The Emergence of Architecture – The Case of ATLAS, CERN

Philipp Tuertscher
Vienna University of Economics and Business Administration
philipp.tuertscher@wu-wien.ac.at

Raghu Garud
Pennsylvania State University
rgarud@psu.edu

Markus Nordberg
CERN
markus.nordberg@cern.ch

DRAFT

November 15th, 2007

Presented at the Conference
Studying Path Dependencies of Businesses,
Institutions and Technologies

Berlin

The Emergence of Architecture – The Case of ATLAS, CERN[1]

ABSTRACT

In this paper, we are looking at the development of ATLAS, a complex technological system that is one-of-a-kind and involves several nascent technologies. The focus of our study is the early phase of design when the structures incorporated in the technology have not yet stabilized and the architecture is still emerging. The study of this phase shows the emergence of technological and organizational paths as different groups and technologies involved in the process interact with each other. Whereas prior literature has emphasized an unambiguous understanding of architectures, we believe that it is important to consider conflicting interpretations in order to understand how architectures emerge. In particular, we study how controversies resulting from conflicting interpretations get resolved as the different technological components and groups interact with each other. We suggest that this is the process whereby new technological paths emerge over time.

[1] An earlier version of this paper was presented at the 2007 EGOS Colloquium track on Path Dependence and at Berlin Workshop on Paths of Developing Complex Technologies. We would like to thank Simone Ferriani, Guido Moellring, Arie Rip, Cornelius Schubert, Joerg Sydow for their valuable comments which helped to improve the draft of this paper.

INTRODUCTION

The development of many complex technologies is organized in large innovation systems. In such systems, different contributors – e.g., corporations or universities – are working on components of the larger system. The rationale for this division of labor is based on design principles from early work on design science. In his seminal work on the architecture of complexity, Herbert Simon (1962) emphasized that decomposition of complex problems into smaller subtasks which are easier to fulfill is a useful strategy to reduce complexity and facilitate problem solving.

However, in order to guarantee a working system, the distributed efforts eventually need to be integrated. Integration requires coordination across boundaries throughout the distributed effort (Clark & Fujimoto, 1990; Ulrich, 2003). For the development of a complex technological system, this means that the many parts of the final system – usually consisting of many different technologies – need to fit seamlessly. The integration does not require physical fit alone; it also comprises effective interplay and exchange of information between the various components which is usually required for complex systems to perform the expected functionality. Successful integration requires coordination throughout the innovation process. The need for coordination starts even before the actual development when the design problem is split up into subtasks and allocated to different contributors (Simon, 1996). During this early phase, it needs to be specified where the boundaries of the various components should be drawn, which functions they need to perform and what the inputs and outputs of each system should be.

Given the uncertainties involved in innovation, coordination can be very problematic. Since many requirements of the system need to be identified first and features of the system are yet to be developed, it is virtually impossible to find the optimal decomposition of the technological system *ex ante*. Moreover, the development of some components may change the requirements for other parts of the system, rendering them obsolete as the design evolves.

As a consequence, many decisions in the design process are of preliminary nature only (Mihm, Loch, & Huchzermeier, 2003).

What is the process whereby innovation systems are coordinated in a manner that allows for such emergence? In order to address this question, we discuss in the following approaches for self-organization proposed by literature on modularity. Proponents of this approach suggest that the development of nearly decomposable systems can be coordinated by pre-specifying interfaces between interdependent components. We briefly discuss the limitations of this perspective, in particular with respect of emergent architectures. The fact that architectures need to be created in fist place has been largely neglected by extant literature.

We briefly describe our perspective on the interactive emergence of technological architectures. While social processes such as politics play an important role in structuring an innovation project, also properties of technologies determine the structure of technological architecture. However, neither of both factors can explain the emergence of architectures independently. People generate social constructions of technology using resources, interpretive schemes, and norms embedded in the larger institutional context (DeSanctis & Poole, 1994; Orlikowski, 1992).

Finally, we present a longitudinal case study of the development of a complex innovation project, the ATLAS experiment at CERN. Whereas prior literature on modularity has focused on clearly understood specifications, we argue that it is important to consider the case of ambiguous specifications in order to understand how architectures emerge. In particular, we describe how controversies resulting from conflicting interpretations got resolved as the different technological components and groups were confronted with each other. We conclude with the finding the ongoing negotiations create interlaced knowledge that facilitates coordination of the distributed development efforts as the architecture emerges.

MODULAR ARCHITECTURES

Existing literature on management of technology and innovation points to the role of architecture for the coordination of innovation systems, in particular to modular architectures (Baldwin & Clark, 2000). A modular architecture has certain advantages over other designs with respect to coordination (Sanchez & Mahoney, 1996). It consists of modules which can be used as self-contained building blocks for higher level systems. Each module only interacts with another through standardized interfaces that define functional, spatial, and other relationships between the components (Langlois & Roberts, 1995; Sanchez, 1995; Schilling, 2000). This way, modular architectures allow for a range of variations in its loosely coupled components without having to change the designs of other components (Garud & Kumaraswamy, 1995). As long as the components conform to the standardized interface specifications, there is no need to manage directly or monitor the processes of the individual component developers. While the modular components can be developed autonomously by different groups, the coordination of the overall development process is embedded in the interface specifications provided by the modular architecture (Sanchez, 1995).

Modularity enables a shift of coordination from relying on the overt exercise of managerial authority to effective coordination embedded into a modular architecture. The reduced need for managerial coordination not only lowers cost and difficulty of managing complex processes (Sanchez, 1995), it also allows different parts of a large design to be worked on concurrently, shortening the end-to-end time needed to complete a given process (Garud & Kotha, 1994). In addition to this advantage in terms of coordination, modularity increases the range of manageable complexity since the scope of interaction between elements or tasks is limited by agreed-upon standard interface specifications (Baldwin et al., 2000). Finally, it has been argued that modularity is a means for accommodating uncertainty. Modular architectures are flexible since the standardized interface specifications allow for a range of variations in components (Ulrich & Eppinger, 2000) which may be substituted in response to market and technological changes (Garud et al., 1995; Sanchez & Mahoney, 1994).

Proponents of modularity assume that specifications of components and interfaces are complete and understood by people involved in the innovation process. This is usually the case in mature industries such as automobile and aircraft industries in which technologies are based on dominant designs that have emerged over decades. The architecture of dominant designs is known and taken-for-granted by specialists involved in the development process. In many innovation systems, however, architectures need to be created in first place. In the case of nascent technologies or complex one-of-a-kind systems, for example, interfaces are just emerging and are difficult to standardize due to a lack of maturity of the field (Wood & Brown, 1998). Yet, there is no extant literature as to how architectures emerge. Precisely this question we address in this paper. To shed light on this question, we study the development of ATLAS, a complex technological system that is being developed at CERN, the European Center for Nuclear Research. The focus of our research is on the early phase of design when the architecture of the technology has not yet stabilized. The aim of our study is to contribute to our understanding of complex innovation processes by an in-depth analysis as to how architectures of complex technological systems emerge.

How do novel architectures emergence?

Current research on modularity has focused on the natural partitioning of complex technological systems, implicitly assuming that architecture is an inherent structure of complex systems to be "discovered". Baldwin and Clark (2003) suggest that technological systems have "natural transfer bottlenecks" which are optimal locations for decomposing a complex technological system into components. Recent contributions made efforts to provide tools and heuristics such as the design structure matrix (Baldwin & Clark, 2005; Eppinger, Whitney, Smith, & Gebala, 1994) for identifying natural bottlenecks between partitions of complex systems. Such tools providing "design rules" (Baldwin et al., 2000) seem to be useful to organize and improve the architecture of existing systems; however, partitioning of novel systems is more problematic. Since many interdependencies between components of emerging technological systems are latent (Alexander, 1964) and become visible only with the progress of development work, it is unfeasible to identify natural boundaries within novel systems, even if they may seem obvious after the complex system has emerged.

Moreover, it is difficult to discover the "inherent" architecture when both the technologies required to build the different components as well as possible configurations of the overall system are emergent. The possibility to identify interdependencies between components *ex ante* depends on the novelty of both the system and the technologies used for its components. Even if emerging architectures are to a large extent building upon existing technologies, the novel way of linking existing components may cause interferences that had not been observed in prior systems (Barry & Rerup, 2006). Rather than assuming that rational choice enables designers to identify an inherent structure in technological systems, we would like to draw attention to other factors that influence the emergence of architectures. In particular, the structure of complex systems is not only based on technological determinism but also shaped by the social context in which it is embedded. In their work on complex technological systems, Sydow, Windeler, Schubert and Moellring (2007) suggested that the interplay of technological and social factors should be taken into account for the understanding of emergence of architecture.

For example, institutional pressures from existing technologies in form of dominant designs (Tushman & Murmann, 2003) can lock designs into existing paths and prevent from adopting innovative, more appropriate architectures (Postrel, 2002). Similarly, architectures embedded into a larger technological ecosystem are shaped by external standards imposed by institutional authorities or networks (Langlois & Robertson, 2003). But also social structures are likely to shape emergent architectures. While organizational units originally have been structured to reflect interdependencies of existing designs (Sosa, Eppinger, & Rowles, 2004), novel architectures are likely to involve interdependencies across specialized units. Boundaries between organizational units can impose knowledge barriers that inhibit the identification of interdependencies between components (Henderson & Clark, 1990) and make the adoption of a more appropriate architecture more difficult (Carlile, 2004).

It is important to note, however, that technological systems are not simply determined by technological and social institutions. Strategic agency of social actors makes it possible to deviate from existing structures and constitute a novel technological path, for example by changing interdependencies. The type of interdependence is not necessarily an inherent property of a task environment as it was considered until recently (Thompson, 1967). Instead, actors may construct and modify dependencies, for example by changing the sequencing of sub-tasks from linear to parallel (Garud et al., 1994; Grant, 1996), and thereby change the system architecture. It is important to appreciate this role of strategic agency and mutual constitution of technology and social structure (Sydow & Windeler, 1998).

Moreover, technology is considered to have interpretive flexibility. Research on structuration found that differences in interpretation of technology by people result in different technological and social structures (DeSanctis et al., 1994). Therefore, it may not come as a surprise that many different solutions are possible when the architecture of an emergent innovation system is defined. Even if designs are highly constrained by the laws of physics and logic – research on modularity would suggest very inflexible structures in that case (Baldwin et al., 2003: 35) – fundamentally different designs can emerge. This indeterminacy can be illustrated with two sister experiments conducted at CERN. Although both designs were based on the same physics requirements, they have completely different architectures using fundamentally different technologies (cf. CERN, 1994a; CERN, 1994b).

When multiple groups contribute to a design, the specification of architectures often becomes a process of political negotiation and some perspectives often win out over others.

(Anand & Watson, 2004; Carlile, 2004). This political influence can be very high when certain communities are privileged to lead the specification process whereas other communities are strictly regimented and expected to follow (Brown & Duguid, 2001). But social processes not only play a role if groups use their political power in pursuit of their selfish interest. Even when all actors involved are willing to cooperate and try to achieve the best solution possible for the overall system, political decisions remain an issue. Since designers in a complex innovation project need to be deeply knowledgeable in their respective domains but cannot know everything about how the other things are made, this division of cognitive labor (Aoki, 2001) results in groups having different thought worlds (Dougherty, 1992).

Ongoing negotiation of controversies

The different technological, institutional and social contexts of the actors involved in a complex innovation project result in collaborations consisting of multiple communities. There are bound to be controversies as members from different communities clash (Garud & Ahlstrom, 1997). Creative tension arises from controversy, but in order to make the innovation project successful, the communities need to transform the obstacles posed by incommensurate thought worlds into a synthesis from the ongoing juxtaposition of dialectic forces.

In the emergence of architectures, conflicting perspectives materialize in controversies about interfaces and system boundaries. Key design decisions revolve around issues such as detailing the boundaries of each component and establishing stable interface specifications to ensure smooth functioning between modules. There can be a number of interdependencies between components that make it difficult to identify one best solution, partly because the components are sharing a common context such as space. For example, more generous space for one system usually implies a more confined space for the other system. If each group has the goal of optimize the performance of its component, there are bound to be conflicting claims how the geometrical boundaries between components need to be specified. Similar controversies are likely to emerge around other specifications such as thermal or electromagnetic properties of interdependent components.

The resolution of such controversies is characterized by “contests of unfolding” (Knorr-Cetina, 1999), a gradual but continuous process of confrontation with other groups who propose alternative concepts. When groups make claims about the architecture required for superior performance of their component, they need to present justifying reasoning and evidence in order to convince competing groups that the proposed specification is favorable for the overall system too. In this interactive process, many different points of view and component-specific perspectives are articulated and eventually result in a reconfiguration and recontextualization of the available knowledge. The conflicts are resolved only if the communities negotiate and explicitly agree on shared understanding of the technological architecture.

We suggest that specifications are not always clearly specified and understood as assumed by research on modularity (Baldwin et al., 2003; Sanchez et al., 1996). By contrast, we conceptualize modular architectures as boundary objects (Star & Griesemer, 1989) spanning different communities. From this perspective, modular architectures have meaning which is both commonly recognized and yet differently applied by each group. As a consequence, subtle differences in the understanding of the specifications may exist and conflicts between the different components may arise. While this interpretive flexibility offers degrees of freedom for the different groups to accommodate their diverse contexts, specifications which are too ambiguous may create a lot of cycling across component boundaries (Baldwin et al., 2003). Ongoing negotiation helps to understand what the other communities meant by a specification, and translate that understanding into technical terms that one could work with.

Second, ongoing negotiation is required even if the specification is agreed upon and understood since the architecture continues to evolve while the system is being developed. Modules are likely to change over time and in some cases it may eventually be necessary to renegotiate and redefine some specifications as the assumptions on which the specifications were based often turn out to be wrong. This can be either due to a change of the environment but also because an envisaged technology turned out to be unfeasible for a specific purpose. In this case it would not be reasonable and not responsible to stick to the agreed upon specification. However, again, the perceived urgency to renegotiate and change specifications may vary across different groups as the problems to meet the specification may arise unilaterally. Such renegotiation of specifications may have a local impact on neighboring components only; however, in some cases local changes may also have an impact on the overall architecture.

For many complex technological systems, there is no apparent optimal architecture. Rather such systems are built upon specifications that are selected arbitrarily from several possible solutions, sometimes involving ambiguous and potentially contradicting specifications. In order to understand how such systems emerge, it is important to recognize that modular architectures are not defined in a deterministic way by processing readily available information according to a clearly specified set of design rules. On the contrary, modular architectures are constructed in an interactive process that involves individuals and groups who possess local, partially overlapping and potentially contradicting knowledge of interdependencies between different system components. In order to understand how a modular architecture emerges, it I necessary to adopt a perspective which is able to capture how individuals and groups interact with each other and the emerging technology and how the resulting controversies are resolved in a process of ongoing negotiation.

In this paper, we are looking at the development of a complex technological system that is one of a kind and involves several nascent technologies. The focus of our study is the early phase of design when the structures incorporated in the technology have not yet stabilized and the architecture is still emerging. Whereas prior literature has emphasized on the unambiguous understanding of architectures, we believe that it is important to consider conflicting interpretations in order to understand how architectures emerge. In particular, we study how the controversies resulting from conflicting interpretations get resolved as the different technological components and groups interact with each other.

RESEARCH DESIGN

A short note on the special characteristics of our research setting is in place. Our research investigates the development and construction of a complex technological system, a detector to measure subatomic particles in a high-energy physics (HEP) experiment. This research setting is particularly useful for our purpose for two reasons.

First, the absence of traditional organizational structures and hierarchy controls for alternative explanations as to how technological architectures emerge such as centralized planning and coordination. Coordination in HEP experiments has been described as arising "without centralized decision making, without a centralized control hierarchy with some individuals taking in information about the experiment, deciding what needs to be done, and issuing commands to other individuals who then perform the tasks. Experiments are nested hierarchies of lower level units in which the relationship between physicists and the technology is intense, particular and detailed" (Knorr-Cetina, 1995: 125).

Second, the ATLAS detector has not evolved from a dominant design but represents a one-of-a-kind technological system. This is true for both the architecture of the overall detector as well as the technologies and design of the various components comprising the complex technological system. This qualification of our research site guaranteed that the technological architecture of the experiment did not exist before and enabled us to study the emergence of architecture from the very beginning.

An exploration of the dynamics of emergent architectures requires us to take a process perspective (Mohr, 1982), which typically involves longitudinal analysis (Van de Ven, 1992). Consequently, in studying the emergence of the technological architecture of ATLAS, we relied upon the archival records that were created from the very beginning of the collaboration. In many cases, the same controversy has been discussed in different kinds of meetings such as working groups, review panels, or plenary meeting. Although the descriptions of the events were often accounted from different perspectives, they were mutually confirming, thereby generating confidence in the quality and depth of the data.

Overall, we analyzed several hundreds of pages of meeting minutes and correspondence covering the development period of ATLAS from 1991 to 2003. We also read the Letter of Intent, Technical Proposal and the Technical Design Reports of the ATLAS components. These detailed descriptions of the technological concepts and design consideration represented snapshots of the ATLAS architecture for the years 1992, 1994, and 1997.

We were fortunate to obtain the right to use the CERN archives, which provided us access to internal memos, technical reviews and other similar material. In addition, personal notes written by participants of the meetings offered insights on the micro details of the controversies that had not been apparent from the meeting minutes. The meeting minutes were complemented with conversations that were captured in electronic mailing list archives. This type of conversations was taking place electronically among the globally dispersed development groups. As a result, there is no lack of material about what happened, who was involved and the context within which the architecture emerged.

Finally, we had the opportunity to interview 20 scientists and engineers that had been involved in the collaboration throughout the genesis of ATLAS. The interviews were used for two reasons. First, respondents helped us to make sense of many of the technical details which were discussed in the technological controversies and also provided us with additional information about the social context which was sometimes to subtle to identified in the formal documents.

Second, the interviews were used to identify critical incidents in the emergence of the ATLAS architecture. Focusing on critical incidents reported by respondents "facilitates investigation of significant occurrences (events, incidents, processes or issues) identified by the respondent, the way they are managed, and the outcomes in terms of perceived effects" (Chell, 1998: 56). The objective is to gain understanding of the incident, taking into account cognitive, affective and behavioral elements.

An incident needs to be sufficiently complete in itself to permit inferences and predictions to be made about the process. To be critical the incident must occur in a situation where the purpose or intent seems fairly clear to the observer and where its consequences are sufficiently definite to leave little doubt concerning its effects (Flanagan, 1954: 327). Indeed, we discovered controversies in the emergence of ATLAS, such as the 'air core toroid decision' or the 'inner detector cooling review', which were recalled and identified as critical incident by a majority of respondents.

To analyze the data, we constructed a database that contained the events that had unfolded including both, the critical incidents identified by our respondents as well as other controversies that were captured in our data. The database included the sources from which we had identified these events and our interpretations (Miles & Huberman, 1984). We tracked changes in the technical architecture before, during and after each event.

Consistent with a process perspective, we considered each event as an important occurrence within a larger flow of events (Van de Ven, 1992), an approach which helped us generate a deep and consistent understanding of the unfolding process. As we were interested in the resolution of technological controversies, we focused on the critical events identified by our respondents and paid particular attention to the motivations and justifications of the actors involved. In particular, we examine how the controversies were handled and what the desired outcome was by looking at the discourse in the meetings and other documents before, during and after the critical incidents.

Whereas prior literature on modularity has focused on clearly understood specifications, we argue that it is important to consider the case of ambiguous specifications in order to understand how architectures emerge. In particular, we studied how controversies resulting from conflicting interpretations got resolved as the different technological components and groups were confronted with each other in meetings and review panels. A longitudinal observation of this discourse about technology finally resulted in insights about how the architecture was shaped. In other words, the discourse on the micro level helped to explain emergence that eventually is evident at the macro level of the technological architecture (Poole & DeSanctis, 1992).

As Knorr-Cetina (1999: 190) has noted in her work on HEP experiments, "discourse is (…) the activity through which much of the construction work is not only exhibited, but done." Taking this into account, we are confident that by examining the critical controversies that were captured in real time, we were able to observe the essence of the processes involved in the emergence of a complex technological system.

EMERGENCE OF ARCHITECTURE: ONGOING NEGOTIATION OF ORDER

ATLAS is one of five particle detector experiments being constructed at the Large Hadron Collider, a new particle accelerator at CERN in Switzerland. The detector will be 45 meters long and 25 meters in diameter, and it will weigh about 7000 tons when the construction is completed in 2007. The project involves roughly 2000 scientists and engineers from 151 institutions in 34 countries. By measuring phenomena that involve highly massive particles that were not measurable using the earlier lower-energy accelerators, ATLAS is expected to shed light on new theories of particle physics beyond the Standard Model, the most comprehensive model of particle interactions available today.

This ambitious project has its roots in UA1 and UA2, two very successful HEP experiments conducted at CERN in the 1980s. They became very famous for the discovery of the W and Z bosons, which led to the Nobel Prize for physics being awarded to CERN researchers in 1984. While the two experiments were still running, a core group of people was driving the development of future HEP detector concepts[2]. These protocollaborations – as they are called by Knorr-Cetina (1995, 1999) – met on HEP conferences where they proposed and discussed new concepts in workshops. In 1989, CERN started to fund groups of institutes for conducting R&D and prototyping for future detector designs. Until the early 1990s, 4 rather independent collaboration clusters emerged from this funding initiative. Each collaboration cluster was consisting of multiple R&D projects and was built around a new detector concept.

A critical event happened in 1992 when CERN decided that it was not useful to continue with 4 experiments. The collaborations were asked to join forces to develop and construct 2 independent detectors. The independent groups were forced to think about how their design concepts could be merged into a new architecture. The various protocollaborations had already developed idiosyncratic technological paths: although the functional requirements for each HEP detector are essentially the same, each group has opted for different technological solutions for the components performing a comparable function. Each group has developed a specific architecture that specifies how the various components are aligned and interact with each other. And not only the technological concepts had to be merged, the configuration of institutes had to match as well in order to make a successful proposal.

In this phase, scientists experienced tensions between technological preferences and social constellations. However, two of the collaborations, EAGLE and ASCOT, seemed to

[2] This parallel development of upgrades and future generations of experiments is typical for HEP experiments since development and construction of a new detector requires 10 – 15 years of work. While ATLAS is being developed, some groups in the collaboration are already working on upgrades for the years 2015 and beyond.

have more similarities than the other groups: both opted for a more complex configuration consisting of two magnet systems compared to other detectors with a simpler configuration. Consequently, the two groups merged in order to form the ATLAS collaboration. Making this merger work, however, required a compromise in terms of overall detector design and technology choice.

The creation of socio-technical controversies

Such as prior modern particle detectors, EAGLE and ASCOT included four key components: inner tracker, calorimeter system, muon spectrometer, and magnet system. While a general purpose detector such as ATLAS requires these different subsystems, each specialized on measuring a specific type of particles, the two collaborations largely varied in the emphasis they placed on the different components. While the core features of the EAGLE concept were a very elaborate calorimeter (a hybrid calorimeter consisting of two independent subsystems) combined with a powerful inner tracker, the ASCOT architecture was dominated by the superconducting magnet system required for its powerful muon spectrometer. What can explain these different outcomes given identical goals and physics specifications?

One of the most important reasons for this indeterminacy of the architecture is the fundamental uncertainty with respect to potential technologies. For each of the four components of the detector, several different technologies were considered as potential candidates for the emerging design. However, most options required further development to prove that they meet the requirements of the challenging environment. In this context of uncertainty, preferences and prior experience of EAGLE and ASCOT played an important role. Each group concentrated on components where they had relatively more experience. ASCOT's strength was muon detection. The resulting design was constrained by the characteristics of the muon system designed around a very large superconducting magnet system. EAGLE, on the other hand, was spearheaded by a group of people with a strong background in calorimetry from prior experiments and already in the 1980s came up with the concept of a hybrid calorimeter using independent subsystems as electromagnetic and hadronic calorimeter. Smaller groups, which had strong interest in inner tracker systems using semiconductor technologies, complemented this community.

It was clear that scientists of both EAGLE and ASCOT were closely attached to their proposals. Therefore, the collaborations agreed that the competing options "…can and should be analyzed in a scientific (and not emotional) way"[3]. Consequently, controversies were of rather technical nature. Beyond issues of performance, the competing technologies were evaluated in terms of cost, construction time and technical risks involved as well as impact on other subsystems. Working groups identified areas of consensus and tried to find common solutions. A series of the concept optimization meetings involving the various working groups created mutual knowledge of the strengths and weaknesses of each approach. Based on these insights, representatives of EAGLE and ASCOT started to negotiate how the technologies as well as the social groups could be integrated into a single experiment.

Controversies started to become greater when decisions were to be made for and against specific technologies. In this phase of configuration, the discussions were not at the purely technical level anymore but often involved political arguments too. Particularly problematic was that, by voting against a technology, the experiment might offend its proponents and lose critical knowledge, workforce and funding required for building the detector. Since a large experiment such as ATLAS required all these resources, delicate negotiation had to be performed by the groups whose technology was competing with others.

At this point in time, it became apparent that, besides technological controversies, the involvement of the competing social groups was shaping the architecture of ATLAS. The different compositions of the two collaborations played a particular role in this context. While ASCOT was dominated by two large French and German research institutes, EAGLE was a

[3] The "emotional way" is referring to political controversies going on between competing groups at that time.

numerically even larger community but it consisted to a great extent of small to medium size groups. Whereas big national laboratories have a tendency to work on large systems such as muon spectrometry, smaller institutes are often specialized on areas that do not require large investments in infrastructure, such as chip design for the inner tracker. These smaller institutes that had a strong stake in the EAGLE inner detector had the impression that the large institutes tried to force ASCOT's superconducting toroid on the collaboration and that it would dominate the design at the expense of their technologies.

In the course of the negotiation of a common concept, it turned out that what seemed to be a common understanding of the two collaborations was actually based on two fundamentally different philosophies. Both EAGLE and ASCOT proposed a large toroidal magnet to allow for an independent measuring capability of muons[4] and other highly energetic particles. Such a design provided an independent measurement of muons in addition to the observation of the inner tracker, an essential criterion for the EAGLE group: they argued that two independent measurements would increase the reliability of the experiment. The ASCOT motivation for a separate muon measurement was based on the expectation that muon detection in the inner detector was likely to fail due to the rough conditions of the LHC – its unprecedented energy and extremely high rate of collisions. The energy that gets released in the collisions is 700 % higher than in previous experiments and creates extremely high radiation. Scientists and engineers claimed that there was a substantial risk of failure for the inner detector. A stand-alone capability of the muon spectrometer provided redundancy for this scenario of a damaged inner detector.

The two groups had a different perception of this risk based on their assumptions. EAGLE expected to operate in low to moderate energy levels in the first phase of the experiment since it usually takes several years to run accelerators at their full potential. Therefore, the group emphasized a balanced detector concept featuring an elaborate inner

[4] These particles were considered as very important for the experiment since theories predict that they are related to the Higgs mechanism the physicists need to discover in order to explain physics beyond the standard model.

tracker for optimal particle detection at lower energy levels. The inner detector would sustain this first phase and could be upgraded at a later stage if necessary. ASCOT, on the other hand, expected the accelerator to run at high energy levels rather soon. In their scenario, the risk of damage to the inner detector was much higher. This explains why ASCOT optimized the overall detector layout for the muon spectrometer while the inner tracker was a relatively simple device since it was expected to fail in high energy levels in any case. This example illustrates how the different socio-cognitive bases created a collaborative challenge. Specifically, individuals from EAGLE and ASCOT were driven by different sets of beliefs and evaluation routines. While convergence of ideas and practices within each collaboration facilitated exchange of information and ideas, it made understanding and communication across the boundaries of the collaboration much more difficult because each group was locked into their own though worlds.

Unfolding of the architecture

In the Letter of Intent signed in October 1992, the 88 institutions already involved in the collaboration explicitly agreed on a shared schema; it defined the modular architecture of the new ATLAS detector, specified key properties and functionalities of its components and how these would interact with each other. While this overall detector concept made the basic design consideration explicit, the choice of the specific technologies for the different components was left open. The R&D projects conducted by participating institutes since the late 1980s resulted in various technologies which potentially could be used in parts of the detectors. To develop multiple technological options was considered necessary in the early phase since it was unclear which new technologies were capable of operating in the harsh radiation conditions of the LHC.

The Letter of Intent listed several competing technological options for each subsystem. In most cases a preferred technology was selected as "baseline design" with alternative options that would be chosen if the baseline turned out to be unfeasible. However, in some cases, the different R&D projects were even continued when the feasibility of the baseline seemed guaranteed in order to postpone decisions on technology choices. The community was very reluctant to make technology decisions as there was a constant fear to lose institutions who were favoring alternative technologies. Losing groups was considered a problem since the intellectual and manpower as well as the funding contributed every single institution is critical to the success of the overall collaboration. The community hoped that by continuing the R&D projects and testing prototypes it would become obvious that one technology is superior and that proponents of alternative solutions would be locked into the community by that time (cf. Knorr-Cetina, 1995).

Part of this avoidance and deferral of decision-making was to set up review panels in which the various technologies were evaluated. The review panels did not have any mandates to make a technology choice but to investigate the potential of the competing technologies and make a recommendation to the collaboration board consisting of all ATLAS institute members. The individuals serving on the review panels were nominated by the competing groups and usually were senior scientists, highly respected in the HEP community. The aim of these panels was to mute politics and let the obvious decisions unfold.

The procedure of the review panels can be imagined "similar to tribunal" as a physicist involved described:

> "The panellists who are chosen by the management are the judges and you have two parts against each other. In our case, for the hadronic part [of the calorimeter], we were against the liquid argon people and we were looking for problems of 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. And then, at the next meeting, there would be the answers to these questions. "

The inquisitive questioning in the review panels was challenging for both the opponent and the defending technology. Many of the assumptions, which were taken for granted by the proponents of a certain technology, all of a sudden were put into question by members of the review panel or by representatives of competing technologies. In order to make a convincing case, each group had to be prepared to justify those assumptions and provide evidence such as simulation studies, testing of prototypes, or support their claims with references by external experts or analytical reasoning. While the scientists were putting a lot of effort into justifying their technologies, it was even more demanding to identify potential shortcomings of competing technologies and ask those nasty questions since this required a deep understanding of the concepts proposed by the competing groups.

To some extent, the competing technologies were interchangeable without affecting the overall architecture of ATLAS. Some solutions, however, required an adaptation of the architecture and resulted in a certain development path for the overall detector. Therefore, mastering question and answer procedure in the review panel not only required knowledge of the component the groups were competing for. For example, during the review panel for the calorimeter it was essential for the groups to understand what the proposed specifications meant for neighbouring systems such as the inner detector. As an ATLAS scientist explained:

> "Liquid Argon calorimeter has a lot of advantages and so on, but it is a very expensive technology, it has a lot of channels and so on and so forth. This is kind of forcing you to a small radius and therefore has certain implications on the Inner Detector."

The Liquid Argon calorimeter group prepared to defend against this possible critique and worked on solutions for this conflict between the two interdependent subsystems. Specifically, the solenoid magnet, a part required for the inner detector to identify charged particles, was integrated into the cryostat of the calorimeter. This way, material and space constraints for the inner detector were mitigated, and the Liquid Argon calorimeter group could justify the viability of their proposal.

The questioning and answering forced competing groups to get a deep understanding of their technologies but also developed very good representation of the interdependencies with other systems. They had to have an understanding what important interdependencies were with other systems and, even more important, why those interdependencies were critical for the other system.

Scientists were coming from different contexts, and therefore, had different knowledge and experience with certain technologies. As a result, they also had varying judgements as to which technological option would be most appropriate for a particular component of ATLAS. Specifically, scientists tended to have a lot of confidence in their technologies since they had worked on them for a long time. From their experience, they knew what potentially could go wrong and how to fix it. Because of this familiarity, they were convinced that their proposal was preferable in terms of feasibility and reliability, risk and cost. The confrontation with the perspectives of competing or related systems urged them to scrutinize their own proposal and rethink some of the assumptions they had taken for granted before.

The process of confrontation with other perspectives was facilitated by the eagerness of the competing groups to justify the superiority of their proposal and was mediated by very experienced panel members who often had a strong background in related ATLAS components. From this review panel process, originally established to resolve technological controversies, an interlaced knowledge structure emerged which provided scientists of one system with a useful understanding of how related components, and ATLAS as a whole work.

The interlaced knowledge emerging from the review panels was not for the benefit of the people participating in the panel only but it was shared with the collaboration as a whole. Part of the idea of establishing impartial review panels was that the panel would be able to convince people and to build a consensus. Although there were some private meetings of the panels, in most of them a group of scientists from each detector was present and some of the meetings were completely public. In any case, the experiment was marked by a constant discourse and communication about itself, made up of threads of talks, emails, meeting presentations and transparency exchanges. As Knorr-Cetina (1995) described it, HEP experiments are “mapped into a fine grid of discourse spaces created by intersections between participants”. This narrating and accounting of what the group has been up to and has experienced in its dealings with equipment, data sets, physics calculations and the like makes the collaboration’s interlaced knowledge both collective and dispersed.

The creation of interlaced knowledge throughout the collaboration combined with the deferral of decision-making was characteristic for selecting the relevant options and regulating work at ATLAS. Rather than arbitrarily making decisions, the architecture of ATLAS was left to emerge. What happened routinely is that the system was nudged into a particular direction in a process involving review panels and other forums for discourse. Through the constant unraveling of the features of technical objects, of their details and composition, proponents of competing technologies tried to create a common understanding throughout the collaboration why their design was the solution in terms of performance, cost and risk for the overall detector. Decisions in ATLAS were left to emerge as the obvious, the reasonable or the unavoidable thing to do. As a senior scientist involved in the muon spectrometer expressed, “when decision-making was on an item on the agenda, this often meant that something which was already agreed upon and clear for everyone in the collaboration was made plausible and formally approved”.

The emergence of obvious decisions from the conversations of the review panels had the result that political decision making was not very present in ATLAS. Of course, politics were not completely eliminated but it was muted and waited out. This can be illustrated with the fact that the collaboration did not lose any members, even if their proposal was turned down. If opponents were able to make a convincing argument, scientists were usually willing to give in and suggested that the better design should be implemented.

Changes of architecture as design emerges

The properties of many technologies used for ATLAS components could be estimated only at best since there had been hardly any experience. Their performance and impact on other parts of ATLAS could be determined in educated guesses and complex simulations only. As a consequence, it was extremely difficult to determine an overall architecture that would remain stable during the development of the overall experiment. Many specifications, therefore, remained ambiguous and were of preliminary nature only.

A specification of ATLAS that contributed a considerable amount of uncertainty was the size of the superconducting air core toroid and the related muon spectrometer. Originally designed to consist of 12 coils and an inner radius of 5 meters, the cost and risk involved in this gigantic magnet system turned out to be even higher than expected. While scientists were already working on the designs of other ATLAS components based on the original specifications, the ATLAS toroid was scaled down to have 8 coils with an inner radius of only 4,3 meters only, a change of specifications which triggered several controversies with other components.

While the muon spectrometer was benefiting from the downscaling of the magnet system – a smaller radius means that the muon spectrometer has to cover a smaller area, and thus produce fewer expensive muon chambers – the components inside the toroid would have preferred a larger radius. The calorimeter in turn also had to adapt its inner radius since it requires a minimum thickness to effectively absorb and measure the energy of particles. But particularly problematic became the situation of the inner detector which had already suffered from confined space inside the other components. This further reduction in size had as result that the space envisaged for the inner detector electronics was not available anymore. Therefore, they had to be put somewhere else, which required other components to sacrifice space to host the inner detector electronics:

> "What I negotiated with the muons was a gap at the end of the barrel calorimeter, so the services came out and went straight out, or 50 percent of them did, between the inner layer and the second layer of the muon chambers, where there is more space to put a panel. And that has had an impact on the muons, they had to have a gap. We also have to cool the inner detector services in that region so that it does not warm up the muon electronics and the muon chambers themselves which is quite critical."

Moreover, the cables that were running to the inner detector electronics from their millions of readouts introduced additional material and also created heat, both negative side-effects for the Liquid Argon calorimeter. To change the location of the electronics was not trivial: the cables were routed next to power supplies and strong magnetic fields. To protect the electronic cables from picking up noise it was required to introduce effective shielding, which in turn required additional space and material.

Eventually, the need for more space inside the inner detector triggered small changes in the specification of the ATLAS architecture. For example, the gap between the barrel and the endcap calorimeter was extended twice in order to allow for additional space for shielded cables and cooling. While services such as cooling and cables were less of a problem for peripheral components since they could be routed outside the detector, services for the inner detector were invading the space of both calorimeter and muon spectrometer. However, since this increased need for services was partly due to constraints imposed by the outside detectors, it was not too difficult for the inner detector group to justify the extra space for cables and cooling.

Besides space limitations, the restriction to use as little material as possible created controversies between the inner detector and surrounding components. The inner detector and the calorimeter, for example, had conflicting requirements: while it was clear from the perspective of the inner detector engineers that additional material for cooling and shielding was required, the calorimeter group would have preferred to have even less material in front of the calorimeter. The amount of material used in the inner detector was causing reactions of particles before they could enter the calorimeter, and thus could cause a bias in the measurement of the calorimeter. The calorimeter group, therefore, tried to lobby for an inner detector design with fewer components.

Specifically, there was a discussion that the inner detector should have been built without a transition radiation tracker (TRT). However, the inner detector scientist could demonstrate in complex simulations that the TRT was crucial for achieving the physics performance specified for ATLAS and convinced the collaboration in a review panel that the amount of material thus could be justified. The collaboration found that it would be better to adapt the design of the calorimeter instead to account for the increased amount of material. The design of the Liquid Argon calorimeter was changed to introduce an additional measurement layer for the detection of electron and photon showers caused by material before the calorimeter.

The role of justification and creation of interlaced knowledge becomes more important as the architecture is forced to change itself. It enabled adaptive coordination when problems occurred which could not be solved locally inside a module but required support and compromises across various components. Interlaced knowledge enabled scientists of different components to interrelate heedfully and anticipate interference with other parts of ATLAS.

For example, it was the review of the cooling system of the inner detector that identified potential risks and technological problems inherent in the cooling system proposed by the inner detector. The cooling system was based on binary ice, a coolant consisting of ice crystals in a cooling liquid which was pumped through a complex system of pipes. Scientists of other components, such as the Liquid Argon calorimeter, who tried to avoid large amount of material of binary ice cooling pointed to problems related to water based cooling systems: water leaks in the cooling system could destroy parts of the detector. Interestingly, although this risk was highest for the inner detector itself, it was not recognized as a threat. Like all other ATLAS components, the inner detector was designed by scientists and engineers having local knowledge and incentives. Specifically, the inner detector engineers were concerned about getting the heat out of the densely packed inner detector and were focusing on the superior performance of binary ice cooling only.

Interestingly, other groups than the inner detector designers perceived the risk of water leaks more clearly. What seemed to be the better choice from the inner detector perspective was perceived to be worse from the point of view of the calorimeter. Scientists and engineers working on other components were less emphasizing the cooling performance of binary ice but were also considering negative aspects. In particular, they were dissatisfied with the amount of material introduced by the binary ice cooling which provided their motivation to find arguments against the binary ice cooling. Their different perspective enabled them to draw attention to the risks involved in this technology and propose an evaporative cooling system. The resolution of this disagreement resulted in a design that not only used less material, but also minimized the risk of water leaks in the inner detector.

The complex task of designing the inner detector given to the difficult space and material constraints caused further changes in the ATLAS architecture. The pressure to use more efficient electronics forced the inner detector group to change the design of the pixel system and to adopt a new type of semiconductors. The controversies caused by this change were mainly driven by conflicting schedules of the detector development. Since this technology was new, further testing and R&D work was required which resulted a delay of the pixel system of the inner detector. An inner detector scientist explained:

> "We thought having three systems, the pixels, the silicon strip detector and TRT, to have them all ready at the same time, it is going to be very difficult because the pixel changed… Then we changed the system: we put the silicon strip detector together with the TRT and then also put in the endcap. We then put in the pixel system in the end. This means that the schedule of the pixel system was decoupled from the schedule of the two other systems."

This example illustrates that conflicting schedules caused that the architecture of the systems was changed to decouple elements from the inner detector and integrate them in other components of ATLAS. In order to let this change happen, the agreed upon interfaces had to be renegotiated at a very late stage of building ATLAS. A change at this stage was only possible with the support of other groups who understood the problematic and potential of a delay of the inner detector development. This awareness of the progress and status of the various components was created by the reports and justifications in the numerous review panels, working groups and plenary meetings that connect the different groups involved in ATLAS.

What we saw in ATLAS was that this negotiation did not happen in the early phase only; rather, we observed a process of ongoing negotiations, distributed ordering across the collaboration throughout the development and construction phase. Radii of the different detectors were changed, lengths of the barrel, as well as allowance of material in each subsystem. As it turned out that one of the negotiated interfaces was unfeasible or insisting on the specification would exact an unreasonable burden on one of the systems, the collaboration was open to renegotiate the previously agreed-upon boundaries.

While this observation suggests that coordination is not embedded into the pre-specified system architecture as proposed by the modularity literature, the case of ATLAS demonstrates that decentralized coordination exists in the collaboration nevertheless. However, the need for managerial intervention is not embedded in a clearly specified, stable modular architecture. Rather, self-coordination of the collaboration is enabled by interlaced knowledge allowing for mutual adjustments of the interdependent actors and groups involved in the ATLAS collaboration.

DISCUSSION

Our longitudinal analysis of the emergence of ATLAS suggest that, in the case of the development of complex technological systems, architectures are neither as clear and understood, nor as stable as it has been assumed by the modularity literature so far.

Consequently, the proposition that the information structure embedded in a modular architecture enables coordination of decentralized R&D work in complex innovation systems (Baldwin et al., 2003; Sanchez et al., 1996) needs to be reconsidered.

Although a very rudimentary architecture of the experiment was created in the very beginning of ATLAS, this structure was far from being clear and understood by all participants involved. On the contrary, the architecture has emerged over time; it has experienced significant changes during the early development phase of six years. Almost on a daily base, geometrical and functional boundaries of interdependent components were challenged, and renegotiations of specifications continued even in the construction phase.

The coordination of the dispersed development and construction work of the more than 140 institutes involved thus could not be provided by simply pre-specifying the system boundaries and interfaces. Moreover, since the ATLAS collaboration did not have any central decision making or governance body, there needs to be an alternative explanation of how the decentralized activities of 2'000 scientists and engineers were coordinated.

A striking phenomenon throughout the ATLAS case was the ongoing negotiations and the importance of create compromise and consensus. What we observed in ATLAS was that notorious political actors were being tamed into the collaboration. The architecture of the experiment to be built was not defined only by some actors winning out, but through confrontation of the components of the experiment, the groups and physicists who had a big stake and interest in these matters. These confrontations involved technology as well as groups, and technical questions which implicate social resources and tactics. By the public summation of the knowledge everybody had created, the collaboration papered over political strategies, and a negotiated order emerged from the strategy of letting the "obvious" decisions unfold.

Creation of interlaced knowledge

A process of justification accompanied negotiations involved in specifying the ATLAS architecture: in negotiations it was essential to explain ‘why’ choosing a particular design or technology was considered the best solution. Justification of proposals was important both during the initial specification of the architecture, but even more so in controversies about adaptations which were required as the architecture emerged. If a group wanted to change the technological path once chosen, they had to mobilize supporting groups throughout the collaboration. In order to convince the collaboration that the new proposal would result in a better outcome, it was essential to present reasoning and evidence to justify their claim.

Due to the openness of the ATLAS collaboration and the involvement of individuals from multiple groups, the justification got exposed to scrutiny and was always available for inspection by many. In this interactive process, many different points of view and component-specific perspectives were articulated. While negotiations can be seen as reaching a political compromise, this process also contributes to building of common ground on the knowledge level. What we observed in ATLAS was that this ongoing negotiation resulted in a reconfiguration and recontextualization of the available knowledge. By the time controversies got resolved, the competing groups had developed a shared understanding of each other’s components and what implications the technological architecture had. As the various groups involved created such overlaps of their local knowledge bases, they developed what we call “interlaced knowledge” across interdependent components. This observation is in line with current research on epistemology which argues that justification is a key process for creating organizational knowledge (Nonaka, 1994). Justification essentially decides whether claims of individual groups are rejected by others in the organization, or believed to be valid and incorporated into their knowledge bases (von Krogh & Grand, 2000).

By choosing the term “interlaced knowledge”, we differentiate our concept from others such as “common knowledge” (Grant, 1996). We would like to emphasize that, although overlaps between the various knowledge bases exists, this does not imply that there is a common knowledge base shared by all members in the collaboration. On the contrary, Interlaced Knowledge was created on various levels reflecting to the fine grid of discourses in the ATLAS collaboration. The content and degree of sharedness varied throughout the collaboration, depending on the controversies taking place in local working groups, review panels across different components, or plenary meetings involving the collaboration at large.

Interlaced knowledge at ATLAS comprised various levels of detail. On the overall collaboration level, this usually included very basic knowledge about the progress and status of detector components. This knowledge gave individuals an appropriate representation of current issues, what the challenges of related components were, and where controversies were likely to emerge. Such discourse was taking place in the plenary meetings, thereby playing an important role as attention-driving mechanism (Brown & Eisenhardt, 1997). Regularly scheduled updates created an internal rhythm and an internal time pacing in the development process. On this level, interlaced knowledge helped groups of interdependent systems to synchronize the work and keep track of changes in the technological system.

The discourse in working groups and review panels resulted in more detailed knowledge. Controversies usually involved justification across boundaries, for example, during negotiations of interface specifications between two interdependent components. The interlaced knowledge emerging from this level provided the different groups with a deeper understanding of each other's context and requirements. This, in turn, enabled developers of interdependent components to heedfully interrelate to each other when unforeseen changes occurred. The controversy about the material budget of the inner detector showed, for example, how the Liquid Argon calorimeter anticipated the difficulties of the inner detector to build a TRT within the material constraints. The calorimeter's decision to introduce an additional layer of detectors enabled the inner detector to develop a working TRT exceeding the amount of material originally specified.

Interlaced knowledge emerging from technical controversies was articulated and shared in a variety of documentation such as meeting minutes, technical reviews, and design reports. Those documents were presented on many occasions and circulated in electronic mailing lists. The overall interaction was potentially all-to-all, and every scientist or engineer enjoyed free access to all information in ATLAS, thus the knowledge percolated throughout the ATLAS collaboration at large.

Coordination through interlaced knowledge

Our study suggests that the modular architecture was not able to determine the decentralized development of the various components. However, interlaced knowledge emerging from controversies around the technological architecture facilitated the coordination of the development and construction of ATLAS. It is important to note that, although the overlaps of the various groups' knowledge bases often coincided with interfaces explicitly specified in the ATLAS architecture, interlaced knowledge base is more than just the "information structure" (Sanchez et al., 1996) embedded into pre-specified interfaces. Interlaced knowledge is much richer, by contrast, and helps to coordinate complex innovation systems even as they emerge. In ATLAS, interlaced knowledge provided redundancy through the overlapping knowledge, thus enabling the collaboration to employ a "shared division of labor" (Nonaka & Takeuchi, 1995). This type of redundancy is not concerned with inefficiency – as it is the usual connotation of this term – it rather is an advantage when confronting the uncertainty inherent in complex technological systems. A shared division of labor provides additional degrees of freedom; it allows to accommodate uncertainty and to adapt the system as new contingencies emerge.

While the rigid information structure embedded in a modular architecture is likely to lock a system onto a pre-specified development path, interlaced knowledge enables collaborators to renegotiate specifications and break out of the path if necessary. In that sense, the interlaced knowledge is generative in nature: it enables developers to respond to emerging problems by creating novel solutions rather than just adhering to pre-determined specifications. We observed that groups in ATLAS were able to effectively coordinate their actions even in presence of uncertainty, when the architecture was not stable, and specifications were ambiguous among the groups involved. As unforeseen events occurred, for example the delay of the ATLAS pixel detector, interlaced knowledge and the shared division of labor enabled the component to find a workaround in collaboration with interdependent groups. Some solutions even involved the renegotiation of previously specified interfaces resulting in a change of the architecture, as it was the case with the integration of the pixel detector into the endcap.

Modular architecture as boundary infrastructure

Similar to research on modularity (Baldwin et al., 2000; Sanchez et al., 1996; Ulrich & Eppinger, 1995), we too found that the structures facilitating coordination became embedded into the ATLAS architecture. However, whereas the modularity literature emphasizes the importance of blackboxing and information hiding (Parnas, 1972), we suggest that the knowledge *why* particular specifications were introduced needs to be preserved. The agreed upon specifications are merely the tip of the iceberg of interlaced knowledge which emerged from ongoing process of negotiation and ordering. While the mere specification – a very reduced form of knowledge – may be sufficient to coordinate a stable system, the origins of this knowledge as well as important interdependencies are forgotten and overlooked. As a consequence, pre-specified interfaces are not sufficient to coordinate the development of emergent architectures; it may even result in dysfunctional actions among the groups involved.

The specification of the architecture has an important role nevertheless. However, rather than an information structure determining the development of interrelated components, it has the role of a boundary infrastructure (Bowker & Stars, 1999) connecting heterogeneous components. As a boundary infrastructure, a modular architecture can be interpreted from different perspectives, yet its meaning is sufficiently similar to enables the different groups to structure their conversations when negotiating the design (Star et al., 1989). Different groups can use the architecture as a "means of representing, learning about, and transforming knowledge to resolve the consequences that exist at a given boundary" (Carlile, 2002: 442).

The predominant view on modularity underestimates the role of boundary architecture by assuming that specifications components and interfaces will be clear and understood. And indeed, recent research suggests that justification across boundaries decreases as standards are taken-for-granted (Green, 2004). This discouragement of justification is in stark contrast to findings in organization and management research that investigates how organizations can organize in emergent environments. For example the management of high reliability organizations relies heavily on justification in order to create the richer representations of the overall system in order to heedfully interrelate if new events emerge in organizations (Weick & Roberts, 1993; Weick & Sutcliffe, 2001).

Similarly, studies on organizational learning found that in the condition of uncertainty, organizations create more knowledge about the complex interdependencies when people from different perspectives engage in sense-making process and work on a converging representation of how the system might look like (March, Sproull, & Tamuz, 1991). The openness to a variety of possibly irrelevant interpretations is often more valuable than a clear prior model and unambiguous objective (March, 1987).

Of course, the creation of interlaced knowledge from multiple perspectives appears to be a luxury if such flexibility is not required. From an information processing perspective, coordination is more efficient if people can rely on taken-for-granted specifications. This is the perspective of research on modularity: a pre-specified architecture which has evolved historically or by convention allows for distributed information processing (Baldwin et al., 2003). The architecture provides embedded coordination by splitting up the complex task into sub-problems and integrating the parts into an overall solution. This perspective seems appropriate if a mature architecture already exists. However, other mechanisms need to come into play when uncertainty is involved and architectures are emergent.

If specifications are perceived as boundary objects that can be interpreted from different perspectives, modular architectures can enable coordination in such dynamic contexts. The interlaced knowledge, which is created in the ongoing negotiations across boundaries, provides a deeper understanding and appreciation of other components' requirements. Groups are enabled to anticipate latent interdependencies and to heedfully interrelate to interdependent systems when new events in the innovation system occur. This enables the dispersed groups to effectively coordinate their actions even as the architecture and specifications are still emerging.

CONCLUSION

The case of ATLAS illustrated that the emergence of architectures can involve ambiguous specifications that are interpreted differently by various groups involved. Of particular interest was how controversies resulting from conflicting interpretations got resolved as the different technological components and groups were confronted with each other. Our analysis suggests that justification across boundaries creates interlaced knowledge, a structure of local knowledge bases that overlap where interdependencies between components exist. This interlaced knowledge provides a deeper understanding and appreciation of other components' requirements that in turn allows the multiple groups to anticipate latent interdependencies and to heedfully interrelate to each other. As a consequence, the dispersed groups are enabled to effectively coordinate their actions even as the architecture and specifications are still emerging.

We have explored the ATLAS case to sketch out the process involved in the emergence of technological architectures. Through prior negotiation and codification of modular architectures, very complex interdependencies can be hidden behind a small number of pre-specified interfaces. Such information hiding supports an ongoing division of cognitive labor between the different component developers. However, whereas designers in established technological systems can take a common definition of interfaces for granted, the scientists and engineers in emergent systems such as ATLAS cannot build on standards and conventions which evolved historically. For emergent technological systems, architectures need to be created in first place.

REFERENCES

Alexander, C. 1964. ***Notes on the Synthesis of Form***. Cambridge, MA: Harvard University Press.

Anand, N., & Watson, M. 2004. Tournament Rituals in the Evolution of Fields: The Case of Grammy Awards. ***Academy of Management Journal***, 47: 59-80.

Aoki, M. 2001. ***Towards a Comparative Institutional Analysis***. Cambridge: MIT Press.

Baldwin, C. Y., & Clark, K. B. 2000. ***Design rules, Vol. 1: The power of modularity***. Cambridge: MIT Press.

Baldwin, C. Y., & Clark, K. B. 2003. Where Do Transactions Come From? A Perspective from Engineering Design. Cambridge, MA: Harvard Business School.

Baldwin, C. Y., & Clark, K. B. 2005. Between "Knowledge" and "the Economy": Notes on the Scientific Study of Designs. Cambridge, MA: Harvard Business School.

Barry, D., & Rerup, C. 2006. Going Mobile: Aesthetic Design Considerations from Calder and the Constructivists. ***Organization Science***, 17(2): 262-276.

Bowker, G. C., & Stars, S. L. 1999. ***Sorting things out: classification and its consequences***. Cambridge: MIT Press.

Brown, J. S., & Duguid, P. 2001. Knowledge and organization: A social-practice perspective. ***Organization Science***, 12(2): 198-213.

Brown, S. L., & Eisenhardt, K. M. 1997. The Art of Continuous Change: Linking Complexity Theory and Time-paced Evolution in Relentlessly Shifting Organizations. ***Administrative Science Quarterly***, 42: 1-34.

Carlile, P. R. 2002. A Pragmatic View of Knowledge and Boundaries: Boundary Objects in New Product Development. ***Organization Science***, 13: 442-455.

Carlile, P. R. 2004. Transferring, translating, and transforming: An integrative framework for managing knowledge across boundaries. ***Organization Science***, 15(5): 555-568.

CERN. 1994a. ATLAS Technical Proposal, Vol. CERN/LHCC 94-43. Geneva.

CERN. 1994b. The Compact Muon Solenoid Technical Proposal, Vol. CERN/LHCC 94-38. Geneva.

Chell, E. 1998. Critical Incident Technique. In G. Symon, & C. Cassell (Eds.), ***Qualitative Analysis in Organizational Research. A Practical Guide***: 51-72. London: Sage Publications.

Clark, K. B., & Fujimoto, T. 1990. The power of product integrity. ***Harvard Business Review***, 68(6): 107-118.

DeSanctis, G., & Poole, M. S. 1994. Capturing the Complexity in Advanced Technology Use: Adaptive Structuration Theory. ***organization Science***, 5(2): 121-147.

Dougherty, D. 1992. Interpretative barriers to successful product innovation in large firms. ***Organization Science***, 3(2): 179-202.

Eppinger, S. D., Whitney, D. E., Smith, R. P., & Gebala, D. A. 1994. A model-based method for organizing tasks in product development. ***Research in Engineering Design***, 6: 1-13.

Flanagan, J. C. 1954. The Critical Incident Technique. ***Psychological Bulletin***, 51(4): 327-358.

Galunic, D. C., & Eisenhardt, K. M. 2001. Architectural innovation and modular corporate forms. ***Academy of Management Journal***, 44: 1229-1249.

Garud, R., & Ahlstrom, D. 1997. Researchers' Roles in Negotiating the Institutional Fabric of Technologies. ***American Behavioral Scientist***, 40: 523-538.

Garud, R., & Kotha, S. 1994. Using the brain as metaphor to model flexible production systems. ***Academy of Management Review***, 19: 671-698.

Garud, R., & Kumaraswamy, A. 1995. Technological and organizational designs to achieve economies of substitution. ***Strategic Management Journal***, 16: 93-110.

Grant, R. M. 1996. Toward a knowledge-based theory of the firm. ***Strategic Management Journal***, 17: 109-122.

Green, S. E., Jr. 2004. A rhetorical theory of diffusion. ***Academy of Management Review***, 29: 653-669.

Henderson, R. M., & Clark, K. B. 1990. Architectural innovation: the reconfiguration of existing product technologies and the failure of established firms. ***Administrative Science Quarterly***, 35(9-30).

Hobday, M. 1998. Product complexity, innovation and industrial organisation. ***Research Policy***, 26(6): 689-710.

Hobday, M. 2000. The project-based organisation: an ideal form for managing complex products and systems? ***Research Policy***, 29(7-8): 871-893.

Hoegl, M., Weinkauf, K., & Gemuenden, H. G. 2004. Interteam coordination, project commitment, and teamwork in multiteam R&D projects: A longitudinal study. ***Organization Science***, 15(1): 38-55.

Knorr-Cetina, K. 1995. How Superorganisms Change: Consensus Formation and the Social Ontology of High-Energy Physics Experiments. ***Social Studies of Science***, 25: 119-147.

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

Langlois, R. N., & Roberts, J. 1995. Innovation, Networks, and Vertical Integration. ***Research Policy***, 24(4): 543-562.

Langlois, R. N., & Robertson, P. L. 2003. Networks and innovation in a modular system: the microcomputer and stereo component industries. In R. Garud, A. Kumaraswamy, & R. N. Langlois (Eds.), ***Managing in the modular age: architectures, networks, and organizations***. Oxford: Blackwell Publishers.

March, J. G. 1987. Ambiguity and Accounting: The Elusive Link between Information and Decision Making. ***AOS***, 12: 153-168.

March, J. G., Sproull, L. S., & Tamuz, M. 1991. Learning from Samples of One or Fewer. ***Organization Science***, 2: 1-13.

Mihm, J., Loch, C., & Huchzermeier, A. 2003. Problem-solving oscillations in complex engineering projects. ***Management Science***, 49(6): 733-750.

Miles, M. B., & Huberman, A. M. 1984. ***Analyzing qualitative data: a source book for new methods***. Beverly Hills: Sage.

Mohr, L. B. 1982. ***Explaining organizational behavior: the limits and possibilities of theory and research***. San Francisco: Jossey-Bass.

Nonaka, I. 1994. A dynamic theory of organizational knowledge creation. ***Organization Science***, 5: 14-37.

Nonaka, I., & Takeuchi, H. 1995. ***The knowledge-creating company***. Oxford: Oxford University Press.

O'Sullivan, A. 2003. Dispersed collaboration in a multi-firm, multi-team product-development project. ***Journal of Engineering and Technology Management***, 20(1-2): 93-116.

Orlikowski, W. J. 1992. The Duality of Technology: Rethinking the Concept of Technology in Organizations. ***Organization Science***, 3(3): 398-427.

Parnas, D. L. 1972. On the criteria to be used in decomposing systems into modules. ***Communications of the ACM***, 15: 1053-1058.

Poole, M. S., & DeSanctis, G. 1992. Microlevel Structuration in Computer-supported Group Decision Making. ***Human Communication Research***, 19(1): 5-49.

Postrel, S. 2002. Islands of shared knowledge: specialization and mutual understanding in problem-solving teams. ***Organization Science***, 13: 303-320.

Sanchez, R. 1995. Strategic Flexibility in Product Competition. ***Strategic Management Journal***, 16: 135-159.

Sanchez, R., & Mahoney, J. T. 1994. The modularity principle in product and organization design: achieving flexibility in the fusion of intended and emergent strategies in hypercompetitive product markets, ***Office of Research working paper***. University of Illinois at Urbana-Champaign.

Sanchez, R., & Mahoney, J. T. 1996. Modularity, flexibility, and knowledge management in product and organization design. ***Strategic Management Journal***, 17: 63-76.

Schilling, M. A. 2000. Toward a General Modular Systems Theory and Its Application to Interfirm Product Modularity. ***Academy of Management Review***, 25(2): 312-334.

Simon, H. A. 1962. The architecture of complexity. ***Proceedings of the American Philosophical Society***, 106: 467-482.

Simon, H. A. 1996. ***The Sciences of the Artificial***. Cambridge, MA: MIT Press.

Sosa, M. E., Eppinger, S. D., & Rowles, C. M. 2004. The misalignment of product architecture and organizational structure in complex product development. ***Management Science***, 50: 1674-1689.

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

Sydow, J., & Windeler, A. 1998. Organizing and evaluating interfirm networks - A structurationist perspective on network management and effectiveness. ***Organization Science***, 9(3): 265-284.

Sydow, J., Windeler, A., Schubert, C., & Moellring, G. 2007. Organizing networks for path creation and extension in semiconductor manufacturing technologies, ***Academy of Management Meeting***. Philadelphia.

Thompson, J. D. 1967. ***Organizations in action***. New York, NY: MacGraw-Hill.

Tushman, M. L., & Murmann, J. P. 2003. Dominant designs, technology cycles, and organizational outcomes. In R. Garud, A. Kumaraswamy, & R. N. Langlois (Eds.), ***Managing in the modular age: architectures, networks, and organizations***. Oxford: Blackwell Publishers.

Ulrich, K. 2003. The role of product architecture in the manufacturing firm. In R. Garud, A. Kumaraswamy, & R. N. Langlois (Eds.), ***Managing in the modular age: architectures, networks, and organizations***: 117-145. Oxford: Blackwell Publishers.

Ulrich, K., & Eppinger, S. D. 1995. ***Product design and development***. New York: McGraw-Hill.

Ulrich, K., & Eppinger, S. D. 2000. ***Product design and development***. New York: Irwin McGraw-Hill.

Van de Ven, A. H. 1992. Suggestions for Studying Strategy Process: A Research Note. ***Strategic Management Journal***, 13: 169-188.

von Krogh, G., & Grand, S. 2000. Justification in knowledge creation: open issues and research questions. In G. von Krogh, I. Nonaka, & T. Nishiguchi (Eds.), ***Knowledge creation: a source of value***. Houndmills: Macmillan Press.

Weick, K. E., & Roberts, K. H. 1993. Collective mind in organizations: heedful interrelating on flight decks. ***Administrative Science Quarterly***, 38: 357-381.

Weick, K. E., & Sutcliffe, K. M. 2001. ***Managing the Unexpected: Assuring High Performance in an Age of Complexity***. San Francisco: Jossey-Bass.

Wood, S. C., & Brown, G. S. 1998. Commercializing nascent technology: The case of laser diodes at Sony. ***Journal of Product Innovation Management***, 15(2): 167-183.