

# Architecture Alignment in a Large Government Organization: A Case Study\*

R.J. Wieringa, P.A.T. van Eck, and H.M. Blanken

Center for Telematics and Information Technology, University of Twente  
P.O. Box 217, 7500 AE Enschede, the Netherlands  
(roelw|vaneck|blanken)[@cs.utwente.nl](mailto:roelw|vaneck|blanken@cs.utwente.nl)

**Abstract.** In this paper we view IT architecture as the structures present in the entire information technology support used by an organization. Research into IT architecture either is of a strategic nature, yielding no operational guidelines for the practicing IT-architect, or it is part of software engineering, yielding no guidelines that related software architecture to the business environment. In this paper we report on a detailed case study of an operational IT architecture process, in which we investigated the relationship between IT architecture and business context. We analyze this process in terms of a conceptual framework for IT architecture presented earlier. The major findings are that in this case study there is a close relationship between IT architecture and the structure of the IT department, which makes IT architecture design also a problem of organizational design; that application architecture is designed by aligning applications to the business process structure; and that IT infrastructure architecture is designed by aligning it to technological trends rather than to business goals and problems. This paper reports about a single case study done in a series of case studies. We discuss the generalizability of the findings from this case study, and discuss topics for further research.

## 1 Introduction

Architecture alignment is the mutual adjustment of IT-architecture and business architecture. For large information-processing organizations such as banks, insurance companies and government organizations, continuous architecture alignment is crucial to survive. The IT departments of these organizations are faced with frequent mergers, which requires them to integrate several IT structures that are pairwise incompatible, and with frequent reorganizations, which necessitates adjustment of the IT structures to new organization structures. In this context IT architecture design is a complex problem that is never solved but must continually be attended to.

Our research can be positioned between Software Engineering and Business Science. Most current software engineering research [1,2,3] does not provide all

---

\* This research was partially sponsored by the Telematics Institute. See <http://www.telin.nl/NetworkedBusiness/Graal/ENindex.htm>

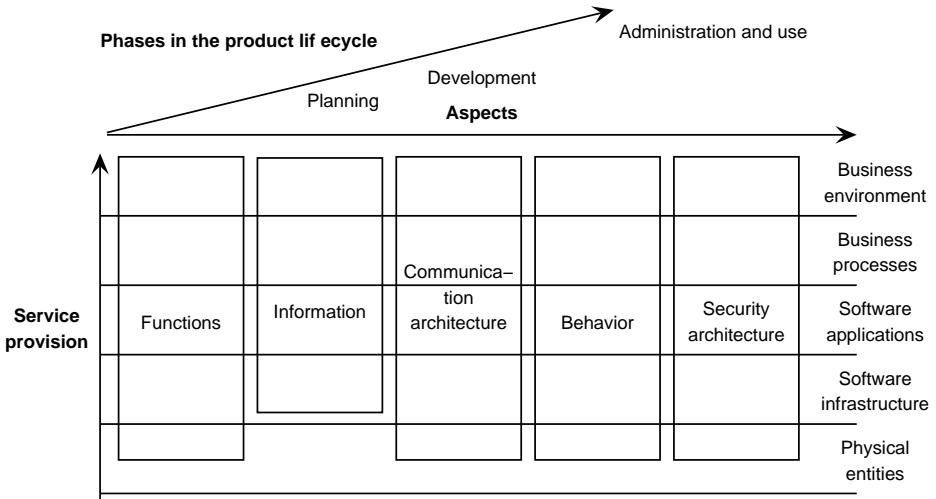
the techniques to deal with architecture alignment, because it focusses on the internal structure of one software system and does not include in its scope the business context of an ensemble of software systems that together realize all information processing needs of an organization. On the other hand, business scientists who study architecture alignment do not come up with operational design guidelines [4,5,6]. Opdahl makes a start with defining guidelines for representing application alignment [7]. However, typical business studies investigate the alignment of business objectives and IT objectives, or management issues such as e.g. the conditions of IT personnel management. In short, the practicing IT architect misses operational guidelines for relating IT architecture to business architecture.

The GRAAL project aims to identify operational guidelines for business-IT architecture alignment for large information-processing organizations. We reach this goal by doing in-depth case study research in these organizations. To be able to compare the analysis results in different organizations (and sometimes even in different parts of the same organization), we developed a conceptual framework for business and IT architectures, which we reported on earlier [8]. In the current paper we report on a case study of architecture alignment performed at a large government organization, that we will call BIGIT (Big Government IT organization). BIGIT has about 37000 users spread out over different locations in The Netherlands. BIGIT provides is part of a large government organization, that we will simply call BIG. BIGIT provides IT development and maintenance services for BIG. BIG is partitioned into a number of departments, one of which we will simply call D. We will focus our case study on the services provided by BIGIT for D.

We should state at the outset that we report on *case study research* [9]. This means that we have observed what happened in an IT department and analyzed our observations. Case study research should not be confused with experimental research. The goal of experimental research is to make inferences about a *population*. This is usually done by taking a sufficiently large random selection of cases from a population and generalize to the whole population by means of statistical methods. Case study research does not generalize to a population, but to *theory* [9]: every extra case study performed in which patterns predicted by a theory are found increases the extent to which we believe the theory is true. Case study research is also different from action research. In action research, a researcher actively tries to influence the outcome of the case at hand by performing actions, and evaluating whether the actions had the effects predicted by theory. In case study research, the case is left undisturbed. This paper reports on an exploratory case study, which is a case study that is not carried out starting from a set of hypotheses, but aims to find hypotheses to be validated in further case studies.

## 2 The GRAAL Framework

In order to analyse cases, we need a conceptual framework and vocabulary in terms of which to describe the cases. An analysis of a number of frameworks for information systems and software development, industrial product development and systems engineering led to a simple framework for software and business shown in figure 1 [10,11]. The core of the framework consists of a stack of layers.



**Fig. 1.** The GRAAL Framework. An additional dimension of our framework, not shown, concerns the refinement level at which we describe entities at each layer.

Each layer consists of entities (hardware, software, people, organizational units, customers). Entities at lower layers provide services to entities at higher layers. At BIGIT, we distinguish, from the bottom up, the layers shown in figure 1. The two lowest layers need explanation:

- Physical entities are, for example, computers, networks, buildings, furniture, etc. An entity is physical if it can be described using physical dimensions such as meters, kilograms, amperes, and seconds.
- The software infrastructure consists of operating systems, network software, middleware, DBMSs. A software entity is part of the infrastructure if it provides services of general use, i.e. not containing specific knowledge of the organization in which it functions.

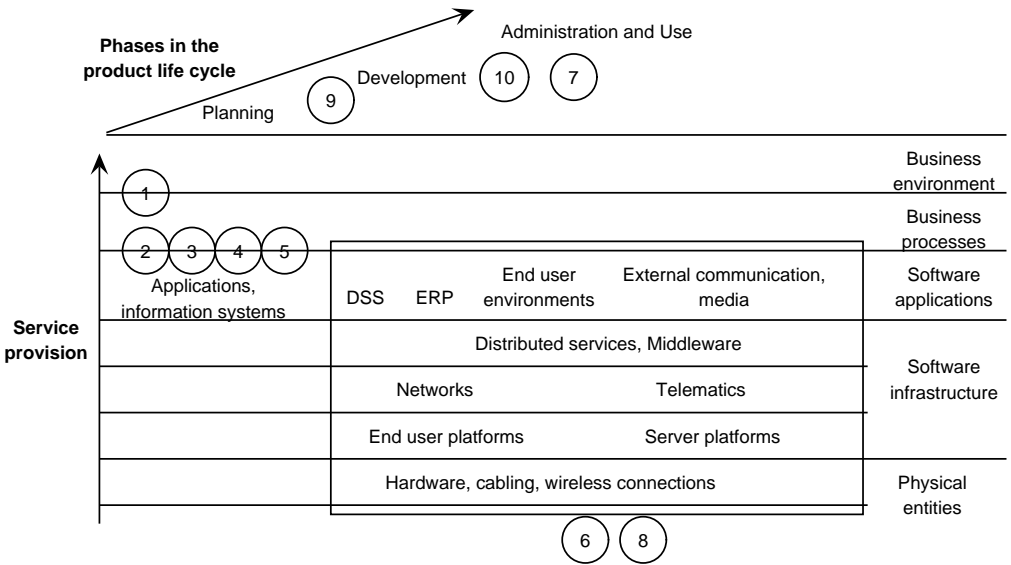
The entities at each layer have properties, which we can classify under *aspects*. Each entity has a *function* for its environment. (In this paper, we use the word

“function” as synonym of “service”.) All entities above the physical layer have an *information* aspect. For example, software manipulates symbols with a well-defined meaning, business processes manipulate data, and customers supply data to BIGIT. All entities *communicate* with other entities, they all have *behavior* over time, and in BIGIT, *security* is important at all layers. In other organizations, other aspects may be important too, such as reliability of interoperability.

Each entity has a *life cycle* that starts with a planning stage and ends with the management of the entity in use. This is complicated by the fact that many entities may exist in different versions, and each of these versions goes through a planning, development and deployment stage. At any point in time, several versions of an entity may exist in different stages of its life cycle.

Finally, a dimension not shown in the diagram concerns the *description* of entities. Each entity can be described by a combination of text and diagrams. These descriptions can be very abstract (little detail) or very refined (a lot of detail). For example, one can describe a business process abstractly by describing its purpose only, and one can describe it in a very refined way by detailing every transaction that is part of it. Space limitations prevent us from comparing this with other frameworks. Comparisons have been published elsewhere [11,8].

### 3 The Case



**Fig. 2.** Analyzed documents. The horizontal dimension has no meaning in terms of the framework of figure 1.

We will focus our case study on the services provided by BIGIT for one of the departments of BIG, called D. We have analyzed ten architecture documents of the IT systems of D, developed and maintained by BIGIT (figure 2). Documents 1 through 5 describe applications and information systems at various levels of detail. Documents 6 and 8 treat *infrastructure domains* in BIGIT, which are the domains of hardware, cabling and wireless, end user platforms, server platforms, etc. as shown in the diagram. Finally, documents 7, 9 and 10 deal with various aspects of the development process of software systems.

Based on these documents, we drafted an initial analysis of business-IT alignment at D and verified this with the IT architects of BIGIT working for D. After adjusting the analysis to feedback that we received, we presented the final result to the BIGIT architects. Here, we present this final result in anonymized form.

## 4 Findings

### 4.1 Applications and infrastructure

The first finding in our case study is that there is a clear separation between applications and infrastructure in terms of acquisition and maintenance. We adapted our framework to this by distinguishing these two layers in figure 1. A software system is an *application* if it provides services for a specific user group in BIG and contains knowledge that is specific for this user group. Applications usually play a role in an particular business process. They are acquired (built or bought) specifically for their user groups and for particular business processes. Failures are felt in this business process only.

A software system is part of the *infrastructure* if it provides services for all users in BIGIT and does not contain any knowledge about its particular context of use. Infrastructure is not acquired for any particular user group or any specific business process. Like the central heating system, infrastructure is expected to be available wherever and whenever it is needed. Failures are felt in the entire user population.

Figure 2 shows that infrastructure at BIGIT is a complex consisting of several layers, ranging from hardware and physical connection technology, to operating systems middleware, telecom services, decision support systems, ERP systems and office software available on all user stations. Not only the top layer of infrastructure is directly accessible to users. Users directly interact with hardware, operating systems and communication technology and in one way or another depend on all infrastructure elements. This puts a high premium on the quality of service of infrastructure.

Application development is event-driven. Usually, there is a change request from a user group, or a management directive to implement a new service. The application development process then generates a steady flow of documents, starting from the change request itself, to an initial requirements definition, top-level architecture design and so on to the finally delivered documents: Source code, executables, technical and user documentation.

Infrastructure development is time-triggered. It takes place periodically, for example yearly, and produces less documents. These documents are not oriented towards a service to be delivered towards users but towards business goals to be satisfied and the available technology to satisfy these goals. We return to this below.

Given this difference between applications and infrastructure, we can expect application architecture to be structured according to user groups and their business processes. Infrastructure is not oriented to any particular user group, and its architecture is structured according to technological domains. There is thus a tension in architecture alignment between applications and infrastructure. This will be confirmed as we continue our analysis.

## 4.2 Conway's Law

Any system whose design is too large to be developed by one person will be developed by a design group. Work will be partitioned into parts that will be allocated to different members of the group. Conway [12] points out that the design group will therefore produce a system of which the architecture is a copy of the communication structure of the design group. Brooks [13, page 111] coins this *Conway's law*. We see Conway's law illustrated at BIGIT.

- The separation between applications and infrastructure is institutionalized by two different departments of BIGIT that deal with these two parts of the IT of BIG.
- The application development part of BIGIT was partitioned into application development departments, one for each department of BIG. So the set of application development departments was isomorphic to the set of user departments.
- The infrastructure development department is organized according to the infrastructure domains shown in figure 2.

The significance of this organization only becomes apparent by a restructuring of BIGIT that happened during the period of our case study research. BIGIT was restructured according to the waterfall model: In the restructured BIGIT, there is a department for requirements engineering, a department for global architecture design, a department for detailed design and implementation, one for system management, and one for user support. In addition, there are some supporting departments, such as Procurement and Finance. This reorganization caused a number of failures.

- BIG is a large organization, with 16 other departments besides department D. The old structure had the advantage that each department X of BIG had an obvious place to go to BIGIT, namely the application development department x of BIGIT that dealt with X. In x there was a lot of domain knowledge about the business of X, knowledge that builds up over years of interaction between X and their IT support department x. In the current

structure, this knowledge is dispersed over BIGIT and for X, the requirements engineering department of BIGIT has no obvious person to go to state their problem. This causes more unnecessary misunderstandings and additional communication overhead between BIGIT and its customers.

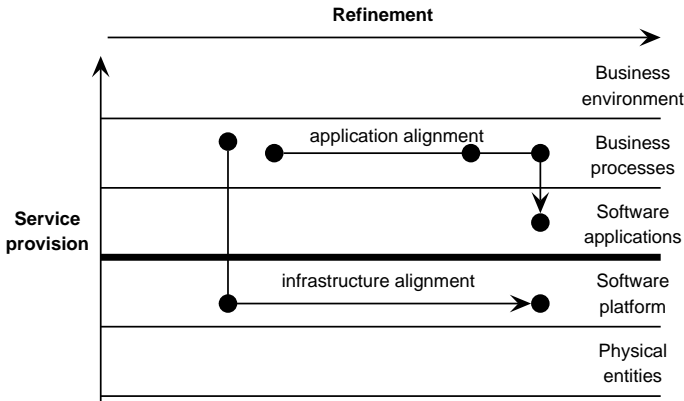
- When BIGIT was structured according to the departments of BIG, each department of BIG was visited by the same people from BIGIT for years. This constituted an early warning system for upcoming change requests. The relevant department of BIGIT was often aware of a change request long before it formally was made, and some preliminary work could be done based on this early information. In the current structure, there is no such early warning system. What is worse, the customer thinks that a project has started when a change request has been sent to the Requirements Engineering department of BIGIT. This department has however its own work processes and scheduling and it takes some time to get the requirements right. When the project goes through Architecture, some missing requirements are discovered, creating more delays, and by the time the project arrives at Detailed Design and Implementation, it could already have been delivered under the old organization structure.
- The close relationship between the application development departments of BIGIT and its customer departments not only created a shared implicit knowledge base and an early change warning system, it also made requirements engineering short and effective. RE was done by application architects, and part of the global architecture design consisted actually of identifying missing requirements and disambiguating requirements that are present.

These observations suggest that Conway's law is not a law of nature—it can be violated—but a guideline for architecture alignment. The hypothesis coming out of these observed failures is that an organization structure that violates Conway's law is less efficient than one that does observe it.

The restructuring of BIGIT indicates an extensions to Conway's law. Conway says that the designed system will have a structure isomorphic to the communication structure, and thus to the work breakdown structure, of the design group. This is valid for new systems. BIGIT illustrates that for existing systems, the reverse also holds: The design group will have a communication structure that is isomorphic to the system under (re)design. For example, the architecture group in the restructured BIGIT internally maintains a structure that separates application architects from infrastructure architects. The application architects maintain a division of work that is isomorphic to the department structure of BIG, and the infrastructure architects maintain a structure that is isomorphic to the infrastructure domain structure.

### 4.3 The Alignment Framework

In a famous paper, Henderson and Venkatraman [5] discussed several strategies for aligning business and IT. If we rotate their framework 90 degrees counter-clockwise we can map it on the GRAAL framework as shown in figure 3. The



**Fig. 3.** The alignment framework. The bullets represent documents produced in the alignment process, discussed later.

horizontal dimension now shows refinement, which, as we explained earlier, is one of the GRAAL dimensions. Descriptions at the left hand side of the diagram have a high level of abstraction (few details) and descriptions at the right-hand side have a high level of refinement (many details). Strategy documents contain descriptions at a high level of refinement, and are positioned at the left-hand side of the diagram.

At BIG, alignment between applications and business was done by first designing the business processes in accordance with the business strategy, and then developing application support to fit these processes. Alignment between infrastructure and business was done by first formulating an IT strategy in accordance with the business strategy, and then implementing the IT strategy in the infrastructure domains.

These two paths do not necessarily lead to points that are in harmony with each other. At BIGIT, there was considerable tension between the two, in terms of mismatch between application requirements on infrastructure, and services offered by the infrastructure.

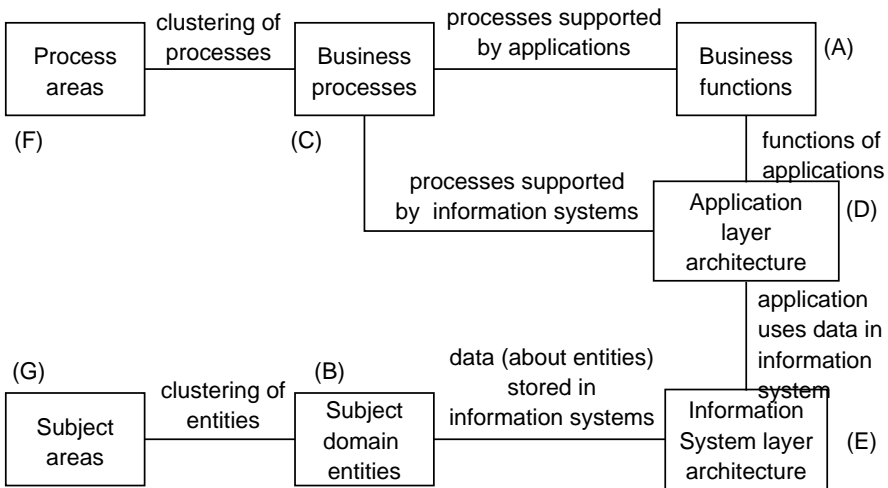
In the next two sections we discuss the two alignment paths in more detail.

#### 4.4 Application Alignment Relations

Application alignment at BIGIT takes place in three steps, that produce three documents (figure 3). Starting from a business strategy (the bullet at the tail of the application alignment arrow), a design of the business processes is produced (the second bullet). This is the basis for a document describing the architecture of the entire application layer (third bullet), which in turn is the basis for the architecture of each individual application (fourth bullet). The actual process is

not as rational as suggested here, but we will see that it does not depart very much from this rational process.

**Application layer architecture** At BIGIT, the application layer is split in two, an information system layer and what we will call a “pure” application layer. The information system layer contains the data storage systems and the pure application layer contains the data manipulation systems. Pure applications are stateless; all state is maintained by the information systems. The term “application” is thus ambiguous because it can refer to pure, stateless applications only, and to information systems plus stateless applications jointly. This ambiguity turns out not to be harmful and we will provide disambiguation where needed.



**Fig. 4.** Application alignment relations.

Figure 4 shows the application layer architecture descriptions produced at BIGIT for their customer D. We only explain the major descriptions.

- A: A list of *business functions* of D is made. For each function, its contribution to the business goal of D is described. Functions are organized into a function refinement tree with the business mission of D at its root and elementary functions at its leaves.
- B: A description of *subject domain entities* of D is made. The subject domain of a business is the part of the world about which the business contains data. (Remember that D is an information-processing organization.) The subject domain description is cross-checked with the business function descriptions by checking whether the entities about which a function needs data, are described in the subject domain description. (This cross-check is not represented in figure 4.)

- C: A description of end-to-end *business processes* of D is made, leading from customer of D to customer of D. This is cross-checked with the function refinement tree (does every process support a function and vice versa?).
- D: The descriptions of functions and processes are used to identify an *application layer architecture*. This is a list of applications and for each application, its role in one or more business processes, and its contribution to business functions.
- E: The application layer architecture and subject domain description are used to design an *information system architecture*. This is a list of information systems and for each information system, the data it stores and the applications it supports.

These descriptions have been produced using a method called Panfox [14], which is itself based on Information Engineering [15]. Observe that application alignment is defined at BIGIT for the functional aspects of our GRAAL framework:

- functions (business functions in figure 4),
- information (subject domain entities, subject areas),
- communication (the application layer and information system layer architectures, process areas),
- and behavior (business processes and process areas).

Functionality, processes, communication and information structures are the main driving forces of application alignment. Security is not a design criterion for the application layer and information system layer architecture, but it is a criterion in the design of individual application architectures, as we will see below.

We observed some problems in implementing the application alignment scheme of figure 4.

- *Subject area partitioning*. No explicit criteria are given for the definition of subject areas. A subject area is defined by BIGIT as a group of entities, the data about which are best managed jointly. But this is begging the question: When is the data about a group of entities “best” managed jointly?
- *Process area partitioning*. Similarly, a process area is defined as a “group of processes whose automated support can best be developed by one automation project”. Again, this begs the question: By which criterion should we decide whether it is “best” to automate the support of a set of projects in one project?
- *Legacy*. When the architecture is designed, there is already a current architecture that by definition is far from perfect (otherwise no application layer architecture design project would have been undertaken). Ignoring current systems altogether would rob the new architecture of any connection with reality and would make it difficult to migrate to the new architecture. But the more of the current architecture is incorporated, the less motivation there is to migrate to the new architecture. So the application layer architecture must come with a management policy for migration in the context of legacy systems. BIGIT had no such policy.

- *Ownership.* Each business process has an owner, namely the unit within D that performs the process. Naturally, this unit sees itself as the owner of the applications used in the process and of the data used and produced by these applications. But if these applications or information systems are used by other units within D as well, there is ambiguity of ownership. One of the users is then selected as the owner and the others are designated clients. The problem observed is that the owning unit is not inclined to adapt applications to the requirements of other client units, that also use the application; and they are not inclined to maintain data in formats required by applications used in other units, unless the owner of an application and information system is rewarded for maintaining them in a shape suitable for (re)use by all units within D. This indicates a direct relationship between application layer and information system layer architecture on the one hand, and managerial reward systems on the other. In BIGIT, this relationship had not been attended to and this created interminable negotiations and revisions of the architecture to align it with the actual ownership and reward structures in D. The lesson learned is that application alignment should include a design of the ownership structure and managerial reward structure of the organization.

These problems do not invalidate the approach to application alignment. However, they do point at a need for further improvement of the approach, including incorporation of legacy, migration, ownership and reward structures.

In addition to problems experienced in designing and implementing the new application layer architecture, we observed some failures in the design.

- The application layer architecture has never been approved by D’s management. Top management support is the major guideline coming out of strategic alignment research [4]. In BIGIT we observed ambiguity about what the current application layer architecture is, it being continually updated by negotiations about ownership and legacy systems. Managers in D were dimly aware of “the” application layer architecture.
- The move to the new structure of BIGIT broke the link between architects and user departments. Architects that used to work exclusively for other departments, now also worked, sometimes, for D. These architects had no commitment to the application layer architecture used by BIGIT for D.
- The function descriptions of D did not distinguish internal functions (one part of D providing a service for another part of D) from external functions (D providing services to its customers). As a result, the function model is tied up to the current structure of D. An external function model represents the essential purpose of the organization independent from the current organization structure, and is simpler than a function description that includes internal functions.
- Related to this, functions tended to be identified with organizational parts of D. This creates a misunderstanding of what functions really are: services provided to the customers of BIG.

- In other places of the function model, functions are process steps. Again, this ties the function model to the current organizational implementation and makes the function model less reusable.
- There is no clear identification criterion of processes. One process is called “mobile work”. But this is really a mobile version of most of the other processes. Some processes had a clear start event, others had not; some had an identifiable customer, others had not. This made the process models unnecessarily complex and reduced (re)usability of the models.

These failures point at a need for support from user management as well as all application architects for the application layer architecture. They also point at the need for operational guidelines for function, process and data modeling.

**Application architecture** For each individual application of D, there is a systematic process at BIGIT to define its internal architecture based on the application layer architecture. Basically, the application layer architecture provides the context diagram for each individual application and information system. Application architecture design is software architecture design [1,2,3] and due to space restrictions we will not discuss this here.

#### 4.5 Infrastructure Design Arguments

BIGIT’s customers are all other departments of BIG, of which D is only one. The infrastructure architecture is motivated in terms of four kinds of phenomena that do not refer to the needs of D but of BIG as a whole.

- *Business goals* of BIG. Examples goals are the improvement of accessibility of BIG’s services, the improvement of customer-friendliness and the facilitation of mobility of BIG’s employees.
- *Business problems* experienced at BIG. Example problems are the high cost of maintenance, the large number of systems for which users must remember passwords, and the lack of OS/390 experts.
- *Current systems* at BIG. These range from ERP systems used to the range of operating systems in use (and their various versions), all classified according to the infrastructure domains shown in figure 2.
- *The current technology market trends* for off-the-shelf applications, office software, document management software, middleware, telematics software, etc., again organized according to technology domain.

These considerations transcend not only the individual business processes at D. They also transcend D. In fact, D must define a business strategy that agrees with BIG’s strategy. The infrastructure alignment arrow in figure 3 starts from a BIG business strategy, which covers all aspects of the business, including some very general IT aspects. This is then specialized into a strategy for the IT infrastructure only. This is in turn refined into a description of the next version of the infrastructure architecture for BIG, and the acquisitions of IT that follow from this. We observed a number of related alignment failures at BIGIT,

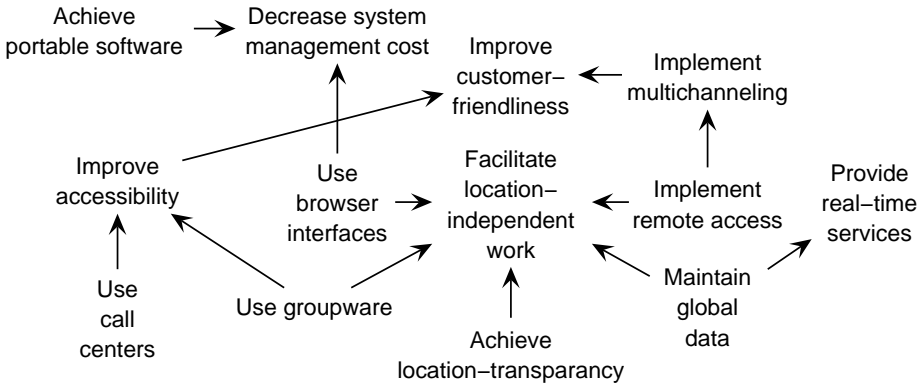
that can all be traced to the following phenomenon: Each infrastructure domain is a technical knowledge area that requires several years of study followed by constant attention to the trade press to understand. So for each domain, there is a technical specialist gathering domain knowledge. At BIGIT, this caused a number of alignment failures.

- *Conway for infrastructure.* The structure of domains is isomorphic to the structure of the set of infrastructure architects. It is impossible to change the partitioning into domains without also changing the job descriptions of these architects. Because domain knowledge is highly technical and only the infrastructure architects have this knowledge, the partitioning into domains is determined not so much by the properties of the infrastructure but by the job satisfaction of the infrastructure architects. A change in the domain architecture will be resisted if this decreases the job satisfaction of infrastructure architects.
- *Technology orientation.* Infrastructure architects follow their part of the technology market daily. This highly technical orientation makes them less sensitive to business strategies and business problems.
- *Domain islands.* The highly specialist nature of infrastructure domain knowledge even tends to isolate domain specialists from each other. The infrastructure architecture document (document 8 in figure 2) has one chapter for each domain. These chapters do not refer to each other and are even written using different typefonts and different diagram conventions, which indicates their lack of alignment.
- *No traceability.* Although business goals and business problems are listed extensively, none of the infrastructure design decisions were related to business goals. There was no inventory of goals expected to be satisfied by the decisions made, nor of problems expected to be solved by decisions made.
- *Technology-driven decisions.* Infrastructure design decisions were made in terms of market developments. The general trend was to add more technology to existing technology. The desire to terminating a line of technology that was more than 20 years old was included in the IT goals but was not followed by an actual decision to do so.

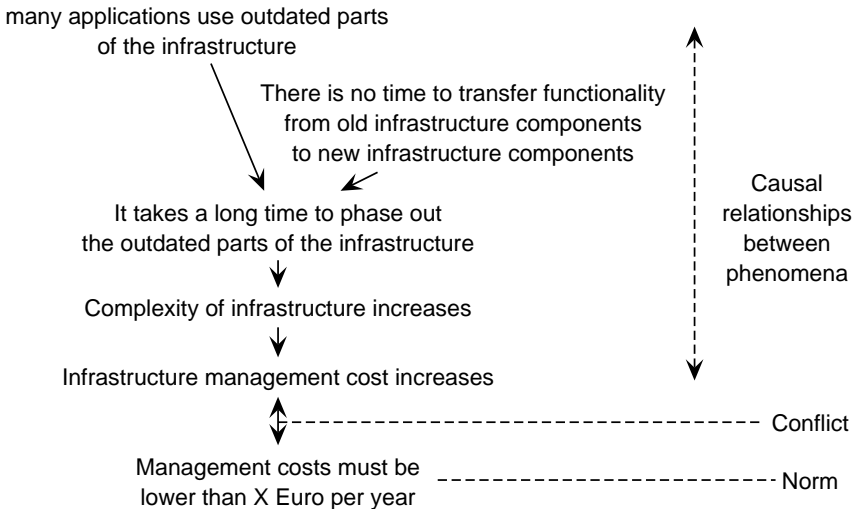
These consequences are all related to the highly specialist technical nature of infrastructure domains, but they do not follow necessarily from it. It is possible to make the link between a technology decision on the one hand, and business goals and problems on the other, explicit. The infrastructure architecture document (nr. 8 in figure 2) contains several of these links, but they are buried in a mass of details about goals, problems, legacy systems and market trends.

Figure 5 shows a fragment of a goal-directed argument. This is a rational reconstruction of part of the reasoning in the infrastructure documentation. Some of the goals are ill-defined, and others are redundant.

Figure 6 shows a problem analysis that links observable phenomena to a relevant norm. Two underlying causes lead, through a chain of causation, to the phenomenon that infrastructure management costs increase. This is a phenomenon that only becomes a problem because it conflicts with a norm, namely



**Fig. 5.** Goal graph. Each node represents a goal chosen by BIG.  $A \rightarrow B$  means that  $A$  contributes to  $B$ , in other words that  $B$  is a reason to choose  $A$ .



**Fig. 6.** problem graph.

that system management costs must be below a certain amount. From this analysis one can derive a goal such as that outdated parts of the infrastructure must be replaced.

Analyses of goals and problems do occur in the infrastructure documentation, but they are sometimes implicit, and in all cases they are hidden behind extensive discussions of technology market trends. We hypothesize that this is the main reason for misalignment of infrastructure to applications. Further case study research is needed to corroborate or refine this hypothesis.

## 5 Discussion and Conclusions

We can derive a number of hypotheses from this case analysis.

- Obeying Conway’s law is more efficient than not obeying it (section 4.2).
- The reverse of Conways law is true when there are already existing systems (section 4.2).
- Application alignment takes place in a different way from infrastructure alignment, and these two can lead to a misalignment of application architecture and infrastructure architecture (section 4.3)
- Applications alignment to business processes is dealt with by classical methods such as Information Engineering (section 4.4).
- For an architecture to be effective, support from user management as well as from all application architects is needed (section 4.4).
- Quality attributes play a role only in mapping an application architecture to the available infrastructure (section 4.3).
- If there is no strong management guidance infrastructure aligns to technology domains rather than to business goals and problems (section 4.5).

These hypotheses have been corroborated in a second case study and are currently being refined in a third.

We also identified a number of topics for further research:

- We need to find operational guidelines for function, process and data modeling (section 4.4). This is the domain of information system development methods. Our case studies, not only the one at BIGIT, show that these methods give a lot of notations and process steps but hardly any operational guidelines about how to actually make modeling decisions. We made a start identifying these guidelines [10,8] but more work needs to be done.
- Similar remarks can be made about guidelines to define application architectures (section 4.4).
- On the organization side, we also need to identify guidelines for the incorporation of legacy systems and the definition of ownership and reward structures that are critical success factors for application alignment (section 4.4).
- We merely scratched the surface of infrastructure architecture design. More work needs to be done to identify guidelines for this (section 4.5).

## Acknowledgements

Thanks are due to the IT architects of the organization called BIGIT in this paper.

## References

1. Bass, L., Clements, P., Kazman, R.: Software Architecture in Practice. Addison-Wesley (1998) [1](#), [12](#)
2. Hofmeister, C., Nord, R., Soni, D.: Applied Software Architecture. Addison-Wesley (2000) [1](#), [12](#)
3. Shaw, M., Garlan, D.: Software Architecture: Perspective on an Emerging Discipline. Prentice Hall (1996) [1](#), [12](#)
4. Chan, Y.: Why haven't we mastered alignment? the importance of the informal organizational structure. MIS Quarterly Executive **1** (2002) 97–112 [2](#), [11](#)
5. Henderson, J., Venkatraman, N.: Strategic alignment: leveraging information technology for transforming operations. IBM Systems Journal **32** (1993) 4–16 [2](#), [7](#)
6. Leifer, R.: Matching computer-based information systems with organizational structure. MIS Quarterly **12** (1988) 63–73 [2](#)
7. Opdahl, A.: Model-supported alignment of information systems architecture. In Kangas, K., ed.: Business Strategies for Information Technology Management. Idea Group Publishing (2003) [2](#)
8. Wieringa, R., Blanken, H., Fokkinga, M., Grefen, P.: Aligning application architecture to the business context. In: Conference on Advanced Information System Engineering (CAiSE 03). Volume 2681 of Lecture Notes in Computer Science., Springer (2003) 209–225 [http://is.cs.utwente.nl/GRAAL/Wieringa\\_etal\\_caise03.pdf](http://is.cs.utwente.nl/GRAAL/Wieringa_etal_caise03.pdf). [2](#), [4](#), [15](#)
9. Yin, R.: Case Study research: Design and Methods. Sage Publications (2003) Third Edition. [2](#)
10. Wieringa, R., Blanken, H., Fokkinga, M.: Documenting the ICT architecture of TSI. CTIT Technical Report 01-42, University of Twente (2001) <http://www.ub.utwente.nl/webdocs/ctit/1/00000071.pdf>. [3](#), [15](#)
11. Wieringa, R.: A survey of structured and object-oriented software specification methods and techniques. ACM Computing Surveys **30** (1998) 459–527 [3](#), [4](#)
12. Conway, M.: How do committees invent? Datamation **14** (1968) 28–31 [6](#)
13. Brooks, Jr., F.: The Mythical Man-Month: Essays on Software Engineering, Anniversary Edition. Addison-Wesley (1995) [6](#)
14. Sanden, W.v.d., Sturm, B.: Informatiearchitectuur: De infrastructurale benadering. Panfox Holding (1997) [10](#)
15. Martin, J.: Information Engineering. Prentice-Hall (1989) Three volumes. [10](#)