VU Research Portal

Justification and Interlaced Knowledge at ATLAS, CERN

Tuertscher, P.R.; Garud, R.; Kumaraswamy, A.

published in

Organization Science
2014

DOI (link to publisher)

[10.1287/orsc.2013.0894](https://doi.org/10.1287/orsc.2013.0894)

document version

Publisher's PDF, also known as Version of record

document license

Article 25fa Dutch Copyright Act

[Link to publication in VU Research Portal](#)

citation for published version (APA)

Tuertscher, P. R., Garud, R., & Kumaraswamy, A. (2014). Justification and Interlaced Knowledge at ATLAS, CERN. *Organization Science*, 25(6), 1579-1608. <https://doi.org/10.1287/orsc.2013.0894>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal ?

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

E-mail address:

vuresearchportal.ub@vu.nl

Organization Science

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

Justification and Interlaced Knowledge at ATLAS, CERN

Philipp Tuertscher, Raghu Garud, Arun Kumaraswamy

To cite this article:

Philipp Tuertscher, Raghu Garud, Arun Kumaraswamy (2014) Justification and Interlaced Knowledge at ATLAS, CERN.
Organization Science 25(6):1579-1608. <https://doi.org/10.1287/orsc.2013.0894>

Full terms and conditions of use: <https://pubsonline.informs.org/Publications/Librarians-Portal/PubsOnLine-Terms-and-Conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2014, INFORMS

Please scroll down for article—it is on subsequent pages

With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

Justification and Interlaced Knowledge at ATLAS, CERN

Philipp Tuertscher

Department of Information, Logistics, and Innovation, VU University Amsterdam, 1081 HV Amsterdam, The Netherlands;
and Institute for Entrepreneurship and Innovation, WU Vienna, 1020 Vienna, Austria, philipp.tuertscher@vu.nl

Raghu Garud

Department of Management and Organization, Smeal College of Business, Pennsylvania State University,
University Park, Pennsylvania 16802, rgarud@psu.edu

Arun Kumaraswamy

Department of Management, College of Business and Public Affairs, West Chester University of Pennsylvania, West Chester, Pennsylvania 19383;
and Department of Strategic Management, Fox School of Business, Temple University, Philadelphia, Pennsylvania 19122,
akumaras@wcupa.edu

We report on a longitudinal study of the emergence of the ATLAS detector, a complex technological system developed at CERN, Geneva. Our data show that the coordination of initial architectural choices was driven by cycles of contestation and justification that resulted in the creation of what we term *interlaced knowledge*—pockets of shared knowledge interwoven within and across subsystem communities at ATLAS. We also found that these justifications were possible because of the presence of a boundary infrastructure that served as a common substrate of knowledge for all ATLAS participants. Together, the boundary infrastructure and interlaced knowledge enabled participants to make co-oriented technological choices, address latent interdependencies, and minimize the incidence and severity of glitches when integrating the various subsystems.

Keywords: collaboration; complexity; coordination; design; innovation; knowledge management; knowledge integration across boundaries; modularity; new product development; technological systems

History: Published online in *Articles in Advance* April 2, 2014.

Introduction

The emergence of complex technological systems is of considerable interest to many scholars of technology and innovation management, and rightly so. The successful development of such systems is key to solving complex scientific puzzles, improving social welfare, and sustaining competitive advantage. In addition, processes underlying the emergence of these technological systems entail experimentation and adjustments that can serve as the basis for organizational transformation.

Although their outcomes can be beneficial, processes by which such complex systems emerge are not straightforward. Consider, for instance, the case of ATLAS, a new particle detector system that was employed recently to discover the Higgs boson.¹ Nearly 3,000 experts hailing from diverse subfields of physics and engineering and working at 175 independent institutions around the world volunteered their time and specialized knowledge to design and deploy this complex system over an 18-year period through scientific collaboration. The system's unprecedented scale and complexity meant that existing knowledge and expertise were insufficient to guide its design and development. Yet choices had to be made among new, competing designs for each component. It was difficult to evaluate the performance of components beyond reasonable doubt. Moreover, the choice of one component affected the choice of others, as well as the ways in which the various components

could be connected. In other words, the system's architecture was not given, thereby precluding coordination of system development through agreed-upon, stable interfaces between various components (what Sanchez and Mahoney 1996 labeled as “embedded coordination”). Furthermore, the collaboration was not governed by a formal hierarchy typically associated with coordinating complex tasks (Zhou 2013). Indeed, the uncertainties and complexities were such that no single “systems integrator” (Brusoni et al. 2001) possessed the knowledge required to mandate choices and direct system development.

How can actors with diverse backgrounds collaborate to develop a complex technological system when its architecture is not given, component technologies are uncertain, and coordination through hierarchy or by a systems integrator is not feasible? To answer this question, we conducted an in-depth longitudinal study of the emergence of the ATLAS detector. Our study led to several findings: First, ATLAS participants could coordinate their activities because of a “boundary infrastructure” (Bowker and Star 1999) comprising, among other assets, texts, tools, and simulation models, which served as a common substrate of knowledge that was transparent and accessible to everyone. All of these together enabled some degree of tacit coordination by making it possible for interdependent participants to interpret and anticipate each other's actions (Puranam et al. 2009,

Srikanth and Puranam 2011). Second, besides enabling such tacit coordination, the boundary infrastructure also allowed ATLAS participants to engage in a process of explication within forums designed to evaluate competing technological options through *cycles of contestation and justification*. Such explication through contestation and justification allowed members of different communities within ATLAS to gain a deeper appreciation of relevant knowledge assets possessed by other interdependent communities. As a result, pockets of shared knowledge emerged across the different ATLAS communities; we call this *interlaced knowledge*. This interlaced knowledge enabled ATLAS participants to *pragmatically decompose* (Simon 1962) the system so that communities could work on different subsystems in a distributed yet parallel fashion. Realizing that any such decomposition could only be provisional given numerous uncertainties and potential latent interdependencies, the different communities continued to communicate and interact with one another, thereby further enhancing interlaced knowledge. This, in turn, enabled them to address latent interdependencies and minimize the incidence and severity of “glitches” (Hoopes and Postrel 1999) when the various subsystems eventually were integrated.

To better place our research in context, we begin by reviewing the literature on the challenges of collaboration during the emergence of complex technological systems. Next, we provide more details on the research site and the specific methods that we used to explore our research question. Then, we describe what happened at ATLAS and highlight our key findings—first the results of our longitudinal study and then the results of additional semantic and scientometric analyses to provisionally validate the findings of our inductive analysis. Finally, the discussion and conclusion sections locate these findings in existing conversations and lay out their implications for relevant literatures.

Literature Review

A rich body of work has emerged that adds depth and nuance to our understanding of the emergence of complex technological systems. For instance, one way to deal with system complexity is to partition the system into components (or modules) that need only interact with one another through agreed-upon and well-specified interfaces (Baldwin and Clark 2000, Garud and Kumaraswamy 1995, Simon 1962, Ulrich 1995). Prior research has noted that such partitioning typically results in a hierarchy of system components, with choices made for components at the apex of the hierarchy influencing choices further down the hierarchy (Clark 1985). Complementing these observations are those offered by scholars who noted that “core” components (i.e., those that are tightly coupled with others) must first stabilize before design parameters for peripheral components

can be set (Tushman and Murmann 1998, Tushman and Rosenkopf 1992).

A parallel stream of research has examined the social dynamics of collaboration shaping the emergence of complex technological systems (e.g., Colfer and Baldwin 2010, Hoopes and Postrel 1999). As Tushman and Rosenkopf (1992) noted, when multiple parties design complex technological systems, decisions are not purely technical in nature but also revolve around the self-interest of involved actors. This is the challenge of cooperation—i.e., aligning interests among interdependent actors (Gulati and Singh 1998, Hoopes and Postrel 1999). Cooperation challenges emerge because different actors find it difficult to balance their personal interests with overall project concerns. Controversies around technological alternatives and the various dimensions of merit that are critical to system performance (Anderson and Tushman 1990, Garud et al. 2002, Tushman and Rosenkopf 1992) can result in excessive haggling and long delays (Farrell and Saloner 1988, Genschel 1997).

Besides cooperation, there are coordination challenges associated with synchronizing the activities of multiple interdependent actors in the design, development, and deployment of a complex system (Camerer and Knez 1996, Hoopes and Postrel 1999, Puranam et al. 2012).² One option is for interdependent actors to conform to an established architecture.³ A second option is for an architecture to be prespecified at the outset by a systems integrator (Brusoni et al. 2001). In either case, a stabilized architecture can allow actors to coordinate their activities, as they need only to build components conforming to agreed-upon interfaces. Sanchez and Mahoney (1996) called this embedded coordination, because each actor knows what must be accomplished without the overt exercise of managerial authority. Relatedly, synchronization of activities occurs by specifying and then adhering to temporal coordination mechanisms such as the use of PERT/CPM (project evaluation and review technique/critical path method) charts (Yakura 2002) or time pacing (Brown and Eisenhardt 1997).

These options are possible once a system architecture has been specified and agreed upon, but they are not available when the architecture, i.e., the components and the connections between them (Henderson and Clark 1990), is emerging. Complicating matters, many potential options for each component and connection may exist, with the choice of any single component impacting options for other interdependent components (Ethiraj and Levinthal 2004). The problem is magnified when the design and development of the system requires diverse and specialized knowledge possessed by different communities of practice (Brown and Duguid 1991, Lave and Wenger 1991). In this situation, no single systems integrator may have the wherewithal to develop an understanding of all system components and their

potential interdependencies, what Postrel (2002) referred to as “trans-specialist” understanding.

How do actors coordinate their activities to develop a system’s architecture under these circumstances? One way to address this challenge is for interdependent actors to develop what Puranam et al. (2012, p. 420) labeled “predictive knowledge,” i.e., knowledge “that enables one agent to act as though he or she can accurately predict another agent’s actions.” Predictive knowledge, then, can serve as a substitute for architectural knowledge (Puranam et al. 2012). Such predictive knowledge becomes all the more possible when “common ground” exists among interdependent actors involved in system development. Common ground has been defined as “knowledge that is shared and known to be shared” (Srikanth and Puranam 2011, p. 850).⁴ As Colfer and Baldwin (2010, p. 20) explained, common ground is required because “contributors need to be able to ‘speak the same language’ in terms of how they explain and interpret designs. They also need to be able to anticipate when, where, and how to look for design information from their counterparts.”⁵

Although predictive knowledge is useful for identifying focal point solutions (Puranam et al. 2012), challenges arise when the underlying technologies are untested. Besides uncertainties about the feasibility and expected performance of these technologies, there are uncertainties about how different groups will come together to support and develop potentially viable choices. In such situations, no “demonstrably correct” solution (Laughlin and Ellis 1986, p. 177) that can serve as a focal point exists at the outset, because latent interdependencies and adjustments (of both the technologies involved and the preferences of the participants) that shape system configuration can occur only as development unfolds (Alexander 1964, Barry and Rerup 2006, Garud and Munir 2008, Pickering 1993). Under these conditions, the design of a complex system becomes a “judgmental task,” i.e., a task “for which there does not exist a demonstrably correct answer” (Laughlin and Ellis 1986, p. 177).

Moreover, when the proposed system is radically new, existing knowledge and expertise may not be sufficient to guide its design and development. So it may not be possible for actors to simply take one another’s expertise for granted and rely on others’ decisions (what Hardwig 1985 called “epistemic dependence”). Indeed, the very notion of what constitutes relevant expertise is contingent and coemergent with the complex system (Garud and Munir 2008).⁶ For instance, early in the process, when specific technologies, components, and the connections (i.e., interfaces) between them have not yet been chosen, it is unclear what specific expertise will be necessary for identifying or resolving conflicts and who will have that required expertise (if at all). Relying on predetermined notions of expertise in such a situation

may lead to suboptimal choices and the premature stabilization of system architecture, potentially generating unproductive interdependencies (Staudenmayer et al. 2005) and a higher incidence of systemwide glitches *ex post* (Hoopes and Postrel 1999).

Under these circumstances, actors belonging to different epistemic communities must jointly and simultaneously accommodate interdependencies across community boundaries to cocreate the architecture (Garud and Karnoe 2003). Such integration of knowledge across “epistemic communities” (Bechky 2003, Knorr-Cetina 1999), which is required to make such accommodations possible, poses nontrivial challenges (Knorr-Cetina 1999). Specifically, it will be difficult for members of a given epistemic community to appreciate or use knowledge developed by other communities without understanding the assumptions and the methods underpinning that knowledge. This difficulty may arise because of semantic and syntactic differences in the language employed by different epistemic communities (Carlile 2002). In addition, it may arise because of pragmatic differences attributable to the differing tools, evaluation criteria, and belief systems of the different communities (Carlile 2002, Garud and Rappa 1994). In fact, attempts to communicate across epistemic boundaries by relying on knowledge transfer alone (i.e., without sufficient understanding of the context in which the knowledge emerged) can generate dysfunctional conflict and divergence rather than cooperation and convergence (Latour 2004, Star and Griesemer 1989). It is for this reason that Bechky (2003) proposed a process of knowledge transformation (as opposed to just transfer) to enable cocreation across epistemic communities (see also Knorr-Cetina 1999).

Given these observations on the critical role of expertise, knowledge, and their integration, we can further refine the broader research question on coordination that we had posed earlier. Specifically, *how can actors from different epistemic communities resolve the knowledge integration challenges that arise in the cocreation of a complex system’s architecture when there is fundamental uncertainty about the system’s components and connections among them?* We address this question through a longitudinal study of the emergence of the ATLAS particle detector at CERN. Our choice of a longitudinal frame is informed by an appreciation that different mechanisms may play a role at different points in time in the emergence of a complex technological system. Even though existing literature has recognized the importance of knowledge integration to reduce the possibility of glitches, how such knowledge integration can be accomplished and how it unfolds over time during a complex system’s design and development remain underexamined.

Research Site and Methods

We had the opportunity to conduct a longitudinal study of the emergence of ATLAS, the large-scale particle

detector located at CERN, that recently was used to detect the Higgs boson, one of the more significant discoveries in physics during the past 50 years. The detector (with a notional cost estimated at one billion dollars) is the culmination of a project begun in 1992 to detect the subatomic particles created by particle beam collisions in the Large Hadron Collider (LHC). By 2007,

construction was complete, and the detector was commissioned for calibration and “data taking.” The focus of this study is on the early design and development phase when the architecture of the ATLAS detector was still emerging.

The detector (see Figure 1), a complex technological system that is 45 meters long and 25 meters in diameter

Figure 1 Overview of the ATLAS Detector

Subsystems	Description
Calorimeter system <ul style="list-style-type: none">• Liquid argon calorimeter• Tile calorimeter	<ul style="list-style-type: none">• The calorimeter consists of metal plates (absorbers) and sensing elements to measure the energy of charged and neutral particles. Interactions in the absorbers transform the incident energy into a “shower” of particles that is detected by the sensing elements.• In the inner sections, the sensing element is liquid argon, i.e., the liquid argon calorimeter. The showers in the argon liberate electrons that are collected and recorded.• In the outer sections, the sensors are tiles of scintillating plastic, i.e., the tile calorimeter. The showers cause the plastic to emit light that is detected and recorded.
Muon spectrometer <ul style="list-style-type: none">• Precision chambers• Trigger chambers	<ul style="list-style-type: none">• The muon spectrometer surrounds the calorimeter and measures muon paths to determine their momentum with high precision.• In precision chambers, gas-filled metal tubes with wires running down their axes are used as sensors. High voltage between the wire and the tube wall allows detection of the traversing muons by the electrical pulses they produce. With careful timing of the pulses, muon positions can be measured to an accuracy of 0.1 mm. The reconstructed muon path determines its momentum and charge.• Trigger chambers are based on a similar principle; however, high time resolution (rather than precision) is the key feature of trigger chambers. Using thin plates or multiple wires as sensors, trigger chambers have a time resolution better than 25 ns.
Inner detector <ul style="list-style-type: none">• Pixel detector• Semiconductor tracker (SCT)• Transition radiation tracker (TRT)	<ul style="list-style-type: none">• The ATLAS inner detector combines high-resolution detectors at the inner radii with continuous tracking elements at the outer radii, all contained in the central solenoid.• The pixel detector generates a set of high-precision measurements that enable the inner detector to find short-lived particles such as B hadrons.• The SCT system is designed to provide a set of precision measurements in the intermediate radial range, contributing to the measurement of momentum, impact parameter, and vertex position.• At larger radii, typically 36 tracking points are provided by the TRT. The TRT is based on the use of straw detectors, which can operate at high rates. They also identify electrons by detecting transition radiation photons created in a radiator between the straws.

Source. ATLAS Experiment © 2013 CERN.

and weighs about 7,000 tons, is housed in an eight-story underground cavern. It consists of subsystems, each of which is a complex system in itself, that identify different types of particles created when proton beams collide, based on energy and momentum measurements.

To design, build, and commission the detector, a tremendous collaborative effort was required over a period of 18 years involving several thousand scientists and engineers with diverse expertise in the areas of physics, engineering, and computer science. Because of the unprecedented size and scope of the project, and uncertainties about the underlying physics, previous generation designs could not be reused for the ATLAS detector. Groups of scientists and engineers, affiliated with 175 independent institutions distributed across 38 different countries (and bound together only by a nonenforceable memorandum of understanding rather than by a traditional hierarchy), collaboratively designed and built the various ATLAS subsystems before shipping them to Geneva for assembly. Eventually, all of these subsystems had to be lowered into a cavern 100 meters underground and installed in the right sequence akin to building a ship in a bottle. Once installed, replacing a subsystem or component was exceedingly difficult and time consuming.

Data Collection

Scholars who studied large-scale projects and organizations similar to ATLAS (Latour 1987, Pickering 1993) suggested focusing on critical events such as technological challenges to understand the core constitutive dynamics of collaborations. Indeed, given the complexities of ATLAS and the uncertainties about the underlying technologies, there were many critical technological challenges pertaining to various components and their connections. For a systematic study of these critical events, we relied primarily on archival data from CERN's archives. These data were generated and archived in real time by participants and represent "unobtrusive measures" (Webb and Weick 1979). We also complemented these 20 years of archival data (1991–2010) with 6 years of contemporary data such as interviews and field visits (2005–2010) (please see Table 1 for a summary).

The archival materials include hundreds of documents such as transparencies presented at meetings, meeting minutes, and internal documents pertaining to central issues and decisions. The materials also include presentations made by scientists to the ATLAS collaboration as a whole and to other scientists at professional gatherings. It should be noted that ATLAS participants routinely documented and archived anything they considered to be relevant to the collaboration's progress. In addition to data from the CERN archives, we gained access to personal notes written by participating scientists and engineers. These personal notes offered insights into the

micro details of system emergence that were not readily apparent in the meeting minutes and other archival sources from CERN. Electronic mailing list archives that captured real-time conversations among ATLAS participants complemented these rich data.

To complement this wealth of archival data and personal notes created by participants in real time, and as a part of a larger study, we carried out 108 extensive in-depth interviews with a cross-section of scientists and engineers from different subsystem communities and groups within the ATLAS collaboration. Following Lincoln and Guba's (1985) guidelines for "purposeful sampling," we interviewed participants who were involved in deliberations around technological challenges. We began by interviewing participants whose names were listed in the meeting minutes and then employed a snowballing technique to identify and interview others. The interview process itself was iterative. For instance, we gathered and analyzed archival material pertaining to events identified by our informants. Then, we sought additional informants based on information found in archival records. We continued this process until we reached data and theoretical saturation (Strauss and Corbin 1998).

The copious data from archival sources and personal notes created by participants in real time already contained many details of the technological challenges associated with the detector architecture's emergence. The interviews we conducted allowed us to gain a deeper understanding of what had transpired. Specifically, they helped us make sense of the technical details, culture, and governance at ATLAS during the period of our study. In addition, they served to validate our identification and interpretations from the archival records of various critical events during the emergence of the detector architecture. Many of our informants had been associated with ATLAS from a very early stage and had no difficulty remembering various technological challenges (such as around the "air core toroid decision" or the "inner detector cooling review," discussed later in this paper). Their accounts were consistent with one another and with the archival records, further increasing our confidence in the reliability and validity of the data.

To get a sense of how ATLAS operates on a day-to-day basis, we attended conferences and participated in meetings during 12 field visits, each typically lasting a week starting in 2006. We also took advantage of chance encounters and engaged in many informal conversations with participants in the ATLAS collaboration. Both during and after these interactions, we also took detailed notes. These field visits gave us a deep appreciation of the culture of collaboration at ATLAS that had emerged and been sustained over many years. The data analysis and results of our study are based on this overall corpus of data and our in-depth understanding of the research site.

Table 1 Sources of Data for This Study

Data source	Description	Use of data
Plenary meeting documents	Summaries of 65 meetings of the ATLAS plenary and supplementary material (presented slides and reports). The ATLAS plenary was the main forum for all-hands discussions concerning physics objectives and results, hardware and software design, and organizational matters. Plenary meetings usually spanned 5 days and were held during ATLAS Week. Our data covers the period 1992–2010.	We used these minutes to identify major events relevant to our study. The “ATLAS Constitution” stipulates that all major decisions must be discussed in the plenary meeting before being voted on by the collaboration board, making the plenary minutes a particularly useful source for generating a chronology of events.
Collaboration board meetings minutes	Minutes of 68 meetings of the collaboration board, the policy- and decision-making body of ATLAS. The minutes extensively list the main issues (including technical as well as social issues) discussed and decisions made between 1992 and 2010.	These data complemented data from plenary meetings with additional details on issues raised in the plenary that were further discussed and eventually voted on by the collaboration board.
Executive board meetings minutes	Minutes of 155 meetings of the executive board. This forum reviewed schedules and milestones and the use of financial and human resources. These minutes cover the years 1994–2010.	These data were used to identify controversies that emerged between various ATLAS subsystems. The executive board had the mandate to bring together coordinators of various subsystem work programs, making it an excellent source for this purpose.
Review panel meetings documents	Minutes of 40 review panel meetings, presentations, and reports discussed during the meetings, as well as the personal notes of a senior scientist who attended all meetings. These forums were set up to choose from alternative designs for the three main detector subsystems and the magnet system. Our data covered all review panels that took place between 1993 and 1994.	These data enabled us to study in detail the dialectical process involved in choosing among competing designs for a specific detector subsystem. Because all review panels took place at about the same time and followed a similar protocol, we were able to conduct a systematic comparison to find similarities and differences between subsystem groups.
Electronic mailing list archives	The electronic mailing list archives containing 128,015 items sent between 1993 and 2005 captured real-time conversations among ATLAS scientists throughout the project.	These data were used to gain additional insights on details of specific incidents and for triangulation. Because these messages were sent in real time when critical events unfolded, these data were not subject to retrospective bias. Moreover, these data offered more details and diverse views of different groups involved than meeting minutes, which sometimes only summarized what all groups involved agreed upon.
Full-text ATLAS Notes	ATLAS Notes documents (2,419 in total) that had been generated in the review panels and task forces across each subsystem. ATLAS documented all of the proceedings of review panels and task forces. Our data covered the years 1992–2005.	We conducted a computer linguistic analysis of the ATLAS Notes. Analyzing this large amount of full-text data allowed us to systematically compare differences in levels of justification of various subsystem groups, thereby probing tentative findings from our qualitative analysis.
Bibliographic records on CERN document server	The CERN Document Server (CDS), an electronic document repository, was used as source of bibliographic data for 2,419 ATLAS Notes. The data covered all ATLAS Notes generated in the years 1992–2005.	Documents stored in the CDS are catalogued according to the MARC standard specified by the Library of Congress, which guarantees a consistent data structure for conducting scientometric analysis, a means of analyzing the knowledge structures of different detector subsystem groups.
ATLAS Project Progress Tracking (PPT) database	The ATLAS PPT database offers detailed data on work packages, milestones, and schedules. It was used as a source for information regarding the progress of the development process within the various subsystems and provided insights on unexpected technical problems and delays. Our data covered the period from the introduction of the database in 1997 until 2007.	We used these data to analyze how the development process unfolded in different detector subsystems and ATLAS overall. Data on unexpected technical problems, delays, and missed milestones allowed us to infer how effectively the various groups were able to coordinate and handle unexpected challenges as the project unfolded.

Table 1 (cont'd)

Data source	Description	Use of data
In-depth interviews	We conducted 84 semistructured interviews with scientists and engineers involved in the ATLAS collaboration at different stages. Some individuals were interviewed on multiple occasions. These interviews typically lasted for 45–90 minutes. In addition, we conducted 24 interviews that began as informal conversations during field visits but, during the conversation, turned out to be valuable sources of data on specific incidents. The content of these interviews was captured in detailed field notes. The interviews were conducted between 2005 and 2010.	These interviews helped us make sense of many technical details and provided additional insights on several controversies and differences in opinions during the emergence of ATLAS. To avoid retrospective bias, these interviews were not used as a primary source of data; rather, they were used to complement our findings based on archival data. We transcribed and coded the interviews using NVivo and cross-checked for accuracy with respondents via email when we found inconsistencies.
Field visits	ATLAS scientists invited us to make field visits to understand and experience firsthand the mechanisms that they have continued to use from their very inception. Accordingly, during 12 field visits, each typically lasting a week, we had the opportunity to observe conferences, participate in meetings, and engage in many informal conversations. These field visits took place between 2005 and 2010.	During our research project, scientists at ATLAS had already begun working on an upgrade of the detector to be installed after 2015. For this reason, we could observe in real time several mechanisms that were similar to the ones unfolding during the early phases of the collaboration.

Data Analysis

Following suggestions made by scholars (Latour 1987, Pickering 1993), we began by examining the technological challenges that participants encountered and how they were addressed. The archival documents we examined revealed that technological challenges had been discussed in several different forums such as working groups, review panels, or plenary meetings. Our analysis of these documents yielded issues such as differences in opinion, conflicts that arose regarding the detector architecture, and the choices that were made. Each issue documented in the meeting minutes usually included referrals to discussions in other meetings and/or documents. To enable us to search for events across our complete database, we digitized archival records using optical character recognition software and indexed our data using a desktop search engine. From this, we were able to identify relevant data not referenced in meeting minutes, such as emails and personal notes of participants. Moreover, the people we interviewed confirmed that the resolution of the technological challenges that the collaboration encountered was critical to the emergence of the detector architecture. Finding such mutually confirming evidence from multiple sources increased our confidence in the quality, depth, and validity of our data and analysis.

We also studied documents such as the ATLAS Letter of Intent, ATLAS Technical Proposal, and Technical Design Reports for individual detector subsystems. These detailed descriptions of technological concepts and design considerations represented snapshots of the detector architecture in the years 1992, 1994, and 1997, providing the development context and enabling us to trace how the detector architecture had emerged over time.

Triangulating across the data sources, we generated a chronology of events as a way to trace and understand the mechanisms that participants used to collaborate over time (see Appendix A for an abbreviated list of events). Consistent with a process perspective, we considered each event as an important occurrence within a larger flow (Van de Ven 1992), an approach that led to a deeper understanding of the unfolding processes. Following recommendations offered by Langley (1999), we also drew diagrams to enable a holistic understanding of the flow of events (see Appendix B for an example).

We coded the data for thematic content (Miles and Huberman 1984) by abstracting raw quotations and text segments from our interviews and from archival materials. In particular, we studied how technological challenges arose and were addressed and resolved as the advocates of different technological solutions confronted one another during meetings and review panels. We performed pattern coding (Miles and Huberman 1984) to identify emergent themes and explanations. Tentative insights emerged as we engaged with the data—e.g., that justification, which enabled coordination, also generated pockets of shared knowledge interwoven within and across different ATLAS subsystem communities.

We decided to further validate these insights by taking advantage of a natural experiment that enabled us to study the pattern unfolding across different ATLAS subsystem communities (Trochim 1989). Specifically, we compared levels of justification and the knowledge structures that emerged over time within and across different ATLAS subsystem communities.⁷ To do so, we used “latent semantic analysis” (Deerwester et al. 1990), a natural language processing technique that makes it possible to extract latent semantic concepts from a large

corpus of texts and to compare the prevalence of specific semantic concepts across different documents. This approach enabled us to explore the prevalence of justifications in the documents produced by ATLAS scientists within each community during the development process. For this, we used a combination of co-word analysis (Callon et al. 1986) and coauthorship analysis (Palla et al. 2007) to identify all knowledge domains relevant for the development of ATLAS (as represented by the ATLAS Notes) and to infer which knowledge assets were held by the participants involved and which knowledge assets they shared and with whom. By doing so, we were able to depict and compare the structure of the distributed yet collective knowledge that emerged over time within and across the communities that were working on different ATLAS subsystems.⁸

On various occasions, we had the opportunity to present our findings and interpretations to ATLAS participants who validated the results of our analysis and our interpretations. They also offered us useful feedback that we incorporated into our analysis. Our interactions with these scientists and engineers increased our confidence in the data we had gathered, the analysis we had conducted, and the inferences we had induced. We share these in the remainder of this paper.

Initial Architectural Choices at ATLAS

The Origins of the ATLAS Collaboration

To appreciate the nature and scope of the challenges that ATLAS participants confronted, it is important to understand the genesis of the ATLAS “collaboration” (as participants called their enterprise). The origins of ATLAS can be traced to two successful CERN High Energy Physics (HEP) experiments (UA1 and UA2) in the early 1980s that led to a Nobel Prize in 1984. Even as these two experiments were unfolding, groups of scientists had initiated a number of independent research and development (R&D) initiatives, with each group exploring a specific technology that potentially could be used for detecting one or more types of particles in future HEP experiments. CERN organized several workshops in which interested groups shared their progress with one another. However, at that time, there was no effort to integrate the various initiatives.

Only when CERN decided to build the LHC did these groups start “protocollaborations” to integrate their disparate efforts into a joint proposal encompassing all the detector subsystems (each consisting of various components and a complex system in itself) required for a new particle detector (Knorr-Cetina 1995).⁹ By 1992, several such protocollaborations had emerged, each interested in developing and building a detector for the LHC. As a result of this decentralized process, there were multiple technological options for each detector subsystem even

within one protocollaboration and many possible configurations of subsystem choices overall.

To rationalize the use of limited resources, CERN sought a merger of protocollaborations. Two groups, EAGLE (Experiment for Accurate Gamma, Lepton and Energy Measurement) and ASCOT (Apparatus with SuperCONducting Toroids), came together to form ATLAS (A Toroidal LHC ApparatuS). The immediate charge to the newly formed collaboration was to create a single new detector architecture. However, *ex ante*, there was no inevitability in the architectural choices for the ATLAS detector system because the collaborators confronted multiple possibilities based on the different approaches and designs that EAGLE and ASCOT brought to the collaboration.

Making design choices and specifying a suitable architecture was challenging for various reasons. The unprecedented levels of energy, radiation, and collision rates required to discover subatomic particles such as the Higgs boson (Stapnes 2007) meant that physicists had to push the frontiers of science and technology. Consequently, prior detector designs, or even subsystems thereof, could not be reused. Indeed, the detector would have to be based partly on technologies that had not yet been developed; scientists could only estimate the performance of these technologies using sophisticated simulation studies or, where feasible, tests on prototypes of individual components. Moreover, a group’s preference for or against a specific approach and their performance estimates relied on assumptions that often were influenced by their experiences with existing technologies and therefore potentially not appropriate for the new approaches being considered. Finally, the sets of assumptions used by different scientists often diverged as a result of differences in their epistemic backgrounds and prior experiences.

These considerations led to ATLAS participants rejecting preexisting notions of who was and who was not an expert on a particular topic and challenging one another to explicate the basis for knowing (in contrast to what Hardwig 1985 called epistemic dependence). Alluding to the driving culture of skepticism, a scientist at CERN noted, “We are a pragmatic community capable to address in a very material way grand and (apparently) immaterial questions, knowing that for every answer we might find, we will open more and unpredicted questions,” and, consequently, “we definitely prefer to be Ministers of Doubt than Kings of Truth: *ubi dubium, ibi libertas*” (Scientist B).

Complicating matters, the two protocollaborations EAGLE and ASCOT had already begun working on different, competing designs. Both protocollaborations seemed to largely agree on the physics goals and wanted to jointly build the detector, but each wanted its own preferred technologies and designs to be adopted. Consequently, many technological decisions such as the choice

and combinations of detector subsystems, triggers, data acquisition procedures, simulation programs, and analysis software had to be made.

Several different configurations were possible depending on specific choices made to accomplish each detector function, giving rise to considerable complexity. Equally important, despite sophisticated simulations of various technological approaches at the component and subsystem levels, the detector itself could be tested in situ only after it had been designed, integrated, and built. In other words, choosing among various technological approaches and configurations at the initial stages involved “judgmental” tasks, not “intellective” ones (Laughlin and Ellis 1986). Not surprisingly, controversies emerged, and resolving them involved a challenge: the collaboration had to move from competing detector designs to a common design by sifting through multiple competing technological options while, at the same time, not alienating groups whose preferred technologies were not chosen. This was by no means an easy task, given the distributed and voluntary nature of the ATLAS collaboration. It is important to remember that the participants in this collaboration all were legally and financially independent physics institutes bound not by legal contracts but by a nonenforceable memorandum of understanding. They all brought in manpower from their home institutions as well as financial resources from their national funding organizations.

In sum, ATLAS was a collaboration among autonomous research institutes coming together “freely” and on an “equal” basis to reach an “understanding” about what they would contribute to the common goal (Knorr-Cetina 1999, p. 164). Consequently, controversies could not be resolved by a centralized decision-making body as in a traditional hierarchy. Not only would it be difficult to enforce such decisions, but such an approach also could alienate groups whose preferred technologies were not chosen and potentially reduce their incentive to cooperate. At an extreme, these groups could even exit the collaboration and, in the process, take with them knowledge, manpower, and financial resources critical to ATLAS’s success.¹⁰

Another important reason why centralized decision making was impractical was that critical knowledge and intellectual resources were distributed among many participants. No individual or group had all the necessary scientific and technical knowledge to make decisions. Indeed, many ATLAS scientists whom we interviewed believed that it was impossible for a centralized group to make appropriate choices and decisions, because each required significant effort and deep understanding of the technological challenges involved.

These two aspects—the importance of consensus among voluntary participants and the fact that knowledge beyond the expertise of any individual or small group of experts was required to design and build the

detector—were reflected in the mechanisms employed for making technology and design choices as well as the governance of the ATLAS collaboration. Before describing the process by which choices were agreed upon, we briefly outline the governance mechanisms that emerged at ATLAS.

Governance at ATLAS

It is not surprising that ATLAS embraced a collaborative form of governance (Adler et al. 2008) that had emerged over the years based on prior high-energy physics experiments and protocollaborations at CERN. Building on these experiences, ATLAS too had adopted a flat, horizontal structure that established links between individual (or groups of) participants with common interests. Activities unfolded within working groups around specific interests and objects such as detector subsystems, measurement of specific particles, or software code. When interdependencies emerged, a variety of groups for coordination were set up as needed. Some of these groups were permanent (e.g., the Executive Board, which reviewed schedules, milestones, and use of financial and human resources), whereas other groups were set up for a limited period of time to handle specific tasks. The organizational structure that emerged was pragmatic and flexible, guided by the interests and tasks relevant to the participants of the collaboration at any given time and without much central control.

Working groups that investigated potential solutions to technical challenges constituted the majority of ATLAS groups. These working groups did not make decisions on behalf of the collaboration, only recommendations. Moreover, these working groups were not supervised by *managers* or *team leaders* but were facilitated by *coordinators* or *conveners* whose mandate was to foster coordination within and across the collaboration’s multiple groups. Similarly, the ATLAS collaboration itself did not have a *president* or corporate-style chief executive officer; instead, the collaboration had an elected *spokesperson*. Having no formal authority over the many participants (initially 414 and increasing to more than 3,000 as of today) who reported back to their respective home institutions, the spokesperson represented and spoke for the collaboration, gently guided attention to topical issues, and served (on request) as an arbitrator by facilitating and steering lengthy consultations (among involved groups and the rest of the participants) toward acceptable solutions.

Given such decentralized governance that eschewed an authority-based hierarchy, how did ATLAS participants coordinate their activities, and where and how were decisions made? The main decision- and policy-making body at ATLAS was the collaboration board, a democratic venue where each participating research institute, regardless of its size and contribution to the collaboration, had one vote. Decisions required a

two-thirds majority vote, placing a strong emphasis on generating a consensus that would engender and sustain cooperation. Even though the collaboration board was the ultimate formal decision- and policy-making body at ATLAS, decisions were usually made after a long deliberative process. This deliberative process involved the ATLAS Plenary, a forum in which all stakeholders and interested participants within the collaboration came together to discuss issues, problems, and progress. Usually, voting in the collaboration board merely formalized the consensus that already had emerged from these discussions. Similarly, although coordinators and the spokesperson were all elected by the collaboration board, such elections occurred only “after nomination of candidates by, and due consultation with, the Collaboration” (according to a 1994 ATLAS internal note, Gen-No-009), further strengthening the collaboration’s emphasis on consensus in decision making.

To facilitate the involvement of all interested participants in seeking solutions and arriving at acceptable decisions, the collaboration sought to make its internal workings and deliberations transparent to everyone. For instance, the studies and recommendations that emerged from various working groups were continually exhibited and disseminated to other groups and the ATLAS collaboration at large in a “grid of discourse” (Knorr-Cetina 1999). Besides the many informal occasions for discourse, there were prearranged events for formal discourse such as the ATLAS collaboration weeks. During the collaboration weeks, meetings of many detector subsystem groups and other working groups were held simultaneously, interspersed with meetings of the ATLAS Plenary for all-hands discussions concerning physics objectives and study conclusions, hardware and software design, and organizational matters. The various working groups presented the outcomes of their meetings and all matters relevant to more than one detector subsystem to the ATLAS Plenary for discussion and scrutiny by everyone who was involved or simply interested. Finally, collaboration board meetings were always scheduled at the very end of each ATLAS collaboration week, so that “the institute representatives [could] participate in all relevant discussions, partake in an emergent consensus, and listen to the ‘mood of the collaboration’” (Knorr-Cetina 1999, p. 184) before voting on important issues that required consideration from many perspectives.

These real-time occasions for discourse were complemented by asynchronous discourse in the form of a variety of documents such as meeting minutes, technical reviews, and design reports as well as extensive communication through electronic mailing lists, databases, and Monte Carlo simulations demonstrating a subsystem’s physics performance that each group made available to the collaboration at large. Moreover, tools for sharing (e.g., databases, Web repositories) and representation

(e.g., common simulation software, conventions for analyses) that partly already existed within the wider HEP community and continued to be augmented as the ATLAS collaboration progressed facilitated access to these assets. Because of such transparent access, any interested participant within ATLAS could contribute to the exploration and solution of any problem faced by the collaboration. Furthermore, all pertinent information to arrive at an informed consensus was available to those who voted as members of the collaboration board. All of these, along with the foundation of physics and engineering knowledge and tools possessed by the participants, constituted a “boundary infrastructure” (Bowker and Star 1999) for everyone involved.

Coordination Through Justification

In their comprehensive review of coordination in emergent organizations, Okhuysen and Bechky (2009) identified five key coordination mechanisms—plans and rules, objects and representations, roles, routines, and proximity. A mapping of these mechanisms with the situation at ATLAS provides additional context to the situation confronted by members of the collaboration during its inception. Plans and rules were ill-defined because there was high uncertainty over how to achieve the objective of discovering new subatomic particles such as the Higgs boson. The emphasis on transparency, discourse, and voting to arrive at a consensus offered ground rules and routines for interaction and decision making. Proximity was an issue, given that participants were geographically distributed. However, the use of communication technologies, comprehensive record keeping, and pervasive access to such records made it possible for participants to stay informed of progress. Undergirding all of these was the boundary infrastructure comprising objects and representations such as simulations that enabled common ground among the geographically distributed participants hailing from different epistemic communities.

Even though these mechanisms offered a foundation for emergent coordination, fundamental issues remained to be addressed in practice. Technological interdependencies had to be considered even though there was uncertainty about what constituted relevant and reliable expertise, i.e., a situation characterized by a lack of epistemic dependence (Hardwig 1985) among the various groups and communities involved (i.e., coordination challenges). In addition, choices had to be made so that groups whose preferred technologies were not chosen would continue to contribute to the collaboration (i.e., representing cooperation challenges).

The collaboration’s approach to dealing with these challenges was to set up review panels—dedicated working groups set up for limited time frames to coordinate the development of major subsystems that eventually could lead to a single detector design. The primary task of these review panels was to recommend to the

collaboration at large “which one of the competing technologies to choose and why” for each subsystem or component thereof (Scientist D). Each review panel included representatives of competing proposals for a particular detector subsystem or component, “neutral” participants working on various other ATLAS subsystems, and, sometimes, independent scientists from outside the collaboration.

Underlying the functioning of these review panels was a core principle that guides ATLAS even today—that competing technological options “can and should be analyzed in a scientific (and not emotional) way” (Scientist F in 1992). The rationale behind this principle was to allow “matters of concern” (Latour 2004) to emerge through a rational enlightened discussion based on “skepticism” (Washington and Cook 2011) by requiring participants to explicate tacitly held beliefs and causal explanations. As one participant explained, this would allow “obvious” decisions to emerge instead of an arbitrary decision being imposed. Not only would this improve the quality of the decisions but also enhance their legitimacy.

Scientist C, who participated in many of these review panels, described the processes that unfolded within the review panel for the calorimeter subsystem:

The panelists were the judges and there were two parties against each other. In our case, for the hadronic part [of the calorimeter], we were against the liquid argon people. We were looking for problems with their approach and they were looking for problems with our approach. The whole thing was relatively formal. We would present our results and our calculations and they would present their results and their calculations. And then, we would ask them nasty questions in writing and they would ask us nasty questions in writing. There would be answers to these questions at the next meeting.

We found many examples of such dialectical processes and intense scrutiny within all the review panels we studied. To facilitate these episodes of contestations and justifications (see also the notion of “trials of strength” described in Latour 1987, p. 78), proponents of each technological option were required to make available all relevant documentation beforehand to allow for informed discussions throughout the review process. Based on the presented material, proponents of a particular technology conducted their own research to develop a deep understanding of competing technologies so as to identify their strengths and shortcomings. These were all articulated in the form of critiques to which each competitor had to respond, justifying their respective approaches by offering evidence from simulation studies and tests of individual component prototypes (where feasible) or by citing support from external experts. These justifications were explicated in the form of presentation handouts and reports to be scrutinized by proponents of competing technologies and review panel members.

As this observation underscores, the cycles of contestation and justification unfolding within review panels resulted in an interpenetration of knowledge between proponents of competing technological approaches for each subsystem. At the same time, the members of the review panels also had to penetrate the black box of technical knowledge to evaluate the arguments being made (see Latour 2004 for importance of critique) to act as informed and engaged judges. So reflexive were these panelists that, in some cases, they even revisited the criteria being used to evaluate competing alternatives. For example, the members of the calorimeter review panel discussed whether or not simulation studies initially selected as the basis for evaluating the technologies would be appropriate. As it turned out, the different groups had based their simulations on many idiosyncratic assumptions, thereby making it difficult to conduct meaningful comparisons between the two competing approaches. The panel members concluded that it would not be prudent to make choices based on these simulations alone, and they decided to delay recommendations until results from prototypes were available.

It is noteworthy that this dialectical process involved a much larger audience than just the competing groups and neutral members within each review panel. Any interested participant—even one not physically present during the reviews—could and did engage. Indeed, many engaged in the process not only to contribute but also because the outcomes of these review panels could impact their own choices. The systemic complexity of the detector was such that choices pertaining to one component in one subsystem could easily impact the choices pertaining to components in another subsystem. Moreover, some components (such as trigger and alignment systems) cut across different subsystems, and certain technologies (such as gaseous detectors) were used in components of different subsystems. In such cases, choices had the potential to impact a number of subsystems that were beyond the purview of a single review panel and necessitated inputs from multiple affected participants and panels.

Recognizing that all relevant participants may not be physically present at these review panel meetings, all interactions that were part of the review process were documented and circulated through electronic mailing lists and made available to all participants in the collaboration on several cross-linked websites. Because of this transparency and engagement, a scientist who was not formally a member of the review panel sent the following email correspondence to the muon spectrometer review panel members in 1993:

I see the following line of reasoning which might have led to this conclusion: (A) For ANY muon trigger scheme, the partitioning in the second coordinate must be of the order of a few centimeters in order for the trigger not to be swamped by background. → The trigger-detector will necessarily measure the second coordinate

with a resolution of a few centimeters. (B) It seems virtually impossible to derive a proper muon trigger from any of the ATLAS precision detector options. → A separate, stand-alone trigger detector is required! (C) Combination of (A) and (B) then seems to lead to the conclusion that the stand-alone trigger device has to measure the second coordinate anyway and hence that no second coordinate is needed from the precision detector.

I fully support the conclusions under (A) and (B). *I do however argue that the apparently very logical conclusion (C) is not well considered and unjustified. It has a number of less favorable consequences which I would like to point out to the panel.* (emphasis added)

This example shows how transparency and easy access to review panel proceedings allowed even individuals who were not formally a part of the review panel to penetrate the technology black boxes to offer valuable new arguments and bring insights to the attention of review panel members. Reflecting on the process, a senior scientist explained the importance of the overall dialectic process and the value of exposing arguments to the scrutiny of both review panel members and the ATLAS collaboration at large:

Very sensible questions were asked, and the more this process rolled on, the more facets were discovered which initially had not even been taken into account. . . . Many things came up which had not been thought of before. This process was scientifically very valuable but also very costly. But if you ask me whether I would endorse such a situation, I do it absolutely. Because the alternative is that you just knock on the table and take a decision. But then you have not thought about everything, not to the same extent to which such a detailed, very painful scientific evaluation process leads you. This is a big advantage. (Scientist D)

The value of such contributions became even more apparent as a specific technology's role in the functioning of the overall detector and the latent interdependencies with other components could only be understood through interactions with those working on other components and subsystems. On these occasions, the coordination process required expertise on interdependent subsystems in addition to knowledge of specific subsystems under review. Scientist O, who was a scientist involved in the calorimeter review panel, explained, for example, how the choice of a specific technology for the inner detector subsystem (recommended by a different review panel) resulted in the use of more material than anticipated by the calorimeter group, thereby potentially introducing errors into the calorimeter measurements. Specifically, particle showers caused by this excess material would result in biased energy measurements in the calorimeter. To account for this error, some calorimeter scientists proposed a new approach to calorimeter design that enabled the detection of particle showers before they entered the calorimeter, thereby

controlling for the measurement bias. Interestingly, this search for latent interdependencies often was driven by some groups seeking to demonstrate that their proposed approach worked better with approaches being considered for other subsystems than competing approaches.

Through this constant unraveling of different technological features, proponents of a particular technological approach explored, to the extent possible, the potential performance and interactions with other detector subsystems. By engaging in these cycles of contestation and justification, they attempted to convince others in the collaboration that their approach was the most appropriate, given the other choices being explored for the ATLAS detector. The long time frames (several review panels explored alternative options for one to two years before making a recommendation) allowed competing groups to mull the steady stream of feedback and eventually come to terms with some alternatives being “more appropriate” than others for the detector design. Over time, after intense scrutiny and deliberations within the review panels and other forums such as the ATLAS Plenary, some technologies became acceptable to a majority, including those who had no intrinsic stake in the technology (Knorr-Cetina 1999, p. 198). Indeed, as a senior scientist noted, “When decision making was an item on the agenda [of the collaboration board], this often meant that something which was already agreed upon and clear for everyone in the collaboration was made plausible and formally approved” (Scientist F).

In sum, the review panels fostered explication through justification and sought to minimize the dysfunctional aspects of politics in the evaluation and decision-making process. This process contrasts with processes found in committee-based decision-making bodies such as standard-setting organizations that often-times involve unproductive haggling and coalition formation (Genschel 1997, Rosenkopf et al. 2001). The constructive confrontation between participants, which were both public and transparent, made it possible to incrementally develop solutions that built on all offered arguments and thus became acceptable to all the groups involved. Indeed, the legitimacy of the choices among competing alternatives was established as these choices were seen to emerge naturally from deliberations.

Equally important, the interpenetration of knowledge between competing groups enabled scientists and engineers who already possessed specialized knowledge of their own technologies to also develop expertise in alternative technologies, both for their own subsystems and interdependent detector subsystems. Such interpenetration of knowledge occurred not only among participants associated directly with the review panels and other working groups but also across the entire collaboration as interested participants from interdependent subsystems got involved in the process. In addition, the outcome of the cycles of contestation and justification was

not merely knowledge transfer. Rather, when competing groups and those working on interdependent subsystems tried to understand the what, how, and why of the various options, they also were able to recontextualize and transform that knowledge to improve or even radically alter their own designs. Such recontextualization and transformation led to the emergence of new options beyond those that were initially proposed for scrutiny by the review panels (see Starbuck 1983 for a notion of design as one that opens up new options).

Emergent Coordination

Over a period of two years, the resolution of early controversies in the ATLAS review panels resulted in choices made among competing alternatives for various subsystems. In late 1994, participants then agreed on preliminary interface specifications among the major detector subsystems, leading to a pragmatic decomposition (Simon 1962) of the detector so that different communities could work separately and in parallel on their respective subsystems.

All participants realized that they were agreeing on these provisional interface specifications for pragmatic reasons and that gaps in their knowledge even at that stage could lead to glitches (Hoopes and Postrel 1999) as their work progressed, which would in turn require adjustments to the specifications.¹¹ Sure enough, a technological challenge emerged around the large superconducting magnet, a component that interacts with several ATLAS subsystems. We describe this in some detail, because it shows how the collaboration functioned. The magnet, initially designed with 12 coils and an inner bore of 10 meters, also set the spatial parameters for the other subsystems. However, over time, the collaboration realized that the costs and risks associated with the deployment of the 12-coil magnet would be far greater than anticipated earlier. Consequently, a Magnet Working Group consisting of magnet experts and participants from different detector subsystems was formed to examine this issue. After thorough deliberation, the group recommended a superconducting magnet with eight coils and an inner bore of only 9.4 meters, which would reduce the costs and risks associated with the magnet and the muon spectrometer subsystem.

However, these new specifications affected other subsystems by leaving only 90% of the originally available space for other detector subsystems inside the magnet. For instance, the space originally kept aside for inner detector electronics was now no longer available. Generating space for these electronics implied compromising the space available for the components of the calorimeter and the muon spectrometer subsystems.

Given the systemic nature of this problem, ATLAS then formed the Global Descoping Task Force to find a solution. All three major subsystem communities were represented in this task force, whose members were

charged “with taking a global perspective rather than acting as a representative for their subsystem” (Ellis et al. 1995, p. 2). Details of the task force meetings were distributed to the collaboration at large to solicit comments, which were distributed to all members of the collaboration. Eventually, after in-depth investigations and intense deliberations, the task force arrived at an acceptable solution. Specifically, gaps would be introduced between the calorimeter and the muon spectrometer subsystems to generate space for additional cables and cooling pipes required for the inner detector subsystem. In the process, the muon spectrometer subsystem was redesigned to give up some space to host the inner detector electronics.

However, this was not the end of the episode. The resolution of the original problem of constrained physical space introduced a new interdependence in terms of electromagnetic fields. The cables installed to connect the electronics with the inner detector subsystem began picking up signal noise due to the magnetic fields generated when they passed close to power supplies. To account for this interdependence, the task force implemented new effective shielding to protect these cables from picking up signal noise. The shielding, in turn, introduced additional materials, generating particle interactions that had to be minimized to reduce disturbance to the calorimeter, creating yet another issue for consideration.

As these observations illustrate, the ATLAS collaboration encountered latent interdependencies that had not been anticipated when the system was pragmatically decomposed. However, each time the collaboration encountered these interdependencies, a fresh cycle of arguments and justification would ensue within concerned forums. To generate solutions that were optimal both locally (at the subsystem level) and globally (at the overall detector level), especially given pragmatic decomposition of the detector architecture, the conveners of these forums made the deliberations as transparent and accessible as those during the early review panels. Accordingly, the different groups presented the new challenges they confronted, the problems that they anticipated, as well as the progress they had made at periodically held plenary meetings. These all-hands conversations played an important role in focusing the attention of all participants and in helping communities working on interdependent subsystems to synchronize their work and track the changes taking place across the ATLAS collaboration.

All of this resulted in further interpenetration of knowledge among communities working on different subsystems, as they had to consider not only the functioning of their own subsystems but also potential interactions with other subsystems and the performance of the overall ATLAS detector. Such interpenetration of knowledge also had another beneficial effect—that of

utilizing knowledge already gained elsewhere within the collaboration to solve emergent problems. As a senior scientist pointed out, at each point in the journey, knowledge generated during the many reviews turned out to be “answers to some other problem encountered later in some other context” (Scientist S).

An example will help illustrate this point. A task force investigating the cooling system for the inner detector subsystem identified potential risks in using binary ice as a coolant.¹² Inner detector engineers had not perceived this risk earlier because they had focused only on the superior cooling performance of binary ice in extracting heat from the densely packed inner detector subsystem. However, because of knowledge overlaps that had emerged, scientists working on other detector subsystems were able to draw attention to the negative impact of binary ice on the performance of their own subsystems. This concern sensitized the task force to the risks associated with binary ice (such as water leakages) and prompted it to propose an evaporative cooling system instead. The resolution of this controversy resulted in a design that not only used less material but also minimized the risk of water leakage within the inner detector subsystem. In other words, because of the interpenetration of knowledge, scientists working on interdependent subsystems were able to utilize the knowledge that had emerged during prior reviews and deliberations to identify robust solutions to unforeseen problems at a later time. Indeed, in some instances, without such interpenetration of knowledge, problems may not have been identified until it was too late.

Further Exploration of Justification and Knowledge Structures

Our inductive analysis of the data revealed that justification driven by contestation resulted in overlaps of knowledge within and across different subsystem communities at ATLAS. To explore this further, we analyzed the role of justification within ATLAS and the knowledge structures that emerged as a consequence. To this end, we took advantage of an opportunity to conduct a natural experiment—i.e., to compare the justification processes within key review panels and associated subsystem communities (a situation of embedded case study; see Yin 1994). Specifically, there was one subsystem community (i.e., the muon spectrometer community) in which the level of scrutiny, justification, and transparency was significantly lower during the early design phase compared to that of others (i.e., the calorimeter and the inner detector communities).¹³ We used this discrepancy to explore whether the different levels of justification influenced the levels of interlaced knowledge that emerged within these communities.

Although the muon spectrometer review panel also held meetings and shared information much like the

other review panels during the early design phase, the nature of the discourse that unfolded during panel proceedings was less reflective and less critical than the others. For instance, whereas other review panels circulated technical details beforehand and allocated most of their time for critical discussions of competing proposals, the muon spectrometer panel merely listened to presentations of competing designs and received them without any systematic questioning or scrutiny. These interactions did not correspond with the norms of justification driven by contestation. Questions, when raised, did not seek explanations of underlying assumptions and rationales. Responses often were taken for granted, and when questions resurfaced in subsequent meetings, questioners were summarily referred to the original responses that could be found in the documentation of earlier meetings (instead of revisiting and scrutinizing the issues carefully). Any scrutiny of consequence came from ATLAS participants who were not formally part of the muon spectrometer review panel. For these reasons, the muon spectrometer review panel was perceived by the collaboration at large (and even by community insiders) to be less transparent than a typical ATLAS review panel (and other working groups). Indeed, after the first few review meetings, the ATLAS spokesperson expressed concern about the lack of scrutiny and discourse within this review panel:

A Panel should not work in isolation but rather in close contact with the collaboration. As most of the people do not know what happens in the Muon Panel they have a right to know Otherwise we will have panic about recommendations being suddenly made and thrown onto the collaboration, without them having a chance of following and appreciating the deliberation process. (ATLAS spokesperson in a 1993 email correspondence)

Despite these concerns, the review panel's process did not change. Indeed, the panel did not even make much effort to justify its final recommendation, even though many thought that important issues had not been explored thoroughly. The groups that had *lost out* (i.e., proposed alternatives to the technology eventually recommended for acceptance by the panel) responded by leading a revolt. One individual involved wrote a formal complaint to the muon spectrometer review panel noting that the “way of justifying—or rather not justifying—was unacceptable” (email correspondence dated 1993). Speaking to legitimacy and cooperation concerns, another influential person within the muon spectrometer community wrote in a 1993 email,

I firmly believe that after many man-years of hard and dedicated work the proponents of all technologies, but especially the losing ones, are entitled to at least one line of comment as to what are, in the eyes of the panel members, their flaws or weaknesses in comparison to the competitors.

As evident from these reactions, it was not the recommendation of or against a specific technology that caused the muon spectrometer review panel to fall short in the eyes of the collaboration. Rather, critics of the review panel were concerned that the lack of scrutiny and transparency made it impossible for other participants of the collaboration to judge from their own perspectives whether the recommendation was based on good scientific reasons and appropriate, both from local (i.e., the specific subsystem) and global (i.e., overall ATLAS) perspectives. Upon further scrutiny after the revolt, it turned out that the technology recommended by the muon spectrometer review panel was marginally superior to competing approaches on paper, but critical problems (e.g., risks associated with production and installation tolerances) were not taken into account even though scientists working on competing proposals were aware of them and attempted to address them in their own proposals.

The protests and revolt eventually led to the rejection of the panel's recommendation¹⁴ followed by significant changes in the processes employed by the muon spectrometer review panel. Additional workshops were organized to discuss the competing proposals and any outstanding arguments. According to review participants and the wider muon spectrometer community, the transparency and intensity with which the muon spectrometer panel engaged in scrutiny and justification increased substantially after the revolt. The eventual outcome was a recommendation that was perceived as being legitimate and, hence, acceptable to all.

Justification and Emergent Knowledge Structures

We conducted additional analysis on the extent to which ATLAS subsystem communities engaged in justification as the project unfolded and also to explore the knowledge structures that emerged as a consequence. To better understand the role that justification played in this process, we compared the muon spectrometer community, which did not sufficiently engage in justification in the early phase, with the calorimeter community, often held up as the role model within ATLAS.

To do so, we first identified the levels of justification during different phases of development for the calorimeter and muon spectrometer subsystem communities using computer linguistic analysis. Second, we analyzed the knowledge structures that emerged for each of these subsystem communities by generating bipartite networks consisting of the individuals involved (several hundred physicists and engineers spread across different research institutes and countries) and their areas of expertise. We then examined the relationship between justification and the knowledge structures over time and between the different subsystem communities by comparing the results of the prior two steps. Finally, we analyzed data on project progress made by these different subsystem communities to explore how the differences in justification

and consequent knowledge structures were related to differences in project progress. Next, we describe these steps and the outcomes of our analysis.

Justification. Our approach to identifying different justification levels among subsystem communities involved conducting a computer linguistic analysis of the ATLAS documents generated in the review panels and other forums pertaining to each subsystem. As we stated earlier, all the proceedings of the various working groups, review panels, and task forces were documented at ATLAS. These texts were important because they (a) provided transparency to the process, (b) could easily circulate among participants of the collaboration so that even those who were not present at the meetings could get involved or scrutinize the rationale for a decision, and (c) generated a memory of how and why a certain recommendation was made. Thus, analyzing these considerable data offered a way for us to probe the levels of justification that had unfolded within these forums.

Altogether, we had access to 2,419 documents generated by ATLAS scientists between 1992 and 2007. Each document was time-stamped and attributable to specific ATLAS subsystems, enabling us to explore differences in justification levels over time within each subsystem community. As a proxy, we identified terms such as “because,” “since,” “therefore,” and “due to” in the text as indicators of justifications offered by authors. We used latent semantic analysis (Deerwester et al. 1990), a natural language processing approach, to measure the importance of these words in the documents pertaining to each subsystem community. An advantage of latent semantic analysis is its ability to overcome the problems that emerge because of synonymy and polysemy,¹⁵ a feature that made it possible for us to capture the semantic network of concepts around justification within a document rather than conduct a simple word count (Landauer et al. 1998).

Figure 2 shows the justification factors for the calorimeter and muon spectrometer communities during different phases of development. The justification factor for a specific subsystem community was the mean value of cosine similarity (Deerwester et al. 1990) of the search vector for justification-related terms (e.g., “because,” “since,” “therefore,” “due to”) and all documents generated by that specific subsystem community, indicating the extent to which that community engaged in justification. As evident from Figure 2, the level of justification was high within the calorimeter subsystem community during the early design phase, and it even increased marginally during detailed development of the subsystem. By contrast, the level of justification in the muon spectrometer community was low initially but increased after the revolt and reached its highest level during detailed development between 1996 and

Figure 2 Justification Associated with Two ATLAS Subsystem Communities

1998. Overall, the pattern in this figure is consistent with our earlier qualitative description of the justification processes as perceived within these two subsystem communities.

Knowledge Structures. We also analyzed the bibliographic records of the 2,419 documents to understand and depict the knowledge structures that had emerged within the calorimeter and muon spectrometer communities. Specifically, we identified each participant's areas of expertise by analyzing the topics appearing in the titles, keywords, and abstracts of the documents he or she had coauthored. We used latent semantic analysis to account for terms that were used synonymously to depict an area of expertise. Based on these bibliographic records, we generated a bipartite network (see Appendix C) of all participants belonging to a specific subsystem community (several hundred people in each case) and their respective areas of expertise relevant to any ATLAS technology, whether related to their own subsystem or to the others.

From these data, we created a unipartite projection of the knowledge structures of the different subsystem communities and analyzed knowledge across ATLAS subsystems using density¹⁶ as a network measure (see Figure 3). Density represents the ratio of connections between all nodes of a network to the number of connections that are theoretically possible (Mitchell 1969). A subsystem community with a knowledge structure connecting many diverse areas of expertise, both within and across subsystems, will have a higher density than a community with knowledge structures connecting fewer areas of expertise (e.g., focused on their own subsystem).

A visual comparison of the knowledge structures during the early design phase in Figure 3 (left panel) shows that the knowledge structure held by members of the calorimeter subsystem community had a higher

Figure 3 Visualization of Interlaced Knowledge of Two ATLAS Subsystem Communities

Notes. The nodes in the networks represent the 456 specialized domains of knowledge involved in the development of the ATLAS detector. The color code indicates the centrality of a particular knowledge domain in a subsystem community (red indicating highest centrality). The edges connecting the nodes indicate overlaps between knowledge domains, i.e., ATLAS scientists working in that subsystem community having specialized knowledge in two or more domains (yellow edges indicating weak and green edges indicating strong connections). As this figure shows, during the early design period (around 1994), the calorimeter community had a more interlaced knowledge structure than the muon spectrometer community, connecting knowledge assets relevant to its own subsystem as well as knowledge assets relevant to others. This difference between the two communities is less apparent toward the end of the development of the ATLAS detector (around 1998).

density than the knowledge structure held by the muon spectrometer subsystem community. Figure 3 also shows that the density of the knowledge structures increased for both communities (right panel). This same pattern is observed in Figure 4, a plot showing the density of the knowledge structures of the two subsystem communities during the early design and the later development phases. In other words, the number of connections between diverse areas of expertise increased, as did connection strength.

A juxtaposition of Figure 2 with Figure 4 suggests a relationship between justification and knowledge structures. For each subsystem, the density of the knowledge structure (Figure 4) parallels the levels of justification (Figure 2). Moreover, for both subsystem communities, knowledge structure density is commensurate with

Figure 4 Density of Interlaced Knowledge in Two ATLAS Subsystem Communities

the respective levels of justification (or lack thereof). Although we cannot ascribe causality from such a comparison, our qualitative and contextualized analysis suggests such an inference. Specifically, as we described earlier, the competing choices for the calorimeter subsystem were heavily scrutinized and debated, frequently involving contributions from members of other subsystem communities. The final recommendation was made after a high level of justification during the early design phase, thereby generating for the calorimeter community a knowledge structure that densely connected knowledge of calorimeter technology with knowledge of interdependent subsystems.

In contrast, the muon spectrometer community did not carry out such intense scrutiny during the early design phase, and justification played only a minor role in the process until after the revolt. Consistent with this context, the knowledge structure that emerged within that community was sparse and more focused on the technology associated with its own subsystem compared to the calorimeter community. However, by the later development phase, the muon spectrometer community also had developed a justification level similar to that of the calorimeter community, and its knowledge structure also became densely connected, including knowledge of other subsystems.

From this analysis, we inferred that justifications (disseminated widely and transparently) had played an important role in the emergence of the ATLAS architecture by generating a knowledge structure with essential points of overlap or interpenetration within and across communities working on different subsystems. Furthermore, as we had qualitatively discussed earlier, not only was the knowledge structure the outcome of the justification processes involved, but it also served as a medium for coordination, enabling interdependent groups to identify emerging problems and develop

solutions as the project unfolded. From this perspective, it was not surprising to see a contrast between the communities working on the calorimeter and the muon spectrometer subsystems in terms of technical problems and project delays.

Indeed, further analysis of the ATLAS Project Progress Tracking database showed that the calorimeter subsystem community had performed the task of developing and constructing their subsystem more effectively. Whereas 9% of all design-related tasks of the calorimeter community were delayed because of unexpected technical problems, the muon spectrometer community was challenged by unanticipated problems in 14% of all design-related tasks. This difference was even more striking when we compared the mean delay of design tasks undertaken by either subsystem community. The average delay of muon spectrometer-related design tasks (6.7 months) was more than double the average delay of calorimeter-related design tasks (2.7 months). In other words, the differences in knowledge structure appear to be reflected in the incidence and severity of glitches (Hoopes and Postrel 1999) encountered by the two communities as their respective work progressed.

Discussion

We began this paper by asking how actors with diverse epistemic backgrounds were able to collaborate to design a radically new and complex technological system when there were significant uncertainties about underlying technologies and a taken-for-granted architecture did not exist that enabled embedded coordination. As we reported, there was insufficient predictive knowledge (Puranam et al. 2012) to facilitate coordination, given the novel technologies under consideration and the resultant uncertainties. In addition, epistemic differences across the subsystem communities (Knorr-Cetina 1999) made it difficult for the collaboration at large to evaluate and agree on choices, especially with different groups offering what they thought were the best options for any specific part of the system. Indeed, it was not possible to demonstrate conclusively (through simulations, experiments, or prototypes) whether the system would work well when components and subsystems were integrated as proposed. In other words, the possibility of glitches (Hoopes and Postrel 1999) was high with potentially disastrous consequences in terms of additional effort and funds required to fix them, the attendant time delays, and a tarnished reputation of the scientific enterprise.¹⁷

There did exist a boundary infrastructure offering a common ground (Srikanth and Puranam 2011) covering foundations of high energy physics and collaboration in large experiments (e.g., importance of consensus, transparency, voting rules). However, this alone was not sufficient to address the challenges that the new system

in the making posed, because there had never been a physics collaboration of comparable scale and technological complexity before.

For all of these reasons, the collaborative community confronted a *judgmental task* in cocreating a common architecture. Accordingly, the collaborators had to go beyond the typical coordination mechanisms proposed for multiparty development of complex technological systems, such as technical committees, task forces, or standards bodies (Farrell and Saloner 1988, Hobday 1998, Rosenkopf and Tushman 1998). In particular, knowledge integration issues were a challenge, given the diverse epistemic backgrounds of interdependent groups. Especially troublesome was the inability to rely on existing knowledge and expertise (what Hardwig 1985 called epistemic dependence), given that the contours of the architecture, and therefore the nature of problems to be addressed, were not known. Participants realized the potential for unproductive conflict that could adversely impact incentives to cooperate if decision making was politicized. Specifically, groups that had volunteered their expertise and technologies to the collaboration could simply decide to withdraw the resources they had brought with them and exit the collaboration, or they could stay in the collaboration but not contribute their best efforts if their offerings were summarily ruled out without careful discussion and evaluation.

So how did participants at ATLAS address these challenges, especially those pertaining to knowledge integration? Our study highlights (a) the importance of justification that led to (b) the emergence of an interlaced knowledge structure that reduced the potential for and severity of glitches, facilitated emergent coordination, and enabled (c) pragmatic decomposition of the complex system. We explicate these insights below.

Justification

Our study of ATLAS revealed justification as an underexplored yet key coordination mechanism when interdependent groups work on the architectural choices and the design of complex systems. We identified this mechanism not as a substitute but as a complement to coordination mechanisms identified in extant literature. For instance, ATLAS had ongoing interactions among participants within task forces, committees, collaboration board meetings, and ATLAS Weeks. Discussions within these forums covered a wide range of topics including the progress that had been made and the bottlenecks that actors had encountered or anticipated. Indeed, these forums can be likened to committees and other collective forums (see also Rosenkopf et al. 2001, Simcoe 2012) that offer interdependent actors opportunities to coordinate their activities by sharing knowledge on who is doing what and when (i.e., schedules and choices). In addition, these forums also served as venues for the

exchange of general domain knowledge such as results of experiments and simulations.

In an earlier study of a high-energy physics community at CERN, Knorr-Cetina (1999) also documented the presence of such forums and noted that they foster a “grid of discourse.” As she noted, “Discourse channels individual knowledge into the experiment, providing it with a sort of distributed cognition or a stream of (collective) self-knowledge, which flows from the astonishingly intricate webs of communication pathways” (Knorr-Cetina 1999, p. 173). At ATLAS, the grid of discourse was such that it was completely transparent to all participants. Such transparent access to these discourses generated a situation similar to what has been called “actionable transparency” (Colfer and Baldwin 2010, Zuboff 1988).

The boundary infrastructure that emerged as a result of this grid of discourse enabled participants to arrive at coordinated inferences (Bowker and Star 1999) as they “co-oriented” (Taylor and Van Every 2000) themselves while negotiating the design. Boundary objects (Star and Griesemer 1989) are key elements of this infrastructure. At ATLAS, these boundary objects included templates for presentations, conventions for presenting physics results using plots, elaborate frameworks for producing Monte Carlo simulations as well as the results created by such simulations, and PERT/CPM charts reporting progress. These may appear mundane, but such standardization helped actors coordinate their activities efficiently and effectively. For example, the Monte Carlo simulations of physics results made it possible for participants to grasp the potential performance of individual components and subsystems and thereby coordinate their activities. By using and exchanging such boundary objects, scientists from different backgrounds enacted a “trading zone,” which allowed them to coordinate tacitly across disciplinary boundaries even if they disagreed on the meaning and the significance of these objects (Galison 1997, Kellogg et al. 2006).

What we found interesting, however, was that participants did not use this boundary infrastructure to engender only tacit coordination (Srikanth and Puranam 2011). Instead, participants also used it as a fabric to voice concerns or contest one another’s claims so that underlying tacit assumptions and beliefs were explicated in the process of making co-oriented choices. These observations are consistent with the findings of science and technology study scholars such as Callon (1995, p. 55), who noted that a whole “invisible infrastructure” is required to spawn meaningful discussion and debate among interdependent actors. Specifically, Callon (1995, p. 55) observed that even the “all controversies, even the most fierce and relentless, depend upon a tacit agreement about what is important and what is not.” In other words, as we observed at ATLAS, the boundary infrastructure can also serve as the catalyst for explication through cycles of contestation and justification.

This boundary infrastructure allowed competing groups of participants to discuss and debate differences among alternatives within review panels. At these panels, many different points of view and subsystem-specific perspectives were articulated within an overall process characterized by others as “contests of unfolding” (Knorr-Cetina 1999, p. 196) and “trials of strength” (Latour 1987, p. 78). To engage in such contests, groups had to generate a deep understanding of their own and others’ technologies (i.e., open up the modularity black boxes) before they could scrutinize competing proposals for their own as well as for interdependent subsystems. In turn, when challenged, they had to present arguments to justify their own proposals (Green 2004, Toulmin 1983). Indeed, even the perspectives, questions, and justifications (whether in support of or against) of groups working on other subsystems than the one under review were considered. Because of these dialectical processes, knowledge was not just exchanged or transferred. As participants explicated their beliefs and underlying assumptions to justify their respective preferences, they and others were able to create new knowledge (Nonaka and von Krogh 2009). In other words, a new synthesis emerged when issues were debated from different vantage points.

These processes led to co-oriented decisions that wound their way through the ATLAS organization to be ratified eventually by the collaboration board. The whole process is analogous to collective induction, “an orderly social combination process of resolution of disagreements by voting, turn taking, demonstration, and formation of emergent hypotheses, rather than any of these four processes alone” (Laughlin and Hollingshead 1995, p. 94). In addition to coordinating technical decisions across interdependent subsystem communities, this process ensured that groups would not lose their motivation to continue cooperating, even if their respective technologies were not chosen. The outcome was a detector architecture that had both technical integrity and legitimacy.

The Emergence of Interlaced Knowledge

As we noted, the dialectical process of inquiry at ATLAS not only served as a mechanism for engaged coordination but also enhanced the boundary infrastructure that all participants could access. At the same time, explication through justification also enhanced the knowledge of the groups that confronted one another within review panels and, as a result, deepened specialized pockets of expertise on each component and subsystem within ATLAS. As panel members and other interested participants working on interdependent subsystems also scrutinized the competing proposals, members of the ATLAS community not only deepened an understanding of their own subsystems but also gained a deeper understanding of other interdependent subsystems. Such

interpenetration of knowledge across communities was not centrally orchestrated but occurred in an emergent fashion driven by members’ appreciation of critical interdependencies based on their own experiences and the interactional expertise (Collins and Evans 2007) that emerged. As many participants interacted with one another in a variety of different forums, knowledge overlaps occurred across different dyads in an idiosyncratic yet redundant way.

Figure 5 offers a stylized representation of the knowledge structure that emerged. In Figure 5, the circle at the base represents the boundary infrastructure on which ATLAS operates. Built on that common substrate are intersecting circles representing the knowledge domains of members of three different communities, A, B, and C. Each community has multiple members (A1, A2, etc.) who possess specialized knowledge that is comprehensible to other members of that community. A2 and B1 may share some knowledge asset, K6, such as a particular sensor that can be used by both communities. B2 and C2 have their own pocket of shared knowledge, K11. But K11 is not shared with members of Community A, and K6 is not shared with members of Community C. The various intersections of the knowledge domains represent interlacing of knowledge within and across communities.

We use the term *interlaced knowledge* to refer to the knowledge structure represented by the distributed overlaps (K6, K7, K9, K11, etc.) considered together. In its structure and distributed nature, interlaced knowledge moves the level of analysis from dyads (Hoopes and Postrel 1999, Postrel 2002, Srikanth and Puranam 2011) to the level of the entire system; that is, no one knows everything, but the whole is constituted from the parts because of the interwoven pockets of shared knowledge among participants.

As interested participants who are concerned about specific aspects of another subsystem open up the black box through contestation and justification, they create targeted and redundant partial overlaps of knowledge across subsystem communities. At the same time, each participant’s expertise remains intact (and perhaps is even deepened because of the cycles of contestation and justification) so that the specialization required to make appropriate design choices for specific detector subsystems and the components thereof is maintained.

Such productive redundancy enables what Nonaka and Takeuchi (1995, p. 14) labeled a “shared division of labor” (see also Garud and Kotha 1994, Morgan 1986). In this sense, it is different from the notion of trans-specialist understanding discussed by Postrel (2002), wherein actors become generalists incapable of solving intricate problems requiring specialized knowledge, but play a role in integrating knowledge across different specialist domains. In ATLAS, the integration was accomplished by the specialists themselves, each crossing a

Figure 5 Interlaced Knowledge and Boundary Infrastructure

Notes. The bipartite graph in panel (a) shows scientists from different subsystem communities who have expertise in various knowledge domains. Moreover, these scientists all share a common boundary infrastructure. The circle at the base of panel (b) represents the boundary infrastructure on which ATLAS operates. Built on that common substrate are intersecting circles representing the knowledge domains of members of three different communities A, B, and C. Each community has multiple members (A1, A2, etc.) who possess specialized knowledge that is easily comprehensible to other members of that community. A2 and B1 may share some knowledge asset (such as a particular sensor that can be used by both communities), K6. B2 and C2 have their own pocket of shared knowledge, K11. But K11 is not shared with members of community A, and K6 is not shared with members of community C. The various intersections of the knowledge domains represent interlacing of knowledge within and across communities.

few epistemic boundaries in a multiplexed way. Because not everyone need generate a deep understanding of all parts of the detector, it rationalizes the costs of crossing epistemic boundaries (for more on escalating costs of knowledge integration across multiple epistemic communities, see Berends et al. 2011, Grant 1996a). And the network of connections across the people and artifacts across the entire collaboration creates a pathway for access to diverse expertise on demand. All of this is supported by the common substrate of knowledge, i.e., the boundary infrastructure that constitutes the lingua franca of the emergent detector system.

Dynamics of Coordination Over Time

Figure 6 summarizes our process model. As we discussed earlier, coordination was accomplished through several mechanisms, especially cycles of contestation and justification that led to the emergence of interlaced knowledge. Interlacing of knowledge, in turn, enabled the pragmatic decomposition of the ATLAS detector

so that parallel and distributed development of various subsystems could occur. Participants, however, realized that any such decomposition could only be provisional and would need to be revisited as development proceeded. Unlike systems where some components are more important than others and occupy the apex of the system hierarchy, in the ATLAS case, many components were equally important. Consequently, it was not possible for the collaborators to wait for components at the apex of the hierarchy to crystallize before setting the parameters for the emergence of components lower in the system hierarchy (Clark 1985, Tushman and Murmann 1998, Tushman and Rosenkopf 1992). Instead, given the uncertainties and complexities, latent interdependencies only surfaced as system development and integration unfolded (Barry and Rerup 2006, Garud and Munir 2008, Staudenmayer et al. 2005).

Realizing the possibility of latent interdependencies emerging with further development, participants at ATLAS continued to maintain rich communications

Figure 6 Dynamics of Coordination Over Time

and interactions within and across subsystem boundaries even after they had pragmatically decomposed the system. This observation lies in contrast to suggestions in the extant literature on modularity that there is limited need for communication across groups working in parallel once a system has been decomposed (Sanchez and Mahoney 1996). However, in the ATLAS case, such communication and interaction made it possible for participants to proactively address unexpected issues that emerged within and across subsystem boundaries as critical dimensions of the system (e.g., mechanical structure, electromagnetic fields, thermal regions) changed over time and as new dimensions gained importance as the project evolved. As these challenges were resolved, the detector architecture itself evolved through adjustments made to facilitate the optimal functioning of the detector as a whole (see also Pickering 1993).

In sum, emergent coordination was made possible by maintaining interlaced knowledge even after pragmatic decomposition. Because the interlaced knowledge structures did not “mirror” the way the technological system was pragmatically decomposed (thereby falling into the “exceptions” category as per Colfer and Baldwin 2010), designers of different subsystems could appreciate the evolving system architecture from different yet co-oriented perspectives (Taylor and Van Every 2000). As a result, problems in one subsystem were sometimes identified by individuals who were not directly involved, as the example of the inner detector cooling task force illustrates. Also, solutions to emergent problems were often found when members of other subsystem communities helped by approaching the problem from a different perspective. In other words, rather than mirroring, the relationship between the technical architecture and knowledge structure generated “interpretative flexibility” (Pinch and Bijker 1987) as participants sought creative solutions to emergent problems.

If, instead, the knowledge structure had been modular, or even “islands of shared knowledge in a sea of mutual ignorance,” as characterized by Postrel (2002, p. 304), the development of the system could have been locked prematurely into a prespecified path, leading to costly glitches and a system with low integrity. Or a change in specifications of one subsystem could easily have triggered changes in other interdependent subsystems causing system development to cycle endlessly (Ethiraj and Levinthal 2004). The presence of interlaced knowledge reduced such possibilities (for a related discussion on the development of complex aircraft systems, see Sosa et al. 2004). As unforeseen events occurred, such as changes to the very form and function of one subsystem, participants whose knowledge was interlaced were able to proactively identify emergent problems caused by latent interdependencies with their own subsystems as well as with other subsystems. Equally important, they were able to find mutually accommodative workarounds for these emergent problems. In this sense, interlaced knowledge represented productive redundancy, allowing subsystem members to mindfully respond (Weick and Roberts 1993, Weick and Sutcliffe 2001) to unexpected situations by developing novel solutions.

All of this offers a perspective on the emergence and the role of architectures in various phases of the development of complex systems that extends the findings of prior studies. Specifically, we found that creating an architecture was important for ATLAS to enable interdependent participants to coordinate. However, in the ATLAS case, this was not a prespecified architecture in its traditional role of providing embedded coordination (Sanchez and Mahoney 1996). Instead, the provisional architecture that emerged because of pragmatic decomposition facilitated coordination because it became part of the boundary infrastructure that flexibly connected heterogeneous components and groups

(Bowker and Star 1999). The continued interaction and engaged discourse across interdependent groups and the collaboration as a whole maintained and deepened interlaced knowledge, thereby allowing participants to proactively address problems as the development of the system progressed. As these problems were resolved, the architecture itself evolved as adjustments were made to accommodate the resistances and affordances of the various components to facilitate the optimal functioning of the system as a whole.

Implications

What can we learn from the ATLAS case? One way to answer this question is to revisit the boundary conditions that characterized ATLAS, thereby locating our observations within a larger theoretical mosaic of ideas on coordination, knowledge integration, and the emergence of technological systems. ATLAS is a case of a complex system whose architecture could not be easily prespecified and decomposed. Coordination challenges arose because of the extreme uncertainty associated with the technology and interdependencies among its components and subsystems. Further complicating matters, multiple groups were involved, each with their own epistemic background and preferred technological approaches to developing each component or subsystem. No one actor had either the authority or the knowledge to simply impose a technical architecture that could then engender coordination. Equally important, it was not possible to first create a prototype of the system and then refine it progressively. All of these issues, in combination, rendered this a judgmental task.

To address this judgmental task, our study identified a variety of complementary mechanisms to address the coordination problem that the participants confronted during the emergence of a radically new, complex system. As different challenges arose at various phases of development, participants deployed different mechanisms to address the specific challenges they confronted. The dynamics of how coordination challenges emerged and were addressed over time is a contribution to the literature.

The cycles of contestation and justification were critical in this regard. Essentially designed to explicate the rationale for the choice of a particular technology, coordination based on contestation and justification differs from coordination based on knowledge that is shared and known to be shared (or common ground, as in see Srikanth and Puranam 2011) and is thus taken for granted. Whereas the latter relies on prior interactions and familiarity to engender tacit coordination, coordination through cycles of contestation and justification draws upon explication of underlying assumptions, logic, and details.

Our study also reveals how the cycles of contestation and justification generated interlaced knowledge. Prior

literature has addressed shared knowledge at the dyadic rather than systemic level. By limiting attention to dyads while ignoring the larger structural context within which these dyads operate, prior studies have not adequately explored how the overall network structure influences mechanisms at the dyadic level (Granovetter 1992, Jones et al. 1997). An innovative methodology (latent semantic and network analysis) allowed us to examine knowledge overlaps not only across dyads but also at the network level. By measuring shared knowledge of all individuals within and across collaborating groups, we were able to highlight the extent of knowledge overlaps as well as the nature of sharedness. Indeed, our results suggest that the interlacing of knowledge created an interconnected tapestry of targeted knowledge overlaps between individuals belonging to different communities, as opposed to Postrel's (2002) "islands of shared knowledge in a sea of mutual ignorance." These productive redundancies in the knowledge structure increased what Kotha and Srikanth (2013) call "visibility," i.e., the situational awareness of all activities throughout the network that enables participants to fully comprehend the issues confronting them and also to figure out where they can obtain the relevant expertise to diagnose and resolve problems.

These findings have implications for addressing a managerial challenge in coordinating complex technological systems identified by Postrel (2002)—to decide what knowledge needs to be shared by whom and to what extent. In ATLAS, participants working on various detector subsystems self-selected to participate in the dialectical processes and the creation of interlaced knowledge based on their own interests. Even participants who were not directly involved in either proposing or evaluating alternative approaches for a specific component or subsystem engaged in this process, usually when they perceived a technological issue to be relevant to their own subsystem, but also for other reasons (e.g., they were simply curious about a specific technology). In other words, interlacing of knowledge through justification occurred in a decentralized manner on a peer and voluntary basis without any manager imposing what needed to be shared by whom. Similarly, resulting knowledge overlaps—the outcome of this process—were not held by a few central individuals acting as managers or integrators; rather, knowledge overlaps were distributed across a multitude of participants spread across various subsystem communities.

The distributed nature and productive redundancies associated with interlaced knowledge create fundamentally different implications for coordination within and across interdependent groups when compared with knowledge overlaps between just a few individuals who act as gatekeepers or brokers (Hargadon and Sutton 1997). For example, knowledge (about problems and how these problems can be addressed) can be accessed

more efficiently by a decentralized structure with a set of direct intercommunity ties than by a central boundary spanner accessing such complex knowledge mostly through indirect ties (Zhao and Anand 2013). Furthermore, instead of specialized knowledge and trans-specialist understanding serving as substitutes (see Postrel 2002), interlacing of knowledge enables a shared division of labor (Nonaka and Takeuchi 1995) among interdependent groups.

Our findings also hold implications for how knowledge itself is created and transformed in the very act of coordination. Mechanisms such as communication and knowledge exchange identified for generating common ground are essentially information or knowledge “transfer” mechanisms (Srikanth and Puranam 2011). In contrast, the cycles of contestation and justification at ATLAS served as mechanisms for knowledge “transformation” (for a discussion of a qualitatively different transformation mechanism, see Bechky 2003). In a process of transformation, existing knowledge has to be recontextualized and translated across epistemic boundaries, resulting in both the sharing of existing knowledge and the generation of new knowledge. In this sense, the cycles of contestation and justification are also different from other committees or collective forums (such as standards bodies) that emphasize exchange and aggregation of information (e.g., through voting).

Finally, the findings from this study can inform innovation processes within collaborative and distributed communities (Adler et al. 2008, Lakhani and Panetta 2007) and self-organized networks (Fjeldstad et al. 2012). Despite the increasing importance of these new forms of organizing for innovation, we still lack an in-depth understanding of how actors engaging with one another within these decentralized systems can coordinate their activities (Garud et al. 2013). For example, communities and networks comprising actors with different backgrounds and motivations can easily fork and fragment (Kogut and Metiu 2001). In this regard, the findings from this study complement research that has explored social and technical solutions to reduce the possibility of fragmentation as collaborative innovation projects unfold (Garud et al. 2008, O’Mahony and Ferraro 2007, von Krogh et al. 2003).

Conclusion

Of continued interest to scholars and practitioners alike are multiparty collaborations to develop and deploy complex technological systems. However, what deserves greater attention are the dynamics of coordination, especially with respect to knowledge integration as a system unfolds over time. In this regard, this study draws attention to justification as a way to coordinate in real time based on a boundary infrastructure that emerges

over time. These cycles of contestation and justification generate a knowledge structure that is interlaced. This interlaced knowledge, along with the boundary infrastructure, serves as a way for collaborators to cocreate a system’s architecture.

In addition, for collaborations involving dynamic systems such as the one we studied, different coordination mechanisms operate at different points in time. For instance, the system architecture, once it emerges, facilitates embedded coordination, thereby enabling the parallel and distributed development of components. Knowledge integration issues must be addressed at various points in time, as latent interdependencies arise.

In this regard, this study shows the importance of maintaining continued interaction across different communities even as they develop subsystems and components in parallel. Such interaction maintains and enriches existing interlaced knowledge. These processes make it possible for the members of the collaboration to proactively deal with issues that emerge as integration unfolds.

Together, these insights extend our existing understanding of other coordination mechanisms and knowledge structures. Overall, they suggest the utility of looking at a suite of coordination mechanisms and how they are deployed dynamically over time. In this way, the insights from this study contribute to the growing literature on coordination through knowledge integration as it pertains to the emergence of technological systems.

Acknowledgments

The authors thank the many scientists at ATLAS, CERN, who generously contributed their time. In particular, they thank Marzio Nesi and Markus Nordberg of CERN. They also thank many members of the academic community who offered their comments on earlier drafts of this paper, including Terry Amburgey, Carliss Baldwin, Hans Berends, Paul David, Oliver Gassmann, Barbara Gray, Stefan Haeffliger, Jukka-Pekka Onnela, Andrea Prencipe, Lori Rosenkopf, Francesco Rullani, Kannan Srikanth, Mary Tripsas, and Georg von Krogh. This paper was presented at various forums including the 2012 Academy of Management Annual Meetings, DRUID Society Conference 2012, 2012 Annual Conference of the Industry Studies Association, the Wharton Technology and Innovation Conference, and the Tilburg Center for Innovation Research Conference on Innovation 2012. The authors thank the participants at these conferences for their questions and comments. They also thank Phanish Puranam and anonymous reviewers at *Organization Science* for their critical and helpful comments on previous versions of this paper. This research was partly funded by the Swiss National Science Foundation [Grant PBSG1-111192] and the U.S. National Science Foundation [Award 0738165].

Supplemental Material

Supplemental material to this paper is available at <http://dx.doi.org/10.1287/orsc.2013.0894>.

Appendix A. Chronology of Events at ATLAS

1981–1983	Two CERN experiments, UA1 and UA2, become well known for the discovery of the W and Z bosons (awarded a Nobel Prize in 1984).
March 1984	Workshop on the feasibility of a Large Hadron Collider. First ideas on novel detector concepts for future colliders based on results of UA1 and UA2.
1988–1989	In a series of workshops, groups of scientists started to converge on technologies that were potentially useful for future detector concepts.
July 1990	Start of CERN Detector Research and Development Program, a funding vehicle for 50 groups of scientists, enabling them to conduct larger R&D projects for future detector concepts.
1990–1992	Several R&D projects clustered into four larger collaborations, each proposing a complex detector consisting of various subsystems.
May 1992	When CERN urged the four collaborations to join forces, two of them, EAGLE and ASCOT, began intense discussions on the possibility of a merger because of some similarities in their designs, e.g., toroid magnet.
June 1992	Controversies around plans to merge EAGLE and ASCOT emerged because of significant differences between the two collaborations.
July 1992	Decision to merge the two collaborations to form ATLAS. Start of work on a joint Letter of Intent for the CERN LHC Committee.
October 1992	ATLAS Letter of Intent was the first official document to be submitted to CERN by the new collaboration. Decisions on architecture and technologies remained open; several alternatives were listed for each potential subsystem.
November 1992	ATLAS magnet panel evaluated competing magnet options in concert with the LHC Magnet advisory group. Decision to drop iron core toroid option in favor of the more radical air core toroid.
February 1993	Start of calorimeter review panel.
July 1993	Start of muon spectrometer review panel.
September 1993	Start of inner detector review panel.
September 1993	Calorimeter review panel recommendation for a hybrid calorimeter consisting of different technologies (considered revolutionary at that time). Panel recommended that a decision on forward calorimeter be delayed by one year to enable additional studies.
December 1993	Muon spectrometer review panel issued its recommendation. The recommendation caused a revolt in the community because it was perceived as not based on justification and it was not transparent.
February 1994	Arbitration among conflicting groups to alleviate controversy in the muon community. To reach consensus, a “conclave” was initiated, which was chaired by an independent convener.
April 1994	Breakthrough in muon spectrometer situation. Entire community agreed to develop a new concept making best use of features of each of the three competing technologies.
April 1994	Position of solenoid magnet adversely affected many subsystems. Collaboration realized that a global study would be required to assess overall physics performance implications.
June 1994	New conceptual design of inner detector took advantage of contributions by new ATLAS members who had joined the collaboration.
September 1994	Calorimeter review panel recommended integrating forward calorimeter. This change of architecture (integration of calorimeter into endcap toroid) created positive performance implications for other subsystems.
September 1994	Muon community agreed on trigger chamber technology for the muon spectrometer.
December 1994	ATLAS Technical Proposal, detailing the architecture and components of the ATLAS detector, was submitted to CERN.
January 1995	Controversies among subsystems around space for electronics and services (cooling, cables, etc.). Several areas were “overbooked” and subsystems were increasing even further.
June 1995	Concerns about increasing cost of the large superconducting toroid magnet system surfaced. This triggered discussions about a simplified or downscaled design.
September 1995	Inner detector review panel recommended eliminating one technology to simplify the inner detector architecture.
September 1995	“Global reconsideration of detector” aimed at cost savings with little loss of physics performance; mandate was given to a “Descoping Taskforce.”
November 1995	Descoping Taskforce recommended (a) reduction of ATLAS dimensions, (b) simplification of subsystems by reducing and homogenizing components, and (c) reoptimizing overall configuration of detector to preserve physics performance as well as possible.
March 1996	Muon community proposed reoptimization of muon spectrometer layout.
September 1996	Collaboration approved changes to the calorimeter and inner detector.

Appendix A (cont'd)

November 1996	Review of competing silicon technology options for inner detector.
December 1996	Approval of Calorimeter Technical Design Report.
February 1997	Concerns about cooling systems suggested by inner detector community. A special task force was set up to review alternatives for inner detector cooling system.
March 1997	Approval of new optimized layout of the muon spectrometer.
April 1997	Approval of Inner Detector Technical Design Report.
June 1997	Approval of Muon Spectrometer Technical Design Report.
September 1997	Inner detector community proposed a staged design of the semiconductor tracker and pixel detector (i.e., some layers were removed that could be replaced and upgraded later if required).
November 1997	Inner detector cooling review panel proposed new cooling system for inner detector.
June 1998	ATLAS Memorandum of Understanding was signed by most funding agencies, which constitutes beginning of construction phase. Construction sites were prepared and tested by producing “module-0” prototypes.
September 1998	Controversy around one institute’s intention to deviate from the design of its deliverable for cost reasons. The collaboration argued strongly against this plan.
June 1999	The inner detector community proposed changes to the pixel detector component due to manufacturing problems of radiation hard silicon sensors.
February 2000	ATLAS Memorandum of Understanding was signed by all funding agencies. Most detector systems had entered the construction phase.
February 2001	A new estimate of cost to completion showed a lack of funding on the order of 20%. A staging scenario with an initial detector consistent with the needs of the initial low-luminosity physics run was integrated into the work schedule.
June 2001	Expected delay of LHC completion allowed additional time for completion of ATLAS; this had implications on design decisions because extra time offered new opportunities.
June 2003	ATLAS cavern (100 m below surface) completed and ready for installation. First components were lowered down and installed later that year.
September 2003	Change in inner detector due to staging of TRT components guaranteed feasibility of inner detector completion. The missing components were compensated for by introducing additional layers in other components (pixel detector).
June 2005	ATLAS high-luminosity R&D projects for future upgrade of ATLAS detector evolved.
February 2008	Completion of LHC.
September 2008	Commissioning of ATLAS detector. Detection of first beam events on September 10, 2008.

Appendix B. Visualization of Flow of Events

Appendix C. Measuring the Knowledge Structure of Subsystem Communities

Notes. The matrix indicating the knowledge domains (K1, K2, etc.) of scientists (S1, S2, etc.) who belong to different subsystem communities (Communities A and B) can be conceptualized as a network representing the knowledge structure of various ATLAS subsystem communities, specifically as a bipartite graph between participants and knowledge domains. As panel (a) illustrates, knowledge domains in this bipartite graph are connected when participants have expertise in diverse knowledge domains. These connections become denser and the knowledge domains become more central when knowledge domains of multiple participants overlap. The connections among various concepts become even more apparent in a unipartite projection of the bipartite graph (see panel (b)). In such a unipartite projection, the nodes are the knowledge domains, whereas the edges connecting the knowledge domains are the participants. We created such network representations for different ATLAS subsystem communities (see (b1) and (b2)). Panel (c) represents the knowledge structure of the overall collaboration, including all subsystem communities. For a more detailed description, see the online appendix, available as supplemental material at <http://dx.doi.org/10.1287/orsc.2013.0894>.

Endnotes

¹The Higgs boson, also known as the God particle, is an elementary particle in the Standard Model of particle physics named after Peter Higgs, who predicted it on theoretical grounds.

²Given our research question, we focus primarily on this aspect of coordination—i.e., the coordination of the development

of complex technological systems by multiple interdependent actors—and the mechanisms for accomplishing such coordination. For a comprehensive review of coordination mechanisms in a broader organizational context, see Okhuysen and Bechky (2009).

³An established architecture may emerge in many ways, e.g., as a dominant design (Abernathy and Utterback 1978,

Tushman and Anderson 1986), defined by a standards body (Farrell and Saloner 1988), or based on prior generations of a system (Sosa et al. 2004).

⁴Common ground can be manifest at different levels of aggregation—for instance, at the dyadic level (Srikanth and Puranam 2011), between communities of knowledge (Hoopes and Postrel 1999), or at the organizational level (Grant 1996b).

⁵This is similar to the notion of “interactional expertise” that emerges from “linguistic socialization” (Collins and Evans 2007, p. 88).

⁶Garud and Munir (2008) described how at first the development of a new battery for Polaroid’s SX-70 camera was assigned to ESB, which had expertise in developing batteries. However, when the combination of the film and the battery started to create problems, Polaroid decided to internalize the design and development of the battery and the film pod to overcome the problem.

⁷We analyzed levels of justification and knowledge structures of communities working on all three major ATLAS subsystems; however, the levels of transparency, scrutiny, and justification were initially perceived (by those involved and the collaboration at large) to be significantly lower within the muon spectrometer subsystem community than within the calorimeter subsystem and inner detector subsystem communities.

⁸A more detailed explanation of our methods pertaining to latent semantic analysis is included in the online appendix.

⁹Although it was known at the outset that different subsystems were required to perform certain functions (e.g., muon spectrometer for measuring muons), it was not clear what approaches/technologies should be used to design the various subsystems or how to configure the various subsystems and the connections between them. In other words, the ATLAS architecture was not given. Indeed, different collaborations at the LHC used different approaches, each with its own distinct detector architecture. For instance, another collaboration, CMS (Compact Muon Solenoid), shared the ATLAS collaboration’s objective (i.e., to detect the Higgs boson), and its detector too had components labeled calorimeter, muon spectrometer, etc.; however, the CMS detector’s architecture (i.e., constituent components, interfaces, and configurations) was fundamentally different from that of the ATLAS detector (ATLAS Collaboration 1994, CMS Collaboration 1994).

¹⁰In fact, an early attempt—the “parity commission” proposed before the merger to ensure that the interests of both proto-collaborations were maintained and brought together—failed and was discarded after being unfavorably perceived by participants as a move to centralize decision making.

¹¹Participants allowed for some flexibility by agreeing to what they called “envelopes,” i.e., acceptable tolerances for agreed-upon specifications for different detector subsystems.

¹²Binary ice is a coolant consisting of ice crystals in a cooling liquid pumped through a complex system of pipes.

¹³This difference was perceived by both community insiders and the collaboration at large.

¹⁴This was an exception in the history of ATLAS, as it was the only case in which the collaboration board declined to accept the recommendation made by a review panel.

¹⁵*Synonymy* is a situation where different words describe the same idea, whereas *polysemy* is a situation where the same words describe different ideas.

¹⁶We also calculated *closeness centrality* (Freeman 1979), an alternative network measure that indicates how easily information located in a network can be accessed by any node in the network, either directly or indirectly. Our analysis using closeness centrality is consistent with our reported findings.

¹⁷At ATLAS, the detector was considered to be its own prototype. It was impossible to create a working prototype that could then be tested, refined, and scaled up before final commissioning. As a result, they had to get it right the first time.

References

- Abernathy WJ, Utterback JM (1978) Patterns of industrial innovation. *Tech. Rev.* 80(7):40–47.
- Adler PS, Kwon S-W, Heckscher C (2008) Professional work: The emergence of collaborative community. *Organ. Sci.* 19(2): 359–376.
- Alexander C (1964) *Notes on the Synthesis of Form* (Harvard University Press, Cambridge, MA).
- Anderson P, Tushman ML (1990) Technological discontinuities and dominant designs: A cyclical model of technological change. *Admin. Sci. Quart.* 35(4):604–633.
- ATLAS Collaboration (1994) ATLAS technical proposal. Report, CERN, Geneva.
- Baldwin CY, Clark KB (2000) *Design Rules, Vol. 1: The Power of Modularity* (MIT Press, Cambridge, MA).
- Barry D, Rerup C (2006) Going mobile: Aesthetic design considerations from Calder and the constructivists. *Organ. Sci.* 17(2):262–276.
- Bechky BA (2003) Sharing meaning across occupational communities: The transformation of understanding on a production floor. *Organ. Sci.* 14(3):312–330.
- Berends H, Garud R, Debackere K, Weggeman M (2011) Thinking along: A process for tapping into knowledge across boundaries. *Internat. J. Tech. Management* 53(1):69–88.
- Bowker GC, Star SL (1999) *Sorting Things Out: Classification and Its Consequences* (MIT Press, Cambridge, MA).
- Brown JS, Duguid P (1991) Organizational learning and communities-of-practice: Toward a unified view of working, learning, and innovation. *Organ. Sci.* 2(1):40–57.
- Brown SL, Eisenhardt KM (1997) The art of continuous change: Linking complexity theory and time-paced evolution in relentlessly shifting organizations. *Admin. Sci. Quart.* 42(1):1–34.
- Brunsoni S, Prencipe A, Pavitt K (2001) Knowledge specialization, organizational coupling, and the boundaries of the firm: Why do firms know more than they make? *Admin. Sci. Quart.* 46(4):597–621.
- Callon MN (1995) Four models for the dynamics of science. Jasanoff S, Markle GE, Petersen JC, Pinch TJ, eds. *Handbook of Science and Technology Studies* (Sage, London), 29–63.
- Callon MN, Law J, Rip A (1986) *Mapping the Dynamics of Science and Technology: Sociology of Science in the Real World* (Macmillan, London).
- Camerer C, Knez M (1996) Coordination, organizational boundaries and fads in business practices. *Indust. Corporate Change* 5(1):89–112.
- Carlile PR (2002) A pragmatic view of knowledge and boundaries: Boundary objects in new product development. *Organ. Sci.* 13(4):442–455.
- Clark KB (1985) The interaction of design hierarchies and market concepts in technological evolution. *Res. Policy* 14(5):235–251.

- CMS Collaboration (1994) The Compact Muon Solenoid technical proposal. Report, CERN, Geneva.
- Colfer L, Baldwin CY (2010) The mirroring hypothesis: Theory, evidence and exceptions. Working Paper 10-58, Harvard Business School, Boston.
- Collins HM, Evans R (2007) *Rethinking Expertise* (University of Chicago Press, Chicago).
- Deerwester S, Dumais S, Furnas GW, Landauer TK, Harshman R (1990) Indexing by latent semantic analysis. *J. Amer. Soc. Inform. Sci.* 41(6):391–407.
- Ellis N, Froidevaux D, Gianotti F, Guyot C, Hoffmann H, Jenni P, Linde FL, Nisati A, Pilcher JE, Rossi L (1995) Report of the Global Descoping Task Force. Report ATL-GEN-95-014, CERN, Geneva.
- Ethiraj SK, Levinthal D (2004) Modularity and innovation in complex systems. *Management Sci.* 50(2):159–173.
- Farrell J, Saloner G (1988) Coordination through committees and markets. *RAND J. Econom.* 19(2):235–252.
- Fjeldstad OD, Snow CC, Miles RE, Lettl C (2012) The architecture of collaboration. *Strategic Management J.* 33(6):734–750.
- Freeman LC (1979) Centrality in social networks: Conceptual clarification. *Soc. Networks* 1(3):215–239.
- Galison PL (1997) *Image and Logic: A Material Culture of Microphysics* (University of Chicago Press, Chicago).
- Garud R, Karnoe P (2003) Bricolage versus breakthrough: Distributed and embedded agency in technology entrepreneurship. *Res. Policy* 32(2):277–300.
- Garud R, Kotha S (1994) Using the brain as metaphor to model flexible production systems. *Acad. Management Rev.* 19(4):671–698.
- Garud R, Kumaraswamy A (1995) Technological and organizational designs to achieve economies of substitution. *Strategic Management J.* 16(S1):93–109.
- Garud R, Munir K (2008) From transactions to transformation costs: The case of Polaroid's SX-70 camera. *Res. Policy* 37(4):690–705.
- Garud R, Rappa MA (1994) A socio-cognitive model of technology evolution: The case of cochlear implants. *Organ. Sci.* 5(3):344–362.
- Garud R, Jain S, Kumaraswamy A (2002) Institutional entrepreneurship in the sponsorship of common technological standards: The case of Sun Microsystems and Java. *Acad. Management J.* 45(1):196–214.
- Garud R, Jain S, Tuertscher P (2008) Designing for incompleteness and incomplete by design. *Organ. Stud.* 29(3):351–371.
- Garud R, Tuertscher P, Van de Ven AH (2013) Perspectives on innovation processes. *Acad. Management Ann.* 7(1):775–819.
- Genschel P (1997) How fragmentation can improve co-ordination: Setting standards in international telecommunications. *Organ. Stud.* 18(4):603–622.
- Granovetter MS (1992) Problems of explanation in economic sociology. Nohria N, Eccles RG, eds. *Networks and Organizations. Structure, Form and Action* (Harvard Business School Press, Boston), 25–56.
- Grant RM (1996a) Prospering in dynamically-competitive environments: Organizational capability as knowledge integration. *Organ. Sci.* 7(4):375–387.
- Grant RM (1996b) Toward a knowledge-based theory of the firm. *Strategic Management J.* 17(S1):109–122.
- Green SE Jr (2004) A rhetorical theory of diffusion. *Acad. Management Rev.* 29(4):653–669.
- Gulati R, Singh H (1998) The architecture of cooperation: Managing coordination costs and appropriation concerns in strategic alliances. *Admin. Sci. Quart.* 43(4):781–814.
- Hardwig J (1985) Epistemic dependence. *J. Philos.* 82(7):335–349.
- Hargadon A, Sutton RI (1997) Technology brokering and innovation in a product development firm. *Admin. Sci. Quart.* 42(4):716–749.
- Henderson RM, Clark KB (1990) Architectural innovation: The reconfiguration of existing product technologies and the failure of established firms. *Admin. Sci. Quart.* 35(1):9–30.
- Hobday M (1998) Product complexity, innovation and industrial organisation. *Res. Policy* 26(6):689–710.
- Hoopes DG, Postrel S (1999) Shared knowledge, “glitches,” and product development performance. *Strategic Management J.* 20(9):837–865.
- Jones C, Hesterly WS, Borgatti SP (1997) A general theory of network governance: Exchange conditions and social mechanisms. *Acad. Management Rev.* 22(4):911–945.
- Kellogg KC, Orlikowski WJ, Yates J (2006) Life in the trading zone: Structuring coordination across boundaries in postbureaucratic organizations. *Organ. Sci.* 17(1):22–44.
- Knorr-Cetina K (1995) How superorganisms change: Consensus formation and the social ontology of high-energy physics experiments. *Soc. Stud. Sci.* 25:119–147.
- Knorr-Cetina K (1999) *Epistemic Cultures: How the Sciences Make Knowledge* (Harvard University Press, Cambridge, MA).
- Kogut B, Metiu A (2001) Open-source software development and distributed innovation. *Oxford Rev. Econom. Policy* 17(2):248–264.
- Kotha S, Srikanth K (2013) Managing a global partnership model: Lessons from the Boeing 787 “Dreamliner” program. *Global Strategy J.* 3(1):41–66.
- Lakhani KR, Panetta JA (2007) The principles of distributed innovation. *Innovations: Tech., Governance, Globalization* 2(3):97–112.
- Landauer TK, Foltz PW, Laham D (1998) An introduction to latent semantic analysis. *Discourse Processes* 25(2–3):259–284.
- Langley A (1999) Strategies for theorizing from process data. *Acad. Management Rev.* 24(4):691–710.
- Latour B (1987) *Science in Action: How to Follow Scientists and Engineers Through Society* (Harvard University Press, Cambridge, MA).
- Latour B (2004) Why has critique run out of steam? From matters of fact to matters of concern. *Critical Inquiry* 30(2):225–248.
- Laughlin PR, Ellis AL (1986) Demonstrability and social combination processes on mathematical intellectual tasks. *J. Experiment. Soc. Psych.* 22(3):177–189.
- Laughlin PR, Hollingshead AB (1995) A theory of collective induction. *Organ. Behav. Human Decision Processes* 61(1):94–107.
- Lave J, Wenger E (1991) *Situated Learning: Legitimate Peripheral Participation* (Cambridge University Press, Cambridge, UK).
- Lincoln YS, Guba EG (1985) But is it rigorous? Trustworthiness and authenticity in naturalistic evaluation. Williams DD, ed. *Naturalistic Evaluation* (Jossey-Bass, San Francisco), 73–84.
- Miles MB, Huberman AM (1984) *Analyzing Qualitative Data: A Source Book for New Methods* (Sage, Beverly Hills, CA).

- Mitchell JC (1969) The concept and use of social networks. Mitchell JC, ed. *Social Networks in Urban Situations* (Manchester University Press, Manchester, UK), 1–50.
- Morgan G (1986) *Images of Organization* (Sage, Beverly Hills, CA).
- Nonaka I, Takeuchi H (1995) *The Knowledge-Creating Company* (Oxford University Press, Oxford, UK).
- Nonaka I, von Krogh G (2009) Tacit knowledge and knowledge conversion: Controversy and advancement in organizational knowledge creation theory. *Organ. Sci.* 20(3):635–652.
- Okhuysen GA, Bechky BA (2009) Coordination in organizations: An integrative perspective. *Acad. Management Ann.* 3(1):463–502.
- O'Mahony S, Ferraro F (2007) The emergence of governance in an open source community. *Acad. Management J.* 50(5):1079–1106.
- Palla G, Barabási A-L, Vicsek T (2007) Quantifying social group evolution. *Nature* 446(7136):664–667.
- Pickering A (1993) The mangle of practice—Agency and emergence in the sociology of science. *Amer. J. Sociol.* 99(3):559–589.
- Pinch TJ, Bijker WE (1987) The social construction of facts and artifacts: Or how the sociology of science and the sociology of technology might benefit each other. Bijker WE, Hughes TP, Pinch TJ, eds. *The Social Construction of Technological Systems: New Directions in the Sociology and History of Technology* (MIT Press, Cambridge, MA), 17–50.
- Postrel S (2002) Islands of shared knowledge: Specialization and mutual understanding in problem-solving teams. *Organ. Sci.* 13(3):303–320.
- Puranam P, Raveendran M, Knudsen T (2012) Organization design: The epistemic interdependence perspective. *Acad. Management Rev.* 37(3):419–440.
- Puranam P, Singh H, Chaudhuri S (2009) Integrating acquired capabilities: When structural integration is (un)necessary. *Organ. Sci.* 20(2):313–328.
- Rosenkopf L, Tushman ML (1998) The coevolution of community networks and technology: Lessons from the flight simulation industry. *Indust. Corporate Change* 7(2):311–346.
- Rosenkopf L, Metiu A, George VP (2001) From the bottom up? Technical committee activity and alliance formation. *Admin. Sci. Quart.* 46(4):748–772.
- Sanchez R, Mahoney JT (1996) Modularity, flexibility, and knowledge management in product and organization design. *Strategic Management J.* 17(S1):63–76.
- Simcoe T (2012) Standard setting committees: Consensus governance for shared technology platforms. *Amer. Econom. Rev.* 102(1):305–336.
- Simon HA (1962) The architecture of complexity. *Proc. Amer. Philos. Soc.* 106(6):467–482.
- Sosa ME, Eppinger SD, Rowles CM (2004) The misalignment of product architecture and organizational structure in complex product development. *Management Sci.* 50(12):1674–1689.
- Srikanth K, Puranam P (2011) Integrating distributed work: Comparing task design, communication, and tacit coordination mechanisms. *Strategic Management J.* 32(8):849–875.
- Stapnes S (2007) Detector challenges at the LHC. *Nature* 448(7151):290–296.
- Star SL, Griesemer JR (1989) Institutional ecology, “translations” and boundary objects: Amateurs and professionals in Berkeley's Museum of Vertebrate Zoology, 1907–39. *Soc. Stud. Sci.* 19(3):387–420.
- Starbuck WH (1983) Organizations as action generators. *Amer. Sociol. Rev.* 48(1):91–102.
- Staudenmayer N, Tripsas M, Tucci CL (2005) Interfirm modularity and its implications for product development. *J. Product Innovation Management* 22(4):303–321.
- Strauss A, Corbin JM (1998) *Basics of Qualitative Research* (Sage, Newbury Park, CA).
- Taylor JR, Van Every EJ (2000) *The Emergent Organization: Communication as the Site and Surface* (Lawrence Erlbaum Associates, Mahwah, NJ).
- Toulmin S (1983) *The Use of Argument*, 2nd ed. (Cambridge University Press, New York).
- Trochim WMK (1989) Outcome pattern matching and program theory. *Evaluation Program Planning* 12(4):355–366.
- Tushman ML, Anderson P (1986) Technological discontinuities and organizational environments. *Admin. Sci. Quart.* 31(3):439–465.
- Tushman ML, Murmann JP (1998) Dominant designs, technology cycles, and organizational outcomes. Staw BM, Cummings LL, eds. *Research in Organizational Behavior*, Vol. 20 (JAI Press, Greenwich, CT), 231–266.
- Tushman ML, Rosenkopf L (1992) Organizational determinants of technological change: Toward a sociology of technological evolution. Staw BM, Cummings LL, eds. *Research in Organizational Behavior*, Vol. 14 (JAI Press, Greenwich, CT), 311–347.
- Ulrich K (1995) The role of product architecture in the manufacturing firm. *Res. Policy* 24(3):419–440.
- Van de Ven AH (1992) Suggestions for studying strategy process: A research note. *Strategic Management J.* 13(S1):169–188.
- von Krogh G, Spaeth S, Lakhani KR (2003) Community, joining, and specialization in open source software innovation: A case study. *Res. Policy* 32(7):1217–1241.
- Washington H, Cook J (2011) *Climate Change Denial: Heads in the Sand* (Earthscan, London).
- Webb E, Weick KE (1979) Unobtrusive measures in organizational theory: A reminder. *Admin. Sci. Quart.* 24(4):650–659.
- Weick KE, Roberts KH (1993) Collective mind in organizations: Heedful interrelating on flight decks. *Admin. Sci. Quart.* 38(3):357–381.
- Weick KE, Sutcliffe KM (2001) *Managing the Unexpected: Assuring High Performance in an Age of Complexity* (Jossey-Bass, San Francisco).
- Yakura EK (2002) Charting time: Timelines as temporal boundary objects. *Acad. Management J.* 45(5):956–970.
- Yin RK (1994) *Case Study Research: Design and Methods*, 2nd ed. (Sage, Thousand Oaks, CA).
- Zhao ZJ, Anand J (2013) Beyond boundary spanners: The “collective bridge” as an efficient interunit structure for transferring collective knowledge. *Strategic Management J.* 34(13):1513–1530.
- Zhou YM (2013) Designing for complexity: Using divisions and hierarchy to manage complex tasks. *Organ. Sci.* 24(2):339–355.
- Zuboff S (1988) *In the Age of the Smart Machine: The Future of Work and Power* (Basic Books, New York).

Philipp Tuertscher is an assistant professor of technology and innovation and a member of the Knowledge, Information, and Networks research group at the VU University Amsterdam. His research explores organizational mechanisms and social practices for collaborative innovation in a variety of

settings. Besides studying large-scale scientific collaborations at CERN, he has been studying innovation processes in collaborative communities such as Linux and Wikipedia.

Raghu Garud is the Alvin H. Clemens Professor of Management and Organization and the research director of the Farrell Center for Corporate Innovation and Entrepreneurship, Pennsylvania State University. His research explores the microprocesses involved in the emergence of novelty, a theme

around which he coconvened a workshop at Crete under the aegis of Process Organization Studies (PROS).

Arun Kumaraswamy studies the management of innovation and growth, and technology catch-up strategies of domestic firms in emerging economies. He coedited the book *Managing in the Modular Age: Architectures, Networks and Organizations* and serves on the editorial board of the *Journal of International Management*.