance*] those spring actuators are fed through the gun. That's the actuator, and the power is right next to it, for each gun. How all that cabling goes together, and stuff like that. You got different types of power supplies in them. So the one over there is a switching power supply: so each one of them can handle powering five targets, one at a time. This one is a single power supply with a separate switch that directs it to which gun it's supposed to go to. All that stuff. […] They don't know what all these things are and what they do. To have them have a better understanding of that would be cool.

Tracing outcomes to properties of users' samples

The broader set of conditions that staff considered that generally eluded users also included numerous properties that users' samples could have. Staff described being aware that users' sample properties could influence tool processing in many ways, depending on the focal process of a tool. Staff had seen many different sample properties interact with their tools, because they had seen many different users' samples. Users, in stark contrast, had often seen, or were only interested in, the variation of the substrate that occurred within their own processes, which was often limited. Therefore, if users did not vary their substrate types or other sample properties, they had no occasion to learn about the role it played in their outcomes. Two sample properties illustrate this contrast: a sample's substrate material or materials stack, and the etch aperture of its features.

Variations in substrate (or stack) material

At the most basic level, the material of a user's substrate affected the performance of many tools. For instance, in a photolithography tool, which uses carefully calibrated doses of light to pattern a photoreactive layer spun on the user's sample, it matters how reflective the surface of one's sample is to light. As Doug, a staff member specializing in photolithography at the Delta facility articulates below, sample properties can change the value of an input parameter (exposure dose) and a dimension of tool performance (expansion of the pattern):

> If you have a transparent substrate versus something really shiny, like aluminum or gold, it changes exposure dose [*and*] how much expansion of the pattern you get.

Adam at Gamma conveys a similar point regarding photolithography and a sample's unique material stack, which was the precise combination of materials layered on top of a substrate:

> A silicon wafer with resist on it is going to behave differently than a silicon wafer with a layer of chromium on it and a layer of nitride on it and then the resist.

Users predominantly worked with a single substrate type or material system. Their understanding of the role that their substrate material or stack played in a given tool's process was defined by the variation included in their design. Users would therefore often not be aware of the role played by properties of their chosen material until they begun to vary it, as Doug indicates:

> Usually where you really run into [*differences between substrate materials*], is users that aren't very aware of what they're doing. Maybe they've always worked with only silicon. And they've never really dealt with anything else. And then suddenly, they have a project that comes up that's SiO_2. And depending on the tool, it can behave very differently.

Doug names a number of implications that substrate type can have on the requisite tool settings to get a similar outcome in photolithography, where – as mentioned above – the shininess of a substrate affects the requisite exposure dose:

> Oftentimes it can be exposure dose: you have to put more dose into it, because you're not getting the reflected light [*from the shiny substrate*]. Another thing that can happen because of that, [*is*] your patterns will come out larger than designed. So you'll get a little more spread.

Doug's example of users struggling to go from Si to SiO_2 indicates that there was a set of parameters and processes typically hidden for users because these were kept constant within their processing. In Doug's words, these users "aren't very aware of what they're doing." Users realized the importance of adapting a given recipe to a sample property when they tried moving to a different substrate, and their recipe parameters no longer yielded the desired results. Staff described many other cases in which the lack of variation in a single process limited what users learned about the full set of conditions responsible for their result, and how staff's own experience with a range of users provided a quality to their knowledge that involved recognizing the role that these typically constant variables played. Seeing many different use cases for their tools meant that staff developed an understanding of those preconditions for outcomes that were hidden to any one user but that became apparent when exposed to the broad variety of projects occurring across the user base and across the tool owner's career.

This point of contrast with users' knowledge led staff to sometimes refer to their knowledge about a tool as "broader" than that of users. In staff's view, experienced users gained deep knowledge of a single path (or set of closely related paths) through the fabrication possibilities that the tools presented. Tool owners, on the other hand, had traversed the same landscape numerous times, taking different routes. In the excerpt below, Laurence, a staff member who works with the E-beam lithography tool at Gamma, notes that even long-time users of the tool typically had a purview that remained limited by their application space:

> A few users become experts. They learn how to do this one thing. [*For example,*] you want to make a transistor, you want to make a waveguide, you want to make this X-ray lens. You know a lot! But about whatever was directly impacting your ability to move ahead or not move ahead.

Staff, on the other hand,

> … have been working on it [*the tool, the technique*] for so long, they have a much broader and deeper understanding of all the interplays that go into doing this. A lot of [*what staff knows*] is just interrelationships with other sample types, substrates and

> things that this [*user*] had never imagined because he doesn't care about that. Why would he consider this or that or the other?

Variations in etch aperture

Another example of a sample property that affects the outcome of a process relates to etching. In plasma etching processes, material is selectively removed from a sample by introducing carefully chosen gases into a vacuum chamber with the sample, bringing them into contact with the sample in an ionized state, during which gas atoms dislodge and bond with atoms on the surface, and then pumping the products of the reaction out of the chamber. It was well-known among staff at the different facilities that the rate at which one's etch occurs is affected in a significant degree not only by material type and chosen gases, but by the "aperture" of the features on one's sample. With aperture, they meant the surface area of the gap through which etch gases could flow into and out of the area to be etched. That is, if one is etching deep, narrow features, the etch process will proceed relatively slowly, because the chemicals and their products have a relatively small opening through which to enter and exit the area to be etched and flow across the sample's surface. But if one's sample has features with larger openings, then the etching will go more quickly, because the gases have more space through which to go in and out from the sample's surface. One's etch rate determines how long one must run the etch, and is therefore in important consideration in honing any etch recipe.

Here as with the case of substrate material, users' exposure to this source of variation is determined by the variation in their process. When staff saw that a user had designed features with significantly varying apertures on a single sample, staff alerted them that they could expect lower uniformity in their etching and proceeded to discuss the degree to which this was tolerable given the overall objectives of the device. However, if a user's sample had no variation in aperture, it was possible that they would not learn about this element of the etching system. What happened in this case is that this variable had, in effect, been held constant for them. At the start of tool trainings, tool owners asked the users a set of questions to ascertain such tool-relevant aspects of their sample and targeted their instructions to what was relevant to the subset of sample properties present in these user's processing.

Tracing outcomes to the formation of byproducts and intermediate products

Staff also recognized that the conditions for the fabrication of a desired feature often also created conditions for the production of unwelcome products and features. Accordingly, staff's understanding of the system by which tools created outcomes also included what users considered byproducts and side-effects. For instance, in the course of an etch, pieces of debris may be dislodged, or byproducts created of chemical interactions involved in the etch. The result could be specks that appear on the user's wafer. A user will interpret a result as "dirt" or "junk." Staff, on the other hand, interpret the result as a byproduct to be expected from the tool. An example of this was provided in the previous chapter, under the heading of "tool troubleshooting." In this case, what users considered a surprising result, staff interpreted as a natural and manageable byproduct of the focal reaction.

Staff's inclusion of byproducts in their accounts of the outcomes of processes extended to kinds of unwelcome structures that can appear in the course of building a feature. Staff had names for these kinds of features, such as "fences" and "reticulation", each of which could be traced to conditions the tool creates. Staff's understanding of the conditions for tool functioning, therefore, included byproducts to be managed throughout.

This approach is exemplified by Craig when he describes how, when confronted with an etch result for which he has no ready explanation, he considers factors that include byproducts of etching ("things that come off" in his words):

> It's seeing something and not understanding it and saying, "OK, I need to understand that. So, let me think deeply about what I'm doing, and what could possibly cause things like this, understanding what I understand about how ions interact, *how things that come off interact*, how light interacts, etcetera, etcetera."

During etching, material is removed from the wafer's surface (it "comes off", in Craig's terms) and can still interact with the ongoing process. It was important to understand the properties of these pieces of debris to account for unexpected outcomes of etch processes. Another way of saying this is that staff developed an understanding of tool performance in which the focal process was understood as one process among many that took place in the tool as it ran, any of which could be positive, negative or neutral for a user, depending on their application.

A fine-grained sense of different outcome dimensions

Besides recognizing a more comprehensive set of conditions for tool performance, staff also maintained a more detailed lexicon of properties along which processing outcomes could vary, and how each related to the conditions created in the tool. A typical situation in which this kind of knowledge was elicited from staff was when users asked staff members for guidance on depositing a film that was of "good quality." Knowing that there are many qualities a film can have, staff members would respond by asking users to specify what quality they were most interested in, as Roger – a staff member specialized in one technique for depositing thin films – models below:

> We would basically ask them, “Okay, so what exactly is [*a*] good film to you?” Is it thickness? Is it the stress in the film? Is it the optical properties? Is it stability, electrical properties?” So, we have to kind of tease out what their end application is. And from that, we can then make the assessment, “Okay, this is what you need. We’re going to recommend that you do it this way or use these conditions.” And that’s something that a biologist wouldn’t have knowledge of because they’ve never really learned the tool extensively. So, they don’t understand what the range of possibilities are.

Staff therefore knew how to distinguish between different outcome variables where users saw only single measures of outcomes. Craig expresses a similar point:

> Different people will have different requirements based on the types of devices they’re making. People think, “Oh, what’s the highest-quality film?” I always look at them and say, “Which quality of the film?’ Because there’s no just one ‘film quality’. And maybe you want poor quality for your device, I don’t know. You tell me.”

As Craig indicates, properties that were considered adverse to some applications were in fact desirable for others. For instance, silicon whose surface has become rough or “grassy” is considered undesirable for many applications. However, for some, it is a feature to be fostered and tuned. Solar cell research, for instance, tunes the roughness of silicon to allow for greater absorption of light. Peter, a process engineer at the Alpha facility, recounted how a user had recently asked him for help in creating such features on his wafer, which gave Peter the opportunity to apply the knowledge he had gained by helping users *avoid* surface roughness, to obtain its opposite. Likewise, James, an equipment engineer at Alpha said, “There’s so many different goals that your optimizations are always changing.” Therefore, those outcome properties that in a given project were considered desirable or undesirable were treated by staff as simply one of many properties to manage.

Conclusion

This chapter has documented how tool owners were positioned to learn from their process development campaigns with users in such a way that they could then meet future process development efforts with a greater ability to finely control the properties of materials. Specifically, this chapter has documented the aspects of their tools' functionality that tool owners said users did not consider in accounting for tool outcomes. In the course of their daily work, staff's knowledge of tools came to encompass an ability to account for observed outcomes with reference to a broad set of conditions that users typically did not consider: mechanical preconditions, numerous sample properties, and process byproducts. These were domains of knowledge that typically eluded users, whose role shielded them from the mechanical intricacies of tools, from variation outside their own process, and from systematic understanding of processes considered secondary to those directly impacting their focal application.

Because tool owners could account for their tools' performance in a more complete sense, understanding its preconditions, they retained and repurposed knowledge gained in process development campaigns more flexibly than users could. Staff could learn from a situated problem in such a way that they developed their knowledge of its constituent elements. In this way, tool owners' role was necessary for knowledge of how to achieve specific processing feats to be repurposed in different ways across diverse projects. In this way we see how staff members developed and disseminated knowledge about how to handle materials to gain specific properties. The next chapter considers another aspect of nanofabrication around which the staff's knowledge clustered: elemental structures that occurred across different structures.

Chapter 7. Parsing Devices into "Simple Structures" to Transfer Knowledge from One Project to Another

Introduction

In Chapter 1, I argued that staff's working knowledge includes knowledge of buildable substructures that can be integrated to form more complex forms of functionality. This chapter supports this idea with empirical material.

This chapter describes a problem-solving interaction between a user and staff early in the user's project at the GNF. The larger thesis I want to explore with this data is that staff thought of user projects in terms of buildable substructures, which staff referred to variously as "simple structures", "baseline structures" or "motifs" that reoccurred across different devices. These substructures were responsible for defined aspects of the overall functionality of a device. Staff knew that these substructures could be built because they had seen them built in the lab before. Staff members drew parallels between user projects over time in terms of these substructures, even if other aspects of the overall devices differed (e.g., their final purpose, their scale). I will first briefly describe how the user in this observation, a researcher from a discipline in the life sciences, defined her project at the start of the meeting and how staff formulated their design suggestions to her. I will dwell mostly on the ways in which Rick and Adam, two staff members at the process meeting, described their goals for the process meeting in a later conversation.

The user defines her device's objectives

One afternoon, a researcher from the life sciences had reserved a time slot in the GNF's weekly open process meeting to discuss a project she was considering. The prospective user (pseudonym Kathy) wanted to develop a device that would handle cells in such a way that their DNA could be extracted, and segments of the DNA could then be captured and identified (“tagged”). For this to work, the device she was proposing would need to be able to isolate individual cells from each other, puncture (“lyse”) each one and then bring the liquid contents from each punctured cell into contact with a biochemically engineered object called a “bead” which would be prepared to include a tagging mechanism. In the researcher’s words, the goal was for the structure to “physically separate” the cells from each other, and then “to associate one bead with one cell.” After lysing and tagging, the liquid that included all the tagged cells (ideally, thousands of them) would be dispensed into a vial for further study at the aggregate level.

To this meeting, Kathy had brought with her papers by researchers at other universities that reported on devices they had built for a similar purpose. Kathy described weaknesses with these past approaches, saying that she wanted to develop a device that avoids these. In an earlier email to the facility’s staff requesting this meeting, she wrote, “I have a conceptual idea, but really no knowledge of what good nanofab solutions might be, so brainstorming is going to be a good first step.”

Staff members suggest fabrication approaches

In the conversation that followed Kathy’s introduction to her project, the five staff members present and Kathy discussed possible designs for the device. Staff members asked questions and proposed possible device structures. Several times throughout the conversation, staff members referred to past projects performed in the lab, or to projects they themselves had done over their career, to suggest possible designs for Kathy’s structure. An exemplary instance of such a suggestion is the following interchange between Rick, the most senior staff person present, and Kathy, about half to two-thirds of the way through the meeting, after they had discussed the different demands that the device would have to satisfy. Rick then suggested the following:

> Rick: We have a [*group*] here, whose name I won't speak, but they're doing [Rick describes a project involving isolating biological material from a liquid]

Kathy protests this definition of her project, saying that "we're doing cells, not molecules."

> Rick: Yes, but you're isolating a small thing in a defined way, and then using statistics to [*analyze the aggregate of these small things*].
>
> Kathy: Yes. At a different scale.
>
> Rick: It's the same philosophical principle although everything else is turned sideways.

Rick's last statement is illustrative of this chapter's thesis: that staff recognized different kinds of resemblances between user projects than users typically did. The excerpt contrasts the perspective taken by the user (i.e., cells and molecules are different objects at different scales) and the perspective taken by the staff member (i.e., they are both small things that one is trying to isolate). Staff members recognized resemblances between user projects based on their reliance on similar structural elements: substructures that played a similar role in the overall functionality of a larger structure.

As the meeting continued, the discussion centered around different device demands: how to evenly space the cells; how to accommodate Kathy's preference to analyze different sizes of cells with the same device yet still sort and isolate them effectively. At the end of the meeting, those present appeared to converge around a design that involved an "arena" of equally spaced "wells" etched in glass, over which the cell-containing solution would be poured, such that the cells settle into their individual wells. Staff recommended that Kathy's group proceed by creating a set of test structures in which they would vary the shape and size of the wells, in order to find the optimal dimensions for their application.

Staff members recount the meeting: “Simple structures”

After the process meeting concluded, I walked with Rick to his office and asked him questions about the process meeting. Rick said that during the process meeting, the staff were trying to get Kathy to talk in terms of “simple structures.” Rick says that Kathy, like many users that ask for staff’s assistance at the start of their project, was rooted in her discipline in a way that was unhelpful for developing a feasible fabrication approach. She was “bogged down by the biology” of her project. Rick continued by saying that in contrast with users, staff were typically not bogged down with the disciplinary knowledge that often hindered users from seeing their projects in buildable terms. He said, “she’s talking about cell segregation, telling us the chemistry, when it’s seems to me after being in there for two minutes, […] that what she wants is different bowls for different cells.” Rick said that when talking with users about their designs, staff tries to get a user to talk about “dumb structures like bowls,” and then staff can say, “here’s all these ways how to make that.” In other words, staff members remembered approaches to fabricating these simple structures – approaches that had been successful in the past in other users’ projects.

Recognizing that Kathy’s project required “different bowls for different cells” led Rick to see resemblances to two projects done earlier in the GNF’s history. Several decades ago, a user had developed a way to fabricate bowls for cells in which to lyse them. Rick described this structure as “a way of isolating a cell in a bowl and see what comes out.” Then, several years later, another user combined the bowls with other facets of a structure to allow for the further analysis of the cell’s contents. After describing these designs, Rick said that “these are the same philosophical concepts” as those that he anticipated Kathy’s device would rely on.

I asked him about other examples of philosophical concepts he encounters. He said there are many: examples include levers, the inclined plane, wheel and axle. These can then be integrated into more complex structures. A transistor, for example, “is just a bunch of electrically different materials stuck next to each other.” A nanofabricated gyroscope relies on a basic structure known as a “comb structure.” He and Adam (who had joined the conversation) then talked about and gave examples of these “philosophical concepts”, which they also referred to as “simple structures,” “motifs,” “primitive structures,” and “baseline structures.” Adam said that a closely related notion was “process lore.” Process lore comprised ways people had developed to deal with a particular part of processing repeatedly. These are “accumulated knowledge.” They are “trues” in the lab. For example, “getting something to stick, getting something not to stick.”

Implications and related to data from other field sites

One way of thinking about a response to the main questions of this thesis – how staff develop and disseminate knowledge among users – is that staff members recognized resemblances between users’ projects that users did not recognize, and that they then used these resemblances to synthesize and share knowledge within and across user communities beyond what users were capable of. In this chapter, I have discussed one kind of resemblance that staff spoke about, namely a set of “baseline structures” that they recognized in the varying projects that users pursued. Interactions between the staff and Kathy show staff proposing a vision of Kathy’s project as a combination of structures that had been built before: wells and channels. Kathy, on the other hand, defined her project as novel compared to the devices that had been developed by her peers for similar purposes. Staff's recognition of the baseline structures required for Kathy's project enabled them to draw parallels with past projects performed in the GNF and, therefore, to establish that Kathy's project was viable, to propose an overall design, and to decide on a first step in fabrication.

These observations align with ways that staff at other facilities described their approach to establishing users' process needs. For example, at the Delta Nanofabrication Facility, a senior staff member described staff's expertise as "making shapes" without needing to know the disciplinary reasoning behind the user's choice of those shapes. He made this comment during a conversation about the quantum computing projects users were doing in the lab. I asked him if the recent surge in quantum computing research had placed any novel demands on the DNF's tool set. He said,

> Josh: It hasn't been anything exciting and new for us. Their designs are important and different. But the process flow they're following is similar to a normal lithography and etch.
>
> Danielle: So, roughly speaking, what do their designs look like?
>
> Josh: I think it's a lot of qubits and things like that. And, I don't know, quantum-confined things. They can show up and say, "Here's the shape I need to make. Help me put that into silicon, or diamond, or whatever." And I don't know where those shapes come from. I assume, somebody's brain somewhere.
>
> […]
>
> That's what we do, is we make shapes. And they don't have meaning, kind of, at our level, because they don't need to. We're not here to do the science of quantum mechanics, we're here to help you make a thing. And so, it doesn't matter if it's a triangle or circle or lines. There are challenges that are different in [*making*] each of those shapes. But that's the challenge and the interesting piece to *us*.

What Josh called "shapes," Rick referred to as "simple structures." Both meant the physical forms of substructures that reoccur across device types and disciplines and that are responsible for a particular aspect of device functionality. They are concrete physical structures with defined properties that can be spoken of without reference to the discipline they'll be used to advance. Because they reoccur in this way, they become the focus of substantial "process lore," an accumulated body of techniques learned in practice and taught in a community of practice (Lave & Wenger, 1991). The data presented in this chapter supports the notion that these kinds of "simple structures" and their associated lore form an important way that staff develop and disseminate knowledge across their facility's user base.

In the preceding empirical chapters, I have endeavored to show how different staff roles at nanofabrication facilities presided over knowledge that was distinct from that of users. This was knowledge necessary for routinely building complex hierarchical structures: control over the properties of materials, substructures, and their integration. In the final chapter that follows, I will summarize my findings, derive conclusions and note implications.

Chapter 8. Conclusion and Discussion

In this dissertation, I have endeavored to portray the work world of an occupation that presides over the materials, tools and techniques that underpin the advanced achievements of a diverse set of user communities. I have documented how one such occupation develops, maintains and extends capabilities that are shared across fields, yet wields them with sufficient precision to advance the leading edge of each. In so doing, my goal has been to name a set of skills overlooked in contemporary accounts of science: the ability to routinely configure known technological possibilities to support novel and demanding ends. After summarizing the findings of this thesis, the discussion turns to some implications of what we have learned.

The guiding question of the thesis has been, what kinds of knowledge do scientific support staff in core facilities possess, and how is this knowledge distinct from that of their facilities' users? I grounded an answer to this question in the account of craft work offered by Douglas Harper (1987). Harper describes the work of Willie, a master metalworker, who meets his rural and remote community's diverse needs for construction and repair. In so doing, he demonstrates what Harper calls "working knowledge." I argued that working knowledge can be understood as a resourcefulness with a set of materials, tools and techniques that a person develops by occupying a central role in a community that relies on these materials, tools and techniques for their diverse objects of use. Accordingly, the person who develops working knowledge presides over a relatively circumscribed set of materials, tools, and techniques that she brings to bear on diverse kinds of structures, the particulars of which are defined largely by the recurring and shifting needs of her community. It is through these engagements with the same materials in shifting configurations that she develops her facility with the enduring properties, recurring problems, and fundamental classes of operations needed to solve problems using her tool set.

At the base of working knowledge lies a fine, embodied understanding of and control over the properties of numerous materials. But an essential feature of working knowledge is that it is applied in the context of structures: systems composed of many simple elements that interact to form successively complex units of functionality. In other words, structures are what Herbert Simon (1962) calls "hierarchically complex systems." Building and fixing such structures involves being able to consider such internally interdependent systems and to recognize – and in some cases, to discover – the unique structural element required for the whole to function. I have called this an ability to recognize and meet the "situated fabrication needs" of complex structures.

Working knowledge is necessary when the structures to be built or repaired are integrated and when fabrication needs are not standardized. I argued that these three features together comprise working knowledge: (1) a fine, embodied understanding of and control over the properties of numerous materials; (2) a central role in a community that cannot rely on standardized solutions for their fabrication and repair problems; (3) an ability to recognize the situated fabrication needs of complex structures.

I have argued that staff members at the nanofabrication facilities in this study occupied a similar role to Willie's in their user communities and therefore developed similar knowledge. These facilities existed to make possible single-batch fabrication of unique structures, and accordingly there was little in the world of academic nanofabrication that was prefabricated. Users had to build their devices from their most basic elements up and lean heavily on the expertise of staff in developing a process with which to do so successfully.

Through experience helping many users over their careers, staff had developed working knowledge of the materials, tools and techniques fundamental to nanofabrication.

Users, by contrast, did not have strong incentives to develop robust knowledge about the elements and techniques of nanofabrication. Their primary interest was to use these elements and techniques to form larger structures that they anticipated would be valuable to their respective fields. Users' distinct contribution to problem-solving, as it became clear in observations and interviews, was their knowledge of their objectives: what exactly did they want their device to do? What modifications made in the course of fabrication could their design tolerate? In other words, their distinct role related to their ability to engage in situated redesign (Rahman & Barley, 2017) in response to staff's targeted probes during process meetings and in response to the experimental results that they themselves obtained in the cleanroom.

Based on this conceptualization of the work of scientific support occupations, we can now specify what we mean by *synthesis* when it is performed by individuals who occupy this role in the division of labor. Staff synthesized knowledge when they used an instance of learning that took place within a distinct *structure* to learn about the constituent *elements* of that structure and their *integration*. In performing this kind of synthesis, staff became repositories for knowledge that would not otherwise be available to those beyond those researchers who had made an original discovery. Following Simon's definition of hierarchically complex systems, we might say that they developed knowledge that clustered around the following domains:

- Ways to handle *materials* to obtain distinct properties
- Ways to shape materials into *structures of intermediate complexity*
- Ways to *integrate* intermediate structures into larger wholes

- Ways to diagnose problems occurring in the fabrication of complex structures

The second guiding research question of the thesis was, how does the staff's distinct knowledge support the dissemination of knowledge among users of their facilities? Users developed solutions to their situated process problems in collaboration with staff. But users typically did not know what purpose these solutions could serve outside of their own project. Staff, on the other hand, recognized resemblances in the fabrication needs of different structures and repurposed solutions developed in the context of other projects accordingly. These resemblances could occur at any level of complexity of a structure, for instance according to materials used, intermediate structures built, integration challenges faced, or a combination of these. Staff members were alert to the versatility of the solutions they developed for any given structure. In effect, staff's working knowledge became a means by which the technical progress made within particular user projects was diffracted into several more fundamental domains of knowledge. From there, staff circulated solutions to fabrication problems by means of technical resemblances, invoked in the context of concrete problem-solving. Based on these findings, I will propose three implications.

Implications for accounts of working knowledge and craft occupations

One contribution this thesis makes is to attempt an explicit definition of *working knowledge*, by identifying a craftsperson's role in a social structure, the problems this person encounters as an occupant of that role, and the abilities this person demonstrates when meeting those problems. This definition makes wider application of the term possible, for instance in contemporary occupational and organizational settings. A second contribution is to move beyond Harper's original conception of working knowledge to show how working knowledge is held and applied by a team of people rather than by a single individual.

Nanofabrication cleanrooms contained more machines than any one staff member could become deeply proficient at. Therefore, as Chapter 4 illustrates, staff divided roles to steward and apply knowledge of what the tools were capable of, what substructures had been built before, and what integration problems they had seen. Staff invited the application of this dispersed knowledge by querying each other around specific problems during process meetings with users. In these same meetings, staff also employed their knowledge of the capabilities of their tools to probe the objectives and tolerances of users' structures, with the goal of anticipating integration problems. In this way, they helped users define their fabrication needs as these were situated within complex structures and multi-step fabrication processes.

These findings drive home the point that staff presided not chiefly over instruments with clear, standardized uses, but rather over an ensemble of machines and a repertoire of techniques from which they configured feasible fabrication approaches to match the objectives of every new user. Therefore, treating staff at such facilities purely as experts on discrete tools is not commensurate with the work they do. A core part of their work involves selecting, applying and combining enduring bodies of technical knowledge to match the unstandardized, diverse, and technically demanding objectives of different lines of research.

Naming working knowledge and the circumstances of its development may enable greater understanding of the role that other typically invisible occupations play in the production of complex goods. For instance, we may expect to find working knowledge among supporting occupations in movie production, such as set designers and special effects engineers (Savage, 2019) and seamstresses, patternmakers and tailors. Similarly positioned might be those occupations that supply artists with specialized materials like pigments, canvases and brushes (Becker, 1982). Other such occupations in scientific settings may include writing experts employed at university writing centers (North, 1984, 1994; Geller, 2007). In addition, there may be resemblances with occupations less typically thought of as craft occupations, but who help individuals with strong goals and little experience assemble feasible plans. For instance, mentors and staff members at accelerators and incubators for new ventures help entrepreneurs develop an approach to their venture but have not received much direct attention in accounts of accelerators' effectiveness (Cohen et al., 2019; Hallen et al., 2020; Krishnan et al., 2021). An interesting question to pursue across these occupations would be how these occupations manage their proximity to highly visible and highly rewarded pursuits, such as movies, the arts, and entrepreneurship, in many cases working a mere "twenty feet from stardom" with equivalent or superior skill to those in the limelight (Neville, 2013). In addition, these are often competitive fields. Because of their central position, individuals with working knowledge may serve clients who compete directly with one another. How do members of these occupations navigate their client networks? How do they build broadly deployable competencies from their engagements in proprietary work? How might their approach enable them to advance the capabilities of their field in ways not open to those engaged in direct competition with one another? Answers to such questions would have implications for scholarly discussions that take place at higher levels of analysis, for instance enabling greater understanding of the roles played by differently incentivized and positioned actors in the development of institutional fields (Whittington et al., 2009; Powell et al., 2012; Owen-Smith & Powell, 2004).

Implications for perspectives on core facilities

In addressing the two research questions of this thesis, this study extends our understanding of how core facilities foster the development and dissemination of knowledge. Early accounts of the benefits of core facilities envisioned interactions between researchers as the primary engine for such knowledge circulation (Lippens et al., 2019; Meder et al., 2016). But as I have shown, the role that researchers play in these facilities does not position them to develop and circulate the kind of knowledge that core facilities are poised to foster. Core facilities are ideally positioned to advance what can be done with the materials, tools and techniques that lie at the base of different lines of research. Researchers' careers predispose them to seek to develop relatively context-independent knowledge claims that have a strong claim to novelty or distinctiveness in their respective fields, rather than to improve their tools in ways that serve a broad set of disciplines (Katz, 2014; Ahalt et al., 2015; Howison & Bullard, 2016).

Shifting our view to other roles in the division of labor at these facilities, the findings show how members of the scientific support staff develop and disseminate technical knowledge among the facilities' users. Technical knowledge was fundamental to the array of projects occurring at the facilities in this study, and support staff leveraged this fundamental knowledge to develop advances that benefited multiple lines of research. Because staff's knowledge related to materials, tools and techniques that occurred across projects, their expertise could advance largely independent of disciplinary anchoring and irrespective of any one project's disciplinary home.

Seen in this light, the distinctive capability of these facilities is to be able to routinely and reliably assemble different complex structures for a range of applications. These organizations specialize in the maintenance and development of versatile, stable subcomponents and the techniques and integration expertise with which to assemble them in shifting combinations for novel, demanding ends. This versatility is a distinctive asset of these organizations, since research and innovation projects are often unprecedented (Owen-Smith, 2018). As the director of one of the facilities articulated when reflecting on her facility's mandate going into the next decade, "We are given the challenge of trying to develop a vision of something that's going to be flexible. Flexible, [*to*] really meet the research challenges of the future, because we don't know what those are yet." I've argued that this flexibility is significantly rooted in the working knowledge of the staff at these facilities, which includes their ability to create continuity with the technical achievements of past generations of researchers.

Implications for support occupations in science

These findings have implications for what we know about support occupations in science (Milojevic et al., 2018; Lee & Walsh, 2021). Past studies have established that those in ideal-typical technician roles are critical for the production of scientific outputs, performing tool-related tasks without which the products of a given lab could not be made (Shapin, 1989; Barley & Bechky, 1994). The findings of this study extend this line of inquiry by documenting how these workers' roles and attendant expertise enable them to link otherwise distant lines of research. Such pathways for the recombination of knowledge have not been recognized by scholars of science and innovation. This is in part because past studies studied technicians employed in single labs (Barley & Bechky, 1994). The oversight is also due to the fact that studies that have mapped the recombination of knowledge in science by using papers and patents to construct their datasets implicitly take the view that only those in investigator roles bridge the boundaries of disciplines (Fleming & Sorenson, 2004; Evans, 2008; Vilhena et al., 2014; Shi et al., 2015; Fortunato et al., 2018). The findings presented here provide an important corrective to this work by showing that scientific support staff who serve multiple projects leverage pathways for recombination that are not apparent to investigators, let alone reflected in their publications or citations. This study therefore extends our knowledge of support occupations in science by documenting how they occupy distinct roles in the flow of knowledge across the landscape of contemporary science.

The work that we see staff members at nanofabrication facilities do also reframes how we conceive of their contribution to the progress of science. Herbert Simon argued that the speed with which hierarchical complexity toward any end could arise depended on the availability and stability of constituent elements. The more stable and abundant the elements and their subsequent assembly into substructures of increasing levels of complexity, the more rapidly the systems built with them could evolve. Because staff members at nanofabrication facilities developed the technical knowledge and sensorimotor skill with which to reliably form and integrate the elements and substructures of nanofabrication, they in effect helped establish this abundance of diverse and stable elements that were available to users, many of whom arrived at a facility having never before touched a wafer or a piece of fabrication equipment. In other words, while it is not inaccurate to say that members of these occupations managed interfaces with complex machines, an important implication is that in that in so doing, they enabled researchers relatively unschooled in technical matters to rapidly attain the level of technical complexity that their research goals required.

Articulating the role, work, and attendant skill of members of these occupations is necessary if universities and science policymakers are to effectively foster science in its modern-day configurations. Contemporary science increasingly relies on specialized support occupations but is uneven in the support and acknowledgement that it provides members of these occupations. A spate of recent research has identified the increasing size of scientific "teams" (Wu et al., 2017; Fortunato et al., 2018; Hall et al., 2018), a trend that has involved a significant shift toward bureaucratic forms of organizing scientific work, with highly specialized divisions of labor (Lee & Walsh, 2021). Illustrating this trend, Milojevic and her coauthors quantitatively analyzed the publication careers of authors in astronomy, ecology and robotics, and identified a role they call the "supporting author scientist." These scientists had Ph.D. degrees and appeared on author lists throughout their career, without ever serving as first author (p. 12616). The implication is that the growth in the average number of authors per paper is not simply the result of greater collaboration among tenure track scientists, "but rather involve[*s*] the recruitment of a special workforce of supporting scientists." (p. 12618)[9]

The shift toward such specialization is not surprising, given the increasing emphasis on productivity in science and the greater capital-intensity of scientific work (Stephan, 2012). But it is an open question whether the career structures that scientific fields offer these workers have caught up with this state of affairs. Universities often offer these individuals only temporary contracts, leading to high attrition (Milojevic et al., 2018). But even when longer-term positions are offered to staff members, as they are at core facilities, turnover is frequent. Those who administer core facilities note that it is difficult to retain skilled staff due to the lack of appealing career paths (Ferrando-May et al., 2016; Meder et al., 2016; Kos-

---

[9] As helpful as such bibliometric studies are in uncovering hidden roles in scientific divisions of labor, they likely underestimate the size of this supporting scientific workforce because many supporting scientific workers are not credited with authorship.

Braun, Gerlach & Pitzer, 2020). High turnover of staff at core facilities results in the loss of valuable expertise and risks the continuity of scientific research.

An avenue for future research would be to examine the factors underlying turnover and retention of staff at core facilities. To address such a question would require data on the career paths of individuals employed at core facilities, as well as those who leave for other jobs. Ideally, such structured data would be complemented with interviews with those who stayed and left, to capture what these career moves mean to them. Understanding this would be a useful step toward creating more fitting job roles and rewarding career structures for these members of the scientific workforce.

Closing comments

As perhaps they should be, the research challenges of the future are unknown. But what is certain is that in addressing these challenges, researchers will reach for the accumulated tools and techniques of their day. Recognizing that scientific support occupations at core facilities preside over the expertise to direct these tools toward precise ends makes plain that they are crucial actors in fostering the flexibility that research universities must exhibit. So, although I would like these findings to shine light on the knowledgeable work that staff do with complicated machines and specialized techniques, this dissertation argues that this technical knowledge culminates in something more. I would like to have people think about support staff not *only* as an interface with these tools, materials and techniques (and they are this too), but as central craftspeople that help their researchers assemble research projects that have a technical complexity and precision to them. By presiding over the knowledge necessary to routinely build complex systems, they are a crucial form of memory at the university that is not only related to operating complicated equipment but in focusing the capabilities of a repertoire of tools and techniques to very fine ends. In so doing, they help researchers access, leverage, and recombine enduring bodies of technical knowledge that have been developed in the technically complex projects pursued in the past. If well-funded, supported and valued, they can be the master builders at the center of a dynamic, interwoven and forward-looking research community.

References

Agar, M. (1996). *The Professional Stranger: An Informal Introduction to Ethnography.* Emerald Group Publishing Limited.

Ahalt, S., Carsey, T., Couch, A., Hooper, R., Ibanez, L., Idaszak, R., Jones, M. B., Lin, J., & Robinson, E. (2015). *NSF Workshop on Supporting Scientific Discovery through Norms and Practices for Software and Data Citation and Attribution* [Technical Report]. National Science Foundation.

Allen, T. (1977). *Managing the Flow of Innovation*. MIT Press.

Bailey, D. E., & Barley, S. R. (2011). Teaching-Learning Ecologies: Mapping the Environment to Structure through Action. *Organization Science*, *22*(1), 262–285. https://doi.org/10.1287/orsc.1090.0511

Barley, S. R. (1986). Technology as an Occasion for Structuring: Evidence from Observations of CT Scanners and the Social Order of Radiology Departments. *Administrative Science Quarterly*, *31*(1), 78. https://doi.org/10.2307/2392767

Barley, S. R. (1989). Careers, Identities, and Institutions: The legacy of the Chicago School of Sociology. In M. B. Arthur, D. T. Hall, & B. S. Lawrence (Eds.), *Handbook of Career Theory* (pp. 41–65). Cambridge University Press. https://doi.org/10.1017/CBO9780511625459.005

Barley, S. R. (1990a). Images of Imaging: Notes on Doing Longitudinal Field Work. *Organization Science*, *1*(3), 220–247. https://doi.org/10.1287/orsc.1.3.220

Barley, S. R. (1990b). The Alignment of Technology and Structure through Roles and Networks. *Administrative Science Quarterly*, *35*(1), 61. https://doi.org/10.2307/2393551

Barley, S. R. (1996). Technicians in the Workplace: Ethnographic Evidence for Bringing Work into Organizational Studies. *Administrative Science Quarterly*, *41*(3), 404. https://doi.org/10.2307/2393937

Barley, S. R., & Bailey, D. E. (2020). Managing the Fears of Studying Technical Work. In S. R. Barley, *Work and Technological Change* (pp. 116–144). Oxford University Press.

Barley, S. R., & Bechky, B. A. (1994). In the Backrooms of Science: The Work of Technicians in Science Labs. *Work and Occupations*, *21*(1), 85–126. https://doi.org/10.1177/0730888494021001004

Barley, S. R., Hofstader, R. and Chapman, K. (1997). Skill Standards in Context: Models of Technician's Work. In R. Hofstader and K. Chapman (Eds.), *Foundations for Excellence in the Chemical Process Industries: Voluntary Industry Standards for Chemical Process Industries Technical Workers* (pp. 13-25). American Chemical Society: Washington, D.C.

Basalla, G. (1988). *The Evolution of Technology.* Cambridge University Press.

Basic Energy Sciences Advisory Committee, Department of Energy. (2018). *A Remarkable Return on Investment in Fundamental Research: 40 years of basic energy sciences at the Department of Energy*. https://science.osti.gov/-/media/bes/pdf/BESat40/BES_at_40.pdf

Becker, H. S. (1982). *Art worlds.* University of California Press.

Bikovski, L., Robinson, L., Konradsson-Geuken, A., Kullander, K., Viereckel, T., Winberg, S., Roman, E., & Tsoory, M. (2020). Lessons, insights and newly developed tools emerging from behavioral phenotyping core facilities. *Journal of Neuroscience Methods*, *334*, 108597. https://doi.org/10.1016/j.jneumeth.2020.108597

Bovenberg, D.E. (2021). Safe Crossings: The Role of Scientific Support Occupations in Attenuating Secrecy in Competitive Scientific Communities. *Technology Management Department working paper, UC Santa Barbara.*

Burt, R.S. (2004). Structural Holes and Good Ideas. *American Journal of Sociology 110*(2), 349–399.

Chang, M., & Grieder, F. B. (2016). Sharing Core Facilities and Research Resources—An Investment in Accelerating Scientific Discoveries. *Journal of Biomolecular Techniques: JBT*, *27*(1), 2–3. https://doi.org/10.7171/jbt.16-2701-004

Cohen, S.L., Bingham, C.B., Hallen, B.L. (2019). The Role of Accelerator Designs in Mitigating Bounded Rationality in New Ventures. *Administrative Science Quarterly 64*(4), 810–854.

Collins, H. M. (1974). The TEA Set: Tacit Knowledge and Scientific Networks. *Science Studies*, *4*(2), 165–185. https://doi.org/10.1177/030631277400400203

Darr, A., & Scarselletta, M. (2002). Technicians, Clients and Professional Authority: Structured Interactions and Identity Formation in Technical Work. *New Technology, Work and Employment 17*, 60-72

Doing, P. (2009). *Velvet Revolution at the Synchrotron: Biology, Physics, and Change in Science.* MIT Press.

Evans, J. A. (2008). Electronic Publication and the Narrowing of Science and Scholarship. *Science*, *321*(5887), 395–399. https://doi.org/10.1126/science.1150473

Ferrando-May, E., Hartmann, H., Reymann, J., Ansari, N., Utz, N., Fried, H.-U., Kukat, C., Peychl, J., Liebig, C., Terjung, S., Laketa, V., Sporbert, A., Weidtkamp-Peters, S., Schauss, A., Zuschratter, W., & Avilov, S. (2016). Advanced light microscopy core facilities: Balancing service, science and career. *Microscopy Research and Technique*, *79*(6), 463–479. https://doi.org/10.1002/jemt.22648

Fine, G. A. (1996). *Kitchens: The Culture of Restaurant Work.* University of California Press.

Fleming, L., & Sorenson, O. (2004). Science as a map in technological search. *Strategic Management Journal*, *25*(8–9), 909–928. https://doi.org/10.1002/smj.384

Fortunato, S., Bergstrom, C. T., Börner, K., Evans, J. A., Helbing, D., Milojević, S., Petersen, A. M., Radicchi, F., Sinatra, R., Uzzi, B., Vespignani, A., Waltman, L., Wang, D., & Barabási, A.-L. (2018). Science of science. *Science*, *359*(6379), eaao0185. https://doi.org/10.1126/science.aao0185

Fu, L., Datta, K. K. R., Spyrou, K., Qi, G., Sardar, A., Khader, M. M., Zboril, R., & Giannelis, E. P. (2017). Phyllosilicate nanoclay-based aqueous nanoparticle sorbent for CO_2 capture at ambient conditions. *Applied Materials Today*, *9*, 451–455. https://doi.org/10.1016/j.apmt.2017.09.009

Galison, P. (1997). *Image and Logic: A material culture of microphysics*. University of Chicago Press.

Galison, P., Hevly, B.W. (Eds.). (1992). *Big Science: The growth of large-scale research*. Stanford University Press.

Gaston, J. C. (1971). Secretiveness and competition for priority of discovery in physics. *Minerva*, *9*, 472–492.

Gavetti, G., & Levinthal, D. (2000). Looking forward and looking backward: Cognitive and experiential search. *Administrative Science Quarterly*, *45*(1), 113–137.

Geertz, C. (1973). Thick Description: Toward an Interpretive Theory of Culture. In *The Interpretation of Cultures* (pp. 3 – 32). Basic Books, Inc., Publishers.

Geller, A. E. (2007). *The Everyday Writing Center: A community of practice*. Utah State University Press.

Gilfillan, S. (1935). *Inventing the Ship.* Follett Publishing Company.

Gioia, D. A., Corley, K. G., & Hamilton, A. L. (2013). Seeking Qualitative Rigor in Inductive Research: Notes on the Gioia Methodology. *Organizational Research Methods, 16*(1), 15–31. https://doi.org/10.1177/1094428112452151

Gould, J. (2015). Core facilities: Shared support. *Nature*, *519*, 495–496.

Haas, M. R., & Park, S. (2009). To Share or Not to Share? Professional Norms, Reference Groups, and Information Withholding Among Life Scientists. *Organization Science*, *21*(4), 873–891. https://doi.org/10.1287/orsc.1090.0500

Hall, K. L., Vogel, A. L., Huang, G. C., Serrano, K. J., Rice, E. L., Tsakraklides, S. P., & Fiore, S. M. (2018). The science of team science: A review of the empirical evidence and research gaps on collaboration in science. *American Psychologist, 73*(4), 532–548. https://doi.org/10.1037/amp0000319

Hallen, B.L., Cohen, S.L., & Bingham, C.B. (2020). Do Accelerators Work? If So, How? *Organization Science 31*(2), 378–414.

Hargadon, A., & Sutton, R. I. (1997). Technology Brokering and Innovation in a Product Development Firm. *Administrative Science Quarterly*, *42*(4), 716–749. https://doi.org/10.2307/2393655

Harper, D. (1987). *Working Knowledge: Skill and Community in a Small Shop.* University of Chicago Press.

Hockberger, P., Weiss, J., Rosen, A., & Ott, A. (2018). Building a Sustainable Portfolio of Core Facilities: A Case Study. *Journal of Biomolecular Techniques: JBT*, *29*(3), 79–92. https://doi.org/10.7171/jbt.18-2903-003

Howison, J., & Bullard, J. (2016). Software in the scientific literature: Problems with seeing, finding, and using software mentioned in the biology literature. *Journal of the Association for Information Science and Technology, 67*(9), 2137–2155. https://doi.org/10.1002/asi.23538

Jokerst, N.M., Donley, C., Cahoon, J., & Jones, J. (2018). *Nanotechnology: A maker's course* [MOOC]. Coursera. https://www.coursera.org/learn/nanotechnology

Jordan, K., & Lynch, M.E. (1992). The Sociology of a Genetic Engineering Technique: Ritual and Rationality in the Performance of the "Plasmid Prep." In A.E. Clark & J.H. Fujimura (Eds.), *The Right Tools for the Job: At Work in Twentieth-Century Life Sciences* (pp. 77-114). Princeton University Press.

Kaplan, S., Milde, J., & Cowan, R.S. (2017). Symbiont Practices in Boundary Spanning: Bridging the Cognitive and Political Divides in Interdisciplinary Research. *Academy of Management Journal 60*(4),1387–1414.

Katila, R., & Ahuja, G. (2002). Something Old, Something New: A Longitudinal Study of Search Behavior and New Product Introduction. *Academy of Management Journal*, *45*(6), 1183–1194. https://doi.org/10.2307/3069433

Katz, D. S. (2014). Transitive Credit as a Means to Address Social and Technological Concerns Stemming from Citation and Attribution of Digital Products. *Journal of Open Research Software, 2*(1), Article 1. https://doi.org/10.5334/jors.be

Kos-Braun, I. C., Gerlach, B., & Pitzer, C. (2020). A survey of research quality in core facilities. *ELife*, *9*, e62212. https://doi.org/10.7554/eLife.62212

Knorr-Cetina, K. (1999). *Epistemic Cultures: How the Sciences Make Knowledge*. Harvard University Press.

Krishnan, R., Cook, K. S., Kozhikode, R. K., & Schilke, O. (2021). An Interaction Ritual Theory of Social Resource Exchange: Evidence from a Silicon Valley Accelerator. *Administrative Science Quarterly*, *66*(3), 659–710. https://doi.org/10.1177/0001839220970936

Lave, J., & Wenger, E. (1991). *Situated Learning: Legitimate Peripheral Participation.* Cambridge University Press.

Lee, Y.-N., & Walsh, J. P. (2021). Rethinking Science as a Vocation: One Hundred Years of Bureaucratization of Academic Science. *Science, Technology, & Human Values*, 01622439211026020. https://doi.org/10.1177/01622439211026020

Leonardi, P. M. (2009). Why Do People Reject New Technologies and Stymie Organizational Changes of Which They Are in Favor? Exploring Misalignments Between Social Interactions and Materiality. *Human Communication Research*, *35*(3), 407–441. https://doi.org/10.1111/j.1468-2958.2009.01357.x

Levinthal, D. A. (1997). Adaptation on rugged landscapes. *Management Science*, *43*(7), 934–950.

Levinthal, D. A., & March, J. G. (1993). The myopia of learning. *Strategic Management Journal*, *14*(S2), 95–112.

Lippens, S., D'Enfert, C., Farkas, L., Kehres, A., Korn, B., Morales, M., Pepperkok, R., Premvardhan, L., Schlapbach, R., Tiran, A., Meder, D., & Van Minnebruggen, G. (2019). One step ahead: Innovation in core facilities. *EMBO Reports*, *20*(4). https://doi.org/10.15252/embr.201948017

Long Lingo, E., & O'Mahony, S. (2010). Nexus Work: Brokerage on Creative Projects. *Administrative Science Quarterly, 55*(1), 47–81. https://doi.org/10.2189/asqu.2010.55.1.47

Lynch, M. E. (1985). *Art and Artifact in Laboratory Science: A Study of Shop Work and Shop Talk in a Research* Laboratory. Routledge and Kegan Paul.

McCray, W. P. (2006). *Giant Telescopes: Astronomical Ambition and the Promise of Technology*. Harvard University Press.

Meder, D., Morales, M., Pepperkok, R., Schlapbach, R., Tiran, A., & Van Minnebruggen, G. (2016). Institutional core facilities: Prerequisite for breakthroughs in the life sciences. *EMBO Reports*, *17*(8), 1088–1093. https://doi.org/10.15252/embr.201642857

Merton, R. K. (1942). *The Sociology of Science: Theoretical and Empirical Investigations*. University of Chicago Press.

Milojevic, S., Radicchi, F., & Walsh, J. P. (2018). Changing demographics of scientific careers: The rise of the temporary workforce. *PNAS*, *115*(50), 12616–12623.

Mody, C. C. M. (2011). *Instrumental Community: Probe Microscopy and the Path to Nanotechnology*. MIT Press.

Mody, C. C. M. (2017). *The Long Arm of Moore's Law: Microelectronics and American Science.* MIT Press.

Mody, C. C. M., & Choi, H. (2013). From Materials Science to Nanotechnology: Interdisciplinary Center Programs at Cornell University, 1960–2000. *Historical Studies in the Natural Sciences, 43*(2), 121–161. https://doi.org/10.1525/hsns.2013.43.2.121

Nadel, S. F. (1957). *The Theory of Social Structure*. Routledge.

National Nanotechnology Coordinated Infrastructure (2016). "About the NNCI." Retrieved on March 18, 2023 from https://nnci.net/about-nnci.

National Nanotechnology Coordinated Infrastructure. (2017). *NNCI Coordinating Office Annual Report (Year 1).* Retrieved April 27, 2023 from https://nnci.net/sites/default/files/inline-files/NNCI%20CO%20Annual%20Report%202017%20Public.pdf

National Science Foundation (2020). *NSF renews investment in National Nanotechnology Coordinated Infrastructure.* (August 24, 2020). Retrieved April 27, 2023, from https://www.nsf.gov/news/news_summ.jsp?cntn_id=301090

Nelson, A. J. (2016). How to Share "A Really Good Secret": Managing Sharing/Secrecy Tensions around Scientific Knowledge Disclosure. *Organization Science*, *27*(2), 265–285. https://doi.org/10.1287/orsc.2015.1040

Neville, M. (Director). (2013, July 26). *Twenty Feet from Stardom* [Documentary, Biography, History]. Gil Friesen Productions, Tremolo Productions.

Nicolini, D., Pyrko, I., Omidvar, O., & Spanellis, A. (2022). Understanding Communities of Practice: Taking Stock and Moving Forward. *Academy of Management Annals 16*(2), 680 - 718. https://doi.org/10.5465/annals.2020.0330

North, S. M. (1984). The Idea of a Writing Center. *College English*, *46*(5), 433. https://doi.org/10.2307/377047

North, S. M. (1994). Revisiting "The Idea of a Writing Center." *The Writing Center Journal*, *15*(1), 7–19.

Nosrat, S. (2017). *Salt, Fat, Acid, Heat.* Simon & Schuster, Inc.

NSF Office of Advanced Cyberinfrastructure, Directorate for Computer & Information Science & Engineering (2019). *Transforming Science Through Cyberinfrastructure: NSF's Blueprint for a National Cyberinfrastructure Ecosystem for Science and Engineering in the 21st Century*. National Science Foundation.

Orr, J. E. (1996). *Talking about Machines: An Ethnography of a Modern Job*. Cornell University Press.

Owen-Smith, J. (2018). *Research Universities and the Public Good: Discovery for an uncertain future*. Stanford Business Books.

Owen-Smith, J., Powell, W.W. (2004). Knowledge Networks as Channels and Conduits: The Effects of Spillovers in the Boston Biotechnology Community. *Organization Science 15*(1), 5–21.

Perkmann, M., McKelvey, M., & Phillips, N. (2018). Protecting Scientists from Gordon Gekko: How Organizations Use Hybrid Spaces to Engage with Multiple Institutional Logics. *Organization Science 30*(2), 298–318.

Powell, W.W., Packalen, K., & Whittington, K. (2012). Organizational and Institutional Genesis: The Emergence of High-Tech Clusters in the Life Sciences. In J.F. Padgett & W.W. Powell (Eds.), *The Emergence of Organizations and Markets* (pp. 434–465). Princeton University Press.

Rahman, H. A., & Barley, S. R. (2017). Situated Redesign in Creative Occupations – An Ethnography of Architects. *Academy of Management Discoveries, 3*(4), 404–424. https://doi.org/10.5465/amd.2016.0039

Savage, A. (2019). *Every Tool's a Hammer*. Simon & Schuster, Inc.

Shapin, S. (1989). The Invisible Technician. *American Scientist, 77*(6), 554–563.

Shi, F., Foster, J. G., & Evans, J. A. (2015). Weaving the fabric of science: Dynamic network models of science's unfolding structure. *Social Networks*, *43*, 73–85. https://doi.org/10.1016/j.socnet.2015.02.006

Sia, S. K., & Owens, M. P. (2015). Share and share alike: How the sharing economy is influencing biotech innovation and what it means for startups. *Nature Biotechnology*, *33*(12), 1224–1228.

Simon, H. A. (1962). The Architecture of Complexity. *Proceedings of the American Philosophical Society, 106*(6), 467–482.

Smith, M., & Smith, J. C. (2020). *Repurposing Therapeutics for COVID-19: Supercomputer-Based Docking to the SARS-CoV-2 Viral Spike Protein and Viral Spike Protein-Human ACE2 Interface*. https://doi.org/10.26434/chemrxiv.11871402.v4

Spradley, J. P. (1979). *The Ethnographic Interview*. Holt, Rinehart and Winston.

Star, S. L., & Griesemer, J. R. (1989). Institutional Ecology, 'Translations' and Boundary Objects: Amateurs and Professionals in Berkeley's Museum of Vertebrate Zoology, 1907-39. *Social Studies of Science*, *19*(3), 387–420. https://doi.org/10.1177/030631289019003001

Stephan, P. E. (2012). *How Economics Shapes Science*. Harvard University Press.

Strauss, A., & Corbin, J. M. (1990). *Basics of Qualitative Research: Grounded Theory Procedures and Techniques*. Sage Publications, Inc.

Suchman, L. A. (1987). *Plans and Situated Action: The Problem of Human-Machine Communication*. Cambridge University Press.

Tuertscher, P., Garud, R., & Kumaraswamy, A. (2014) Justification and Interlaced Knowledge at ATLAS, CERN. *Organization Science 25*(6), 1579–1608.

Traweek, S. (1988). *Beamtimes and Lifetimes: The World of High Energy Physics*. Harvard University Press.

U.S. Department of Energy. (2015). *Office of Science User Facilities, Fiscal Year 2015*. Retrieved on August 10, 2021 from https://science.osti.gov/-/media/_/pdf/user-facilities/Reports/DOE-SC-User-Facilities-FY2015-report.pdf?la=en&hash=D7AC49453FB0EE861717409E3DD1C247994D7495

Van Maanen, J. (1990). Escape from Modernity: On the Ethnography of Repair and the Repair of Ethnography. *Human Studies*, *13*(3), 275–284.

Van Maanen, J. (2011). Ethnography as Work: Some Rules of Engagement. *Journal of Management Studies*, *48*(1), 218–234. https://doi.org/10.1111/j.1467-6486.2010.00980.x

Vertesi, J. (2015). *Seeing Like a Rover: How Robots, Teams, and Images Craft Knowledge of Mars*. University of Chicago Press. https://doi.org/10.7208/chicago/9780226156019.001.0001

Vertesi, J. (2020). *Shaping Science: Organizations, Decisions, and Culture on NASA's Teams*. University of Chicago Press.

Vilhena, D., Foster, J., Rosvall, M., West, J., Evans, J., & Bergstrom, C. (2014). Finding Cultural Holes: How Structure and Culture Diverge in Networks of Scholarly Communication. *Sociological Science*, *1*, 221–238. https://doi.org/10.15195/v1.a15

Walsh, J. P., & Hong, W. (2003). Secrecy is increasing in step with competition. *Nature*, *422*(6934), 801.

Whalley, P., & Barley, S.R. (1997). Technical Work and the Division of Labor: Stalking the wily anomaly. In S. Barley and J. Orr (Eds.), *Between Craft and Science: Technical Work in U.S. Settings* (pp. 20-52). ILR Press.

Whittington, K.B., Owen-Smith, J., & Powell, W.W. (2009). Networks, Propinquity, and Innovation in Knowledge-intensive Industries. *Administrative Science Quarterly* *54*(1), 90–122.

Wolf, E. (2007, June 14). A Personalized Summary of CNF History. Future of Nanotechnology Symposium (Thirtieth Anniversary of Cornell Nanoscale Science and Technology Facility).

Wu, L., Wang, D., & Evans, J. A. (2017). *Large Teams Have Developed Science and Technology; Small Teams Have Disrupted It* (SSRN Scholarly Paper No. 3034125). Social Science Research Network. https://doi.org/10.2139/ssrn.3034125

Zabusky, S. E., & Barley, S.R. (1996). Redefining Success: Ethnographic observations on the careers of technicians. In P. Osterman (Ed.), *Broken Ladders: Managerial Careers in Transition* (pp. 185-214). Oxford University Press.

Zhang, S., & Chen, Y. (2015). Nanofabrication and coloration study of artificial Morpho butterfly wings with aligned lamellae layers. *Scientific Reports*, *5*(1), 16637. https://doi.org/10.1038/srep16637

Appendix: Tables and Figures

Table 1. Membership of NSF-funded networks of academic nanofabrication facilities between 1977 and the present day. Sources: Mody (2017) and National Nanofabrication Coordinated Infrastructure (2016).

1977 - 1993	**1993 - 2003**	**2004 - 2015**	**2016 - 2026**
National Research and Resource Facility for Submicron Structures (NRRFSS, 1977 – 1987); *National Nanofabrication Facility* (NNF, 1987 – 1993)	*National Nanotechnology Users Network* (NNUN)	*National Nanotechnology Infrastructure Network* (NNIN)	*National Nanofabrication Coordinated Infrastructure* (NNCI)
Cornell University	Cornell University	Cornell University	Cornell University
	Stanford University	Stanford University	Stanford University
	Pennsylvania State University	Pennsylvania State University	University of Pennsylvania Community College of Philadelphia
	Howard University	Howard University	North Carolina State University Duke University University of North Carolina
	University of California, Santa Barbara	University of California, Santa Barbara	University of California, San Diego
		Harvard University	Harvard University
		University of Michigan	University of Louisville University of Kentucky
		Georgia Institute of Technology	Georgia Institute of Technology Joint School of Nanoscience and Nanoengineering
		University of Texas at Austin	University of Texas at Austin
		University of Washington	University of Washington Oregon State University
		University of Minnesota	University of Minnesota
		Arizona State University	Arizona State University Northern Arizona University, Rio Salado College, Science Foundation Arizona
		Washington University in St. Louis	Northwestern University University of Chicago
		University of Colorado at Boulder	University of Nebraska at Lincoln

			Montana State University Carleton College
			Virginia Tech Pacific Northwest National Laboratory

Table 2. Key terms in nanofabrication defined.

Term	**Definition**
Device	The user's structure to be fabricated, typically produced on a wafer using nanofabrication process. E.g., a fabricated butterfly wing, an integrated circuit
Nanofabrication process	The set of instructions that, when followed, results in the fabrication of a structure. See Figure 2 for an example of a process
Sample	A term used to refer to the device while it is being fabricated (also referred to as a wafer)
Tool	A piece of equipment used in nanofabrication processing. Common tools include machines for photolithography, etching, deposition and characterization. Parameter specifications for these tools are referred to as recipes.
Recipe	A set of parameters entered into the interface of a tool to control its operation
Photolithography	A class of fabrication steps aimed at transferring a pattern from an original design to the sample using light and light-sensitive chemicals (called photoresist)
Etching	A class of fabrication steps aimed at removing material from the surface of a sample, through either chemical or physical processes
Deposition	A class of fabrication steps aimed at applying a layer of material on the surface of a sample, through a variety of methods
Process integration	Ensuring that each functional element of a device is fabricated in such a way that its fabrication does not adversely affect the ultimate functioning of other parts.

Table 3. A list of definitions of process integration provided by staff members at the Gamma fieldsite during interviews and observations.

A lot of what we have to do in this world is thinking about an entire process from start to finish, all at once, in your head, in all the different areas of the lab and regions of tooling that you would use in the laboratory, to create that device, and thinking about all the possible interactions and what could be going on there. And not all these interactions are obviously positive things.

– Adam, Gamma facility, to prospective users

Process integration is, think of integrating all the processes to make a device. Where you have lithography, dep, etch, all the individual processes that we specialize in. But you need someone to see the whole picture. And to select the materials and the techniques that are conducive to make the device work at the end. Depending on temperature, material compatibility, things like that. So every step along the way has to be thought of very carefully in terms of the end result: is this step going to ruin something along the way? So that's the whole vision, you kind of lay out the whole process, rather than just concentrate on etch or dep.

– Victor, Gamma facility, interview

Integration, where you're looking at the net result of a collection of activities that give you some kind of final product. […] It's a lot like cooking, baking, any kind of mechanical operation. If you're going to make something out of clay, [*and*] the end product is going to be fired at a high temperature, then you would be wise not to make parts of it out of wax. It's a very simple logic but it extends to everything that you're going to be doing.

– Rick, Gamma facility, to prospective users

So, what is a process flow? A process flow is a bunch of these discrete modules all put together in a sensible fashion, in a logical fashion that works. You have to be thinking about all the steps simultaneously, in a way. This is really a difficult thing. So you're sort of taking all this all at once and thinking about, "Okay, if I have to do this implantation at a certain step, but I also want these high-temperature films on this wafer, how do I do that? Because if I take this wafer after I do an implantation on it, and I put it through some other higher-temperature process, well, that implantation, that dopant's going to diffuse everywhere right? I'm not going to get the [*dopant*] that I want, where I want it." So that's that whole thing of structuring that process flow.

– Adam, Gamma facility, to prospective users

…having the skill to understand how a range of tools work, so you know what pitfalls to avoid, and what opportunities you can create because of these things.

– Laurence, Gamma facility, interview

I think it's that interface that makes the incredible things possible. So if I gave you a bunch of modules, it would be like Legos, which are a ton of fun. But don't build your car out of Legos or a boat out of Legos, right? You just can't. You can build something that kind of looks like a car or a boat. But you got to get in there and actually [*integrate*].

– Jason, Gamma facility, interview

Table 4. A list of process integration challenges described by staff members at the Gamma facility, with definitions and illustrative approaches to attenuating these challenges.

Process integration challenge	Description of challenge	Illustrative attenuation approaches
Managing thin film stress	When layering different materials on top of each other, their different thermal expansion properties result in stress at the interfaces of these layers that ramifies throughout the structure. The resulting stress can lead to the materials peeling off from the wafer or the wafer warping to such an extent that it cannot be processed within certain tools.	• Tune film stress with the tool • Select a lower-stress material • Apply a compensating layer to the reverse side of the sample
Attending to thermal budgets and ceilings	Elements of a structure lose their properties when exposed to temperatures above a given level.	Reconfigure the sequence of processing to perform thermally intensive processing before thermally sensitive layers are applied
Maintaining surface electrical properties	Elements of a structure whose electrical properties are central to the functioning of a structure lose their properties when exposed to high currents or ion bombardment during processing.	• Apply sacrificial layers • Use fabrication methods that do not bombard the sample with electrically charged particles
Ensuring chemical compatibility when removing material from constructed stacks	When a structure contains several layers of different materials, the chemicals used to remove one layer may also damage other layers.	• Choose appropriate chemistry • Apply sacrificial layers
Maintaining alignment of features across device layers	When a structure contains several layers, applying these layers may remove or cover up the marks	Apply new alignment marks every new layer

	used for the alignment of features across these layers.	
Monitoring the sample noninvasively throughout the process	Elements of a structure are sensitive to modes of measurement that are nonetheless necessary to inspect the sample during processing.	• Build test structures on representative regions of the sample • Look for non-invasive ways to monitor the sample
Abiding by the cleanroom's system of materials restrictions	The cleanroom maintains strict materials segregation rules, and each process may only proceed from more restrictive tool classes to less restrictive ones (also known as moving from "clean" to "dirty" tools).	Reconfigure the sequence of processing to move from "clean" to "dirty" tools and never in the reverse direction

Figure 1. Scanning electron microscope (SEM) images of (a) cross-section of the surface structure of iridescent *Morpho* butterfly wing and (b) cross-section of the surface structure of a material built by nanofabrication to mimic the blue iridescent coloration of the *Morpho* butterfly wing (Zhang & Chen, 2015)

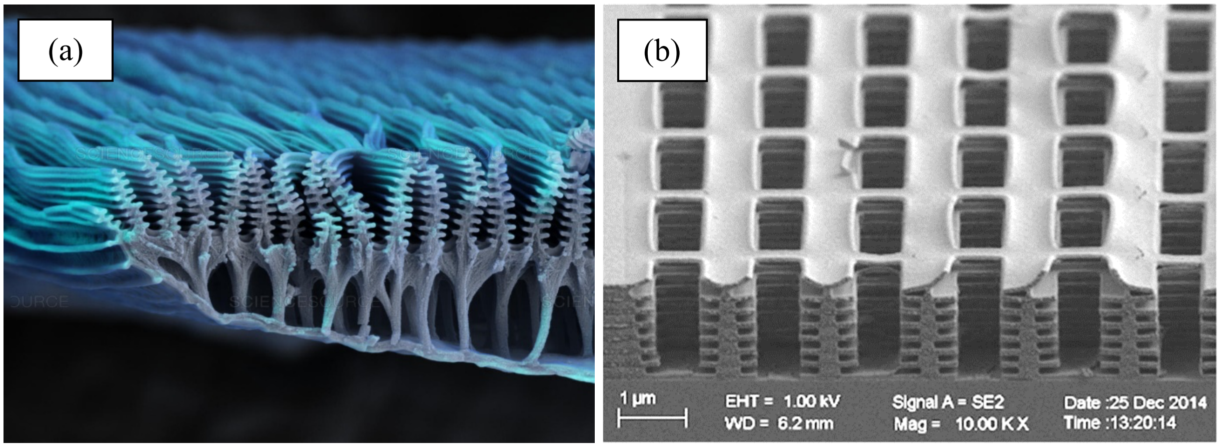

Figure 2. A portion of a finished nanofabrication process developed at a field site.

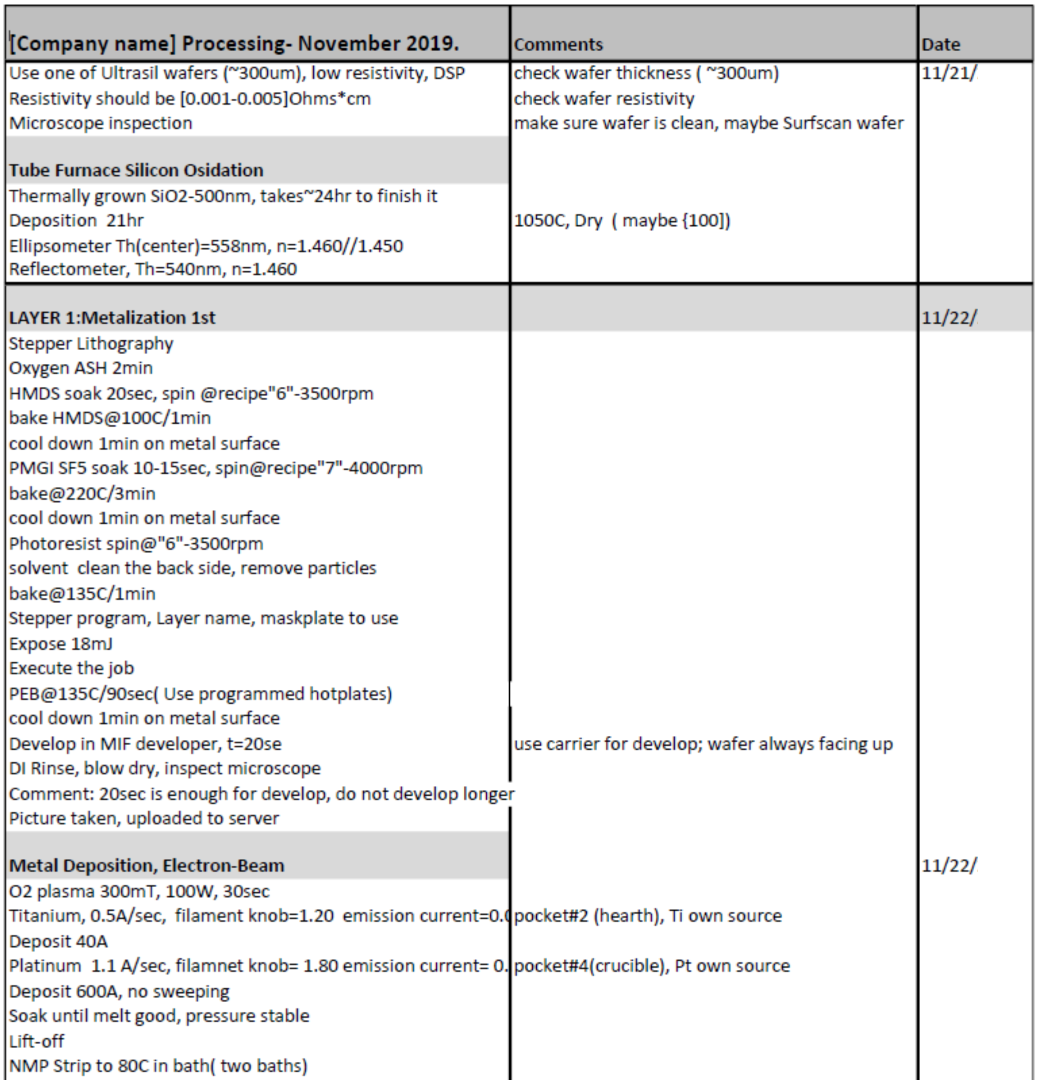

[Company name] Processing- November 2019.	Comments	Date
Use one of Ultrasil wafers (~300um), low resistivity, DSP	check wafer thickness (~300um)	11/21/
Resistivity should be [0.001-0.005]Ohms*cm	check wafer resistivity	
Microscope inspection	make sure wafer is clean, maybe Surfscan wafer	
Tube Furnace Silicon Osidation		
Thermally grown SiO2-500nm, takes~24hr to finish it		
Deposition 21hr	1050C, Dry (maybe {100])	
Ellipsometer Th(center)=558nm, n=1.460//1.450		
Reflectometer, Th=540nm, n=1.460		
LAYER 1:Metalization 1st		11/22/
Stepper Lithography		
Oxygen ASH 2min		
HMDS soak 20sec, spin @recipe"6"-3500rpm		
bake HMDS@100C/1min		
cool down 1min on metal surface		
PMGI SF5 soak 10-15sec, spin@recipe"7"-4000rpm		
bake@220C/3min		
cool down 1min on metal surface		
Photoresist spin@"6"-3500rpm		
solvent clean the back side, remove particles		
bake@135C/1min		
Stepper program, Layer name, maskplate to use		
Expose 18mJ		
Execute the job		
PEB@135C/90sec(Use programmed hotplates)		
cool down 1min on metal surface		
Develop in MIF developer, t=20se	use carrier for develop; wafer always facing up	
DI Rinse, blow dry, inspect microscope		
Comment: 20sec is enough for develop, do not develop longer		
Picture taken, uploaded to server		
Metal Deposition, Electron-Beam		11/22/
O2 plasma 300mT, 100W, 30sec		
Titanium, 0.5A/sec, filament knob=1.20 emission current=0.(	pocket#2 (hearth), Ti own source	
Deposit 40A		
Platinum 1.1 A/sec, filamnet knob= 1.80 emission current= 0.	pocket#4(crucible), Pt own source	
Deposit 600A, no sweeping		
Soak until melt good, pressure stable		
Lift-off		
NMP Strip to 80C in bath(two baths)		

Figure 3. Timeline of data collection between 2018 and 2023 across four field sites, including two phases of analysis and writing. Thick horizontal lines indicate fieldwork; dashed horizontal lines indicate analysis and writing. The relative spacing of events on the timeline at the base of the diagram is informed by typographical constraints rather than by relative temporal distance of events and should therefore only be taken to provide information as to the sequence of events. Dates are provided by which to mark relative time between events.

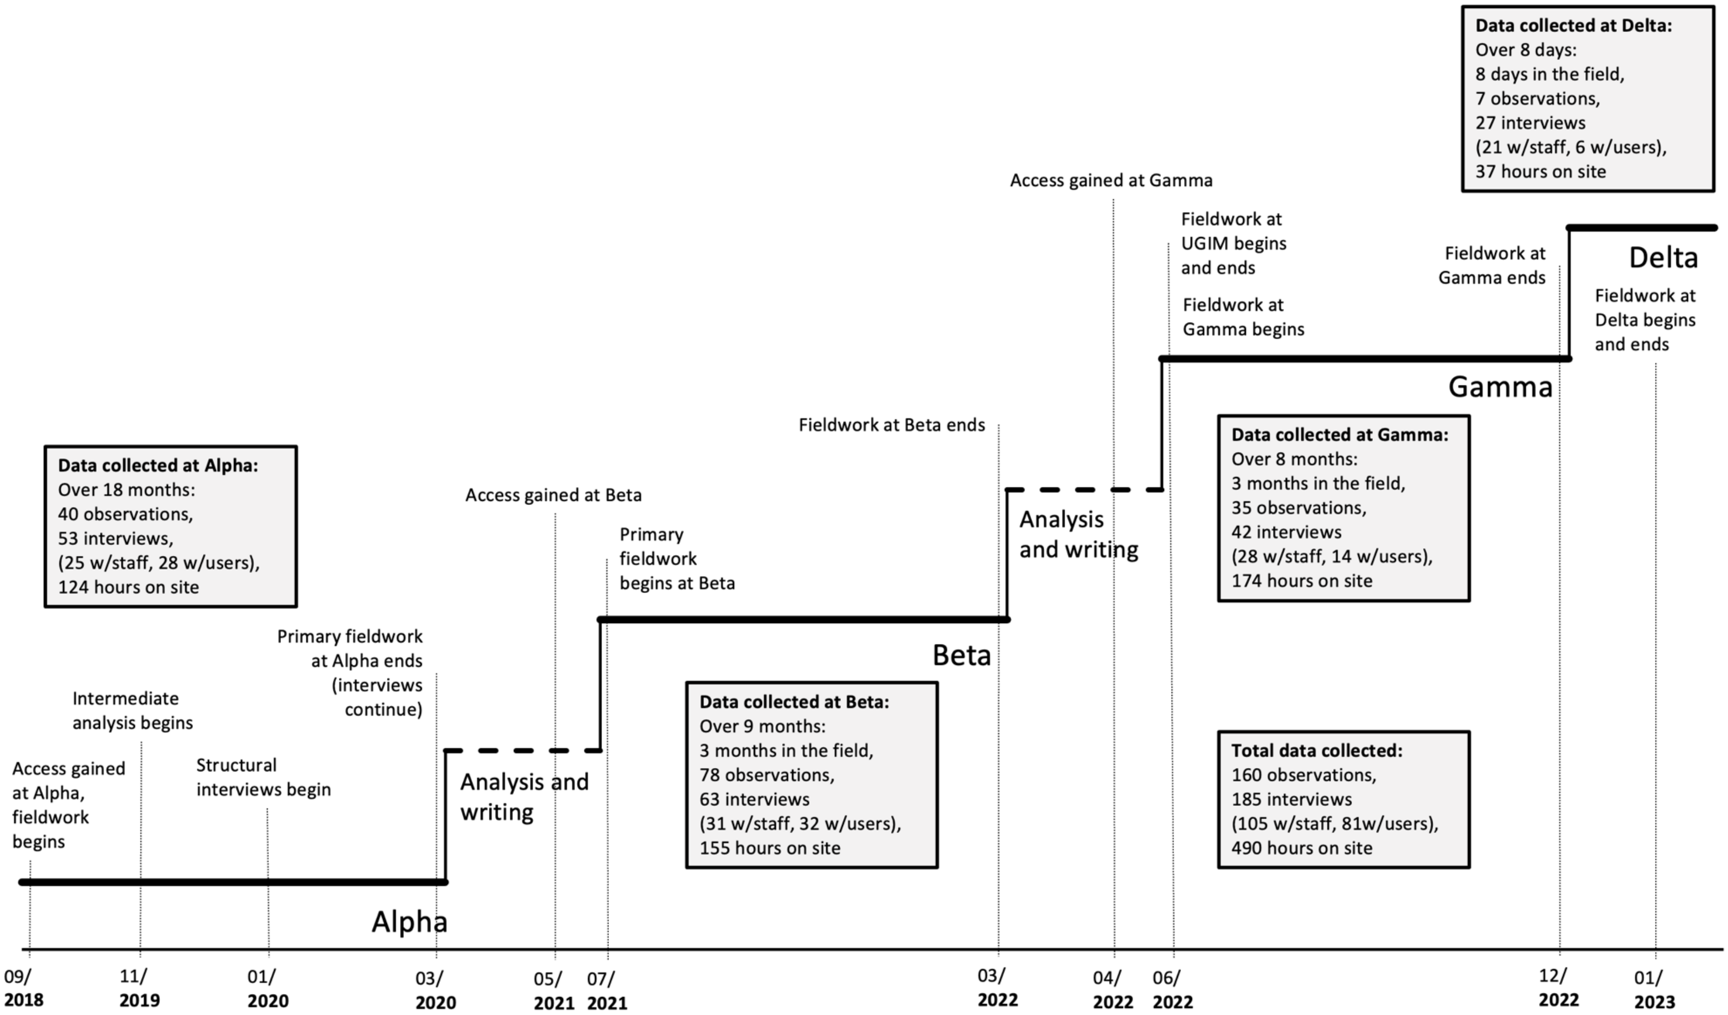

Figure 4. Schematic representation of key elements in process meetings and the relationships between them. Arrows represent actions taken by individuals. Words in boxes represent elements that must be aligned for a fabrication approach to be both suitable given the objectives and constraints of a structure and feasible given the capabilities of tools. Staff members refined their understanding of a project's situated fabrication needs and matched these with the capabilities of specific tools by bringing each of these in relationship to a proposed fabrication approach by way of probes and queries. + and – signify whether a response by user or equipment staff supports or weakens the feasibility or suitability of a proposed approach.

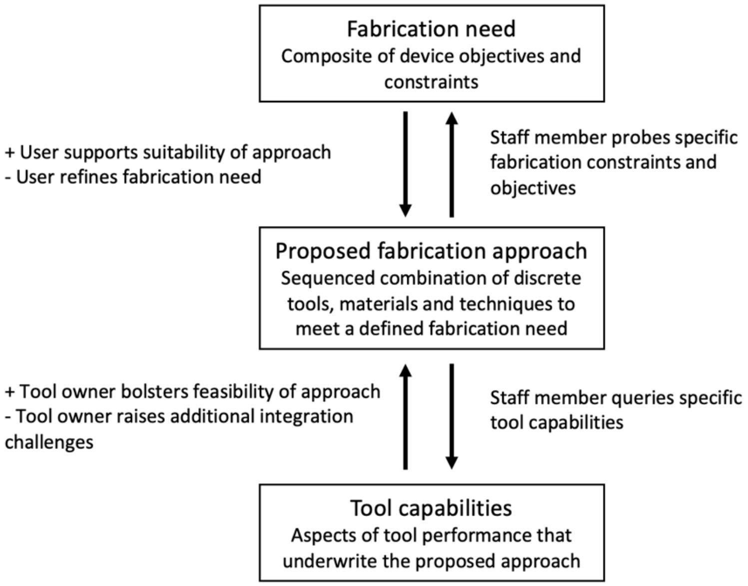

Figure 5. Schematic representation the process by which staff members and users aligned fabrication need, proposed fabrication approaches, and tool capabilities over the course of a given project. Arrows represent actions taken by individuals. Words inside boxes represent elements that must be aligned for a fabrication approach to be both suitable to a given structure and feasible given the capabilities of the facility's tools. The evolution of each element is represented by their changing subscripts. The progression of time is represented by the arrangement of the boxes from left to right. The alignment of the three elements is represented by their vertical proximity to one another. The diagram represents one path such alignment can take. Staff refines their understanding of a project's fabrication need and matches these with the capabilities of specific tools by querying tool owners and users. Their responses lead to revisions of the proposed approach and greater alignment between the three parts of a suitable and feasible fabrication approach. It is also possible that an aligned approach is never reached, due to a mismatch between the constraints and objectives of a given structure and the capabilities of given tools.

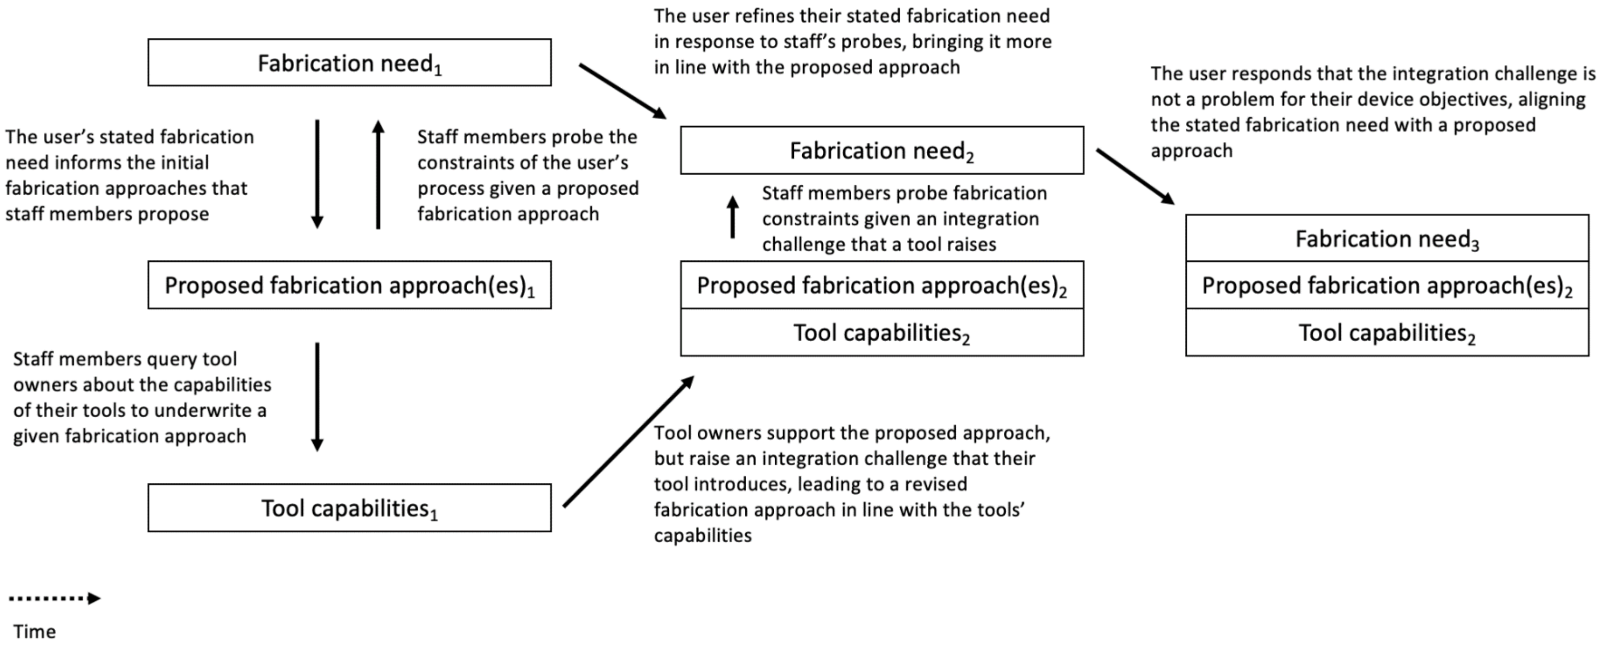

The user's stated fabrication need informs the initial fabrication approaches that staff members propose

Staff members probe the constraints of the user's process given a proposed fabrication approach

Staff members query tool owners about the capabilities of their tools to underwrite a given fabrication approach

The user refines their stated fabrication need in response to staff's probes, bringing it more in line with the proposed approach

Staff members probe fabrication constraints given an integration challenge that a tool raises

Tool owners support the proposed approach, but raise an integration challenge that their tool introduces, leading to a revised fabrication approach in line with the tools' capabilities

The user responds that the integration challenge is not a problem for their device objectives, aligning the stated fabrication need with a proposed approach

Figure 6. A graphical representation of relationships between tool owners' work role, their distinct occupational knowledge, and the outcomes for users and the facilities in which they work. Bolded words indicate the topics of Chapter 5. Unbolded words are the topic of Chapter 6.