

VU Research Portal

Improving Solution Architecting Practices

Poort, E.R.

2012

document version

Publisher's PDF, also known as Version of record

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

citation for published version (APA)

Poort, E. R. (2012). *Improving Solution Architecting Practices*. [PhD-Thesis - Research and graduation internal, Vrije Universiteit Amsterdam].

General rights

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

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

Take down policy

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

E-mail address:

vuresearchportal.ub@vu.nl

2

Resolving Requirement Conflicts through Non-Functional Decomposition

A lack of insight into the relationship between (non-) functional requirements and architectural solutions often leads to problems in IT projects. This chapter presents a model that concentrates on the mapping of non-functional requirements onto functional requirements for architecture design. We build a framework that both provides a model and a repeatable method to transform conflicting requirements into a system decomposition. This chapter presents the framework, and discusses two cases onto which the method is applied. In one case, the method is successfully used to reconstruct the high-level structure of a solution from its requirements. The second case is one in which the method was actually used to create a solution design fitting the stakeholders' needs, and that is reproducible from its requirements.

2.1 Introduction

The primary result of any architectural design process is a blueprint of a solution, identifying the main components and their relationships from different views. A topic that is currently under close scrutiny is the derivation of these architectural components from the functional and non-functional solution requirements. The well-known discipline of Functional Decomposition (FD) can be used as a basis, but will by itself rarely yield a solution that fulfills the non-functional requirements. This is not surprising, since rules of Functional Decomposition only deal with generic best practices for achieving software quality, such as high cohesion and low coupling. FD has no rules to deal with solution-specific quality requirements.

Several approaches exist for deriving a solution's architecture from its NFRs:

CHAPTER 2. RESOLVING REQUIREMENT CONFLICTS THROUGH NON-FUNCTIONAL DECOMPOSITION

- [Boehm and In, 1996] identifies a link between NFRs and a set of product and process strategies to address them. The process Boehm proposes to select strategies is called the WinWin spiral model [Boehm and Bose, 1994]; it is basically a negotiation model.
- A group around Lawrence Chung and John Mylopoulos has done extensive work on an NFR Framework, a process-oriented approach to deal with NFRs [Mylopoulos et al., 1992, Chung et al., 1999]. They introduce the concepts of “soft-goals” and “satisficing”, meaning that goals are set without clear-cut criteria when they are fulfilled. Satisficing is a word for sufficiently satisfying the goals from the stakeholders’ point of view. NFRs are modeled as conflicting or synergistic goals in a softgoal interdependency graph. Design alternatives that realize the NFRs can subsequently be evaluated using tradeoff analysis.
- In [Bosch, 2000] the subject is dealt with by first obtaining a functionality-based architecture, and then applying architectural transformations to satisfy the NFRs. A good example of a detailed method using this iterative approach is given in [de Bruin and van Vliet, 2002].
- Publications of the Software Engineering Institute [Bass et al., 2003] show the development of a framework and tooling towards methodical architectural design, based on NFRs: Attribute Driven Design.
- Another group has developed the Component - Bus - System - Property (CBSP) method for iterative architectural refinement of requirements. In [Gruenbacher et al., 2001], the need is mentioned to group artifacts to create an architecture, but no indication is given how to do this.

All the approaches mentioned above rely on knowledge of the effect of a number of known strategies on quality attributes. Every approach needs a pre-existing catalogue of “product strategies” [Boehm and In, 1996], “operationalizations” [Chung et al., 1999], “tactics” [Bass et al., 2003] or similar. The whole Patterns community is based on the need to classify and document such known strategies [Gamma et al., 1995, Buschmann et al., 1996, Gross and Yu, 2001]. In this chapter, we present a more direct approach, based on first principles rather than a catalogue of pre-existing strategies. We have developed a method for decomposing a solution based on the conflicts in the solution requirements. We have named this method Non-Functional Decomposition (NFD) to highlight the contrast with Functional Decomposition, and to emphasize the importance of Non-Functional Requirements in this process.

NFD proposes a method for grouping and splitting of architectural entities based on requirements, and is complementary to the CBSP approach in that sense.

A clear benefit of the NFD approach is that it focuses on non-functional and other supplementary requirements right from the beginning, yielding a defined trace from those requirements to the solution structure. Moreover, the development process and its requirements are also integrated in the approach, giving a better basis for architectural and project decision trade-offs.

The sequence of this chapter is as follows. First, we will present and discuss some of the shortcomings of the generally accepted model for the architectural design process. We will then develop a refined model of this process. Then we will describe the process for deriving solution structure from supplementary requirements that is based on this model. The succeeding sections then describe two cases: one in which the NFD method was used, and one in which it is applied retrospectively to show its validity. We conclude with a discussion.

2.2 Motivation for Non-Functional Decomposition

Our interest in Non-Functional Decomposition is based on a number of distinct observations from our substantial experience in architectural design.

- *Cohesive force of supplementary requirements:* good architectures tend to cluster functions with similar supplementary requirements in the same subsystem.
- *Divide-and-conquer conflict resolution principle:* if a subsystem has to fulfill conflicting requirements, it is useful to separate the parts that cause the conflict(s).
- *Entanglement of function, structure and building process of software:* these three elements are highly interrelated.

NFD is a framework consisting of both a model of the elements involved in the architectural process, and a method for architecting software-intensive systems based on solution requirements. It is a framework in the sense that it does not venture into the details of achieving specific quality attributes (or other supplementary requirements); there is ample literature available for each conceivable attribute. Rather, it highlights the relationships between these requirements, their conflicts and ways to resolve them. It also helps in making choices about the development process.

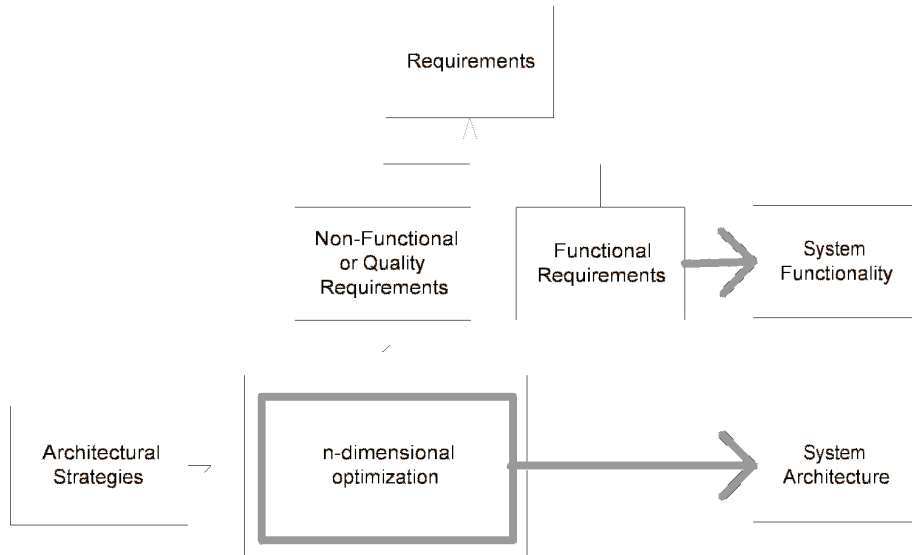


Figure 2.1: Accepted model of relationship between requirements and architecture.

2.3 Model of Requirements and Architecture

2.3.1 Accepted model for architectural design

When studying the available literature on the relationship between requirements and architecture (see §2.1), the following widely accepted model emerges.

System Requirements are usually divided into Functional and Non-Functional Requirements. These Non-Functional Requirements (NFRs) are often referred to as Quality (Attribute) Requirements; these two are treated more or less as synonyms. A generally accepted principle is *the leading NFR principle*: in designing system architectures, the Non-Functional or Quality Requirements are at least as important as the Functional Requirements. In order to satisfy NFRs the software architect applies Architectural Strategies to the system design, such as design patterns, layering techniques, etc. The architect's task then becomes an n -dimensional optimization problem: find the combination of architectural strategies that yields a solution with the best fit to the n NFRs.

The implicit model underlying this reasoning is depicted in Fig. 2.1. Although the simplicity of this view has its merits, in our experience it has some shortcomings. Par-

ticularly, the relationship between quality attributes and non-functional requirements is oversimplified, and it ignores the fact that functional requirements can also be very important in architectural design. It also ignores that NFRs often put constraints on the solution development *process* rather than on the solution architecture, implying that architectural choices are not the only contributors to satisfy NFRs. Conversely, requirements on the development process like project deadlines and budget limitations can have a large impact on solution architecture.

2.3.2 Refined requirements classification for NFD

Our new NFD model, as is illustrated in Fig. 2.2, refines the classification of requirements, and is more detailed on the non-functional aspects. Two major differences come to the foreground: functional requirements are split into primary and secondary functional requirements, and the secondary functional requirements are grouped together with the non-functional requirements. This group is called *supplementary requirements*. Additionally, a distinction is made between two types of non-functional requirements: quality attributes and delivery requirements.

Let us now define the Primary and Supplementary Requirements groups in more detail.

Primary Functional Requirements are demands that require functions which directly contribute to the goal of the solution, or yield direct value to its users. They represent the principal functionality of the solution. The identification of primary requirements (which ones to select) is similar to determining which processes in an organization are primary processes. All primary requirements are functional (there are no non-functional primary requirements), but not all functional requirements are primary requirements, as will be explained in the next section.

Supplementary Requirements represent all other requirements imposed on the solution. They can be functional or non-functional. Supplementary requirements (SRs) are always *about* primary requirements, and usually put constraints on how the primary functionality is implemented. In the NFD model, the Supplementary Requirements are further divided into three subcategories:

1. *Secondary Functional Requirements* (SFRs) require functionality that is secondary to the goal of the solution. Examples are functions needed to manage the system or its data, logging or tracing functions, or functions that implement some legal requirement. Like all other SRs, they usually apply to a particular subset of the primary requirements. For example, “All transactions in module X should be logged”, or “access to data in table Y is subject to authorization

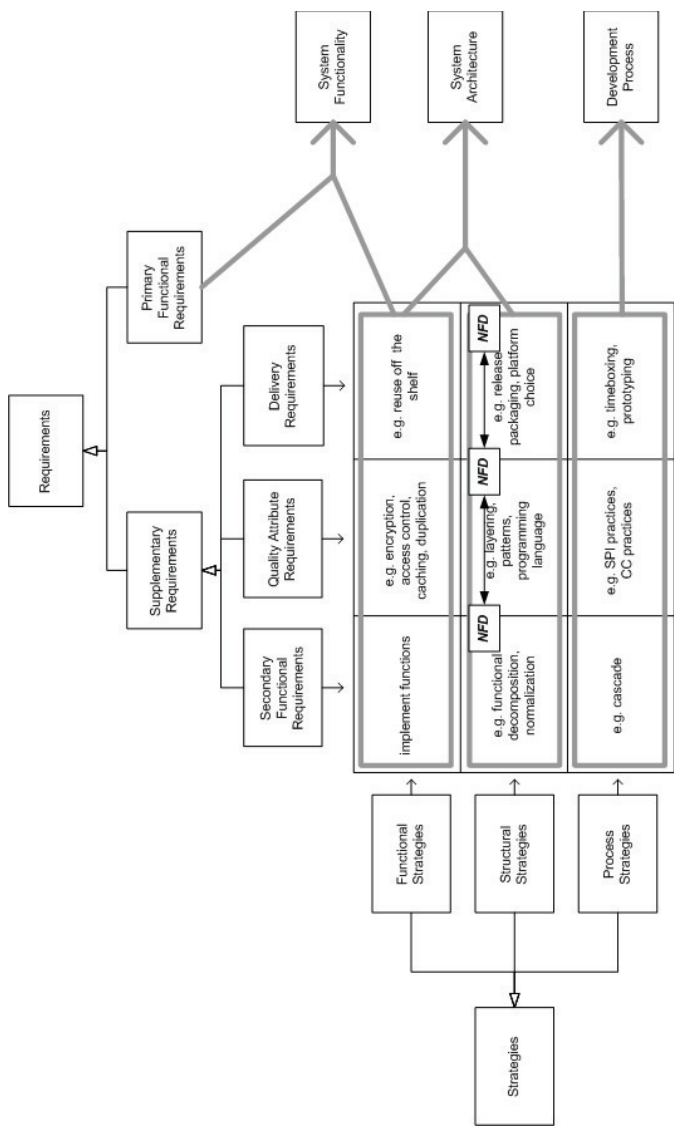


Figure 2.2: The NFD model of the relationship between solution requirements and architecture.

2.3. MODEL OF REQUIREMENTS AND ARCHITECTURE

according to model Z”. SFRs are usually not quantifiable: the solution either has the functionality or it doesn’t.

2. *Quality Attribute Requirements* (QARs) are quantifiable requirements about solution quality attributes. They can always be expressed as a number and a scale, e.g. following Gilb’s notation techniques [Gilb, 2005]¹. Examples of QARs are reliability, usability, performance and supportability. There are many taxonomies available, e.g. SQuaRE [ISO/IEC 25000, 2005].
3. *Delivery Requirements* constitute the third category of supplementary requirements. They put constraints on the solution that cannot be measured by system assessment, and incorporate e.g. managerial issues. Examples of DRs are time-to-market, maximum cost, resource availability and outsourceability. Delivery requirements can be expressed in “-ilities” that make them resemble quality attribute requirements, such as affordability or viability, but they are not about solution quality. However, they can be just as important to solution design as functional or quality requirements.

A solution’s compliance with QARs and SFRs can in principle be measured by anyone having access to the system once it has been realized, *regardless of whether they know about its history or its cost*. Compliance with Delivery Requirements can only be assessed by looking at how the solution was realized.

There is a relationship between Secondary Functional Requirements (SFRs) and functional solutions to Quality Attribute Requirements, which will be discussed later. SFRs can usually be traced back to a high-level quality need, but to express them as a quality requirement would leave too much room for interpretation. For example, the requirement to log system errors over an SMTP interface is an implementation of a manageability need, but to just require that “System management should require at most 0.1 FTE” would allow other, perhaps less desirable solutions. Satisfying Quality Attribute Requirements may also entail adding functionality to the solution, but this time the choice of functionality is at the architect’s decision.

2.3.3 The nature of requirement conflicts

The reason for a classification into primary and secondary requirements is a preparation for the NFD process that leads to solution decomposition exploiting the requirement conflicts. The NFD version of the leading NFR principle cited above is that *in designing solution architectures, the supplementary requirements are more important than*

¹Tom Gilb first introduced these techniques in [Gilb, 1988], and later incorporated them in the “language” notation [Gilb, 2005]

CHAPTER 2. RESOLVING REQUIREMENT CONFLICTS THROUGH NON-FUNCTIONAL DECOMPOSITION

the primary requirements. Primary requirements are never conflicting: if they would, the requirements would be intrinsically inconsistent or the problem statement poorly posed. However, supplementary requirements including secondary functional requirements, can appear to be conflicting, as is explained in the following paragraph.

Requirements on a software system are not intrinsically conflicting, because conflicts arise from limitations in the design strategy domain. Boehm and In have based their software tools for identifying quality-requirement conflicts on this fact [Boehm and In, 1996]. We have further analyzed the common strategies to satisfy supplementary requirements, including quality attributes. Our analysis clarifies that some quality attributes and delivery requirements may be so tightly bound to certain types of strategies, that they are effectively inherently conflicting. This situation arises when a quality attribute can only be achieved by one class of strategies, and when this class of strategies is invariably detrimental to another quality attribute. The Feature-Solution graphs introduced in [de Bruin and van Vliet, 2002] provide a good way to visualize these conflicts. We will illustrate this point with a few examples.

We have categorized the strategies for fulfilling software quality requirements into three types and nicknamed them the three strategy dimensions of solution construction: the *process* dimension, the *structure* dimension and the *functional* dimension.

1. One way to achieve supplementary requirements is by making choices in the software building *process*. Models like the Capability Maturity Model Integration [CMMI Product Team, 2010] and other software process improvement practices generally aim at improving the quality of software. Recommended practices have been documented to achieve certain quantified Safety Integrity Levels [IEC 61508, 1999] or to fulfill certain security requirements [CCPSO, 1999]. These practices tend to make the software construction process more expensive, giving rise to the first example of inherently conflicting requirements: *reliability* versus *affordability* (not an NFD quality attribute, but possibly a delivery requirement).
2. Another way to influence quality attributes or to satisfy other supplementary requirements is by making choices in the *structure* of the software. Examples of software structuring solutions include layering, applying of design patterns, choosing higher or lower level languages, modifying the granularity or modularity of the software, and so on. We have started to explore this area somewhat in [Poort and de With, 2003]. Generally speaking, the structure-based solutions seem to have one common element: more structure (i.e. higher level programming language, more layers, higher granularity etc.) means better modifiability, but less efficient code. This is the second example of inherently conflicting quality attributes: *modifiability* versus *efficiency*.

3. The third way to achieve supplementary requirements is by building *functionality* that is specifically aimed at achieving a quality or delivery objective. Examples are encryption and access control functionality to achieve a certain security goal [CCPSO, 1999], or caching functionality to achieve a certain response time and thus increase usability. Although these types of strategies are not specifically detrimental to other quality attributes, they do increase the size and complexity of the solution, leading to effects such as lower affordability and reliability. The reader should note the difference between secondary functional requirements with underlying quality needs and functional strategies aimed at satisfying quality attribute requirements. In the former case, the functional strategy is raised to the status of requirement and the responsibility of the requirement specifier. In the latter case, the functional strategy is the responsibility of the architect. In practical situations, this dichotomy may be ambiguous and quality needs are translated into functional solutions by an iterative process that involves both the requirements specifier and the architect.

In the above examples of “inherently conflicting” requirements, the conflicts emerge when applying the solution strategies to a single subsystem or component. These conflicts can often be resolved by separating the subsystem or component into different parts, and applying different solution strategies to the respective parts. Viewed from this perspective, modifiability and efficiency need not be conflicting: one can decompose a solution into low-coupled subsystems for modifiability, and then apply strategies for making the code of each subsystem more efficient. Approaches of this kind are put into practice intuitively by experienced architects, and we have modelled them in our Non-functional Decomposition Framework.

2.3.4 Applying solution strategies

The NFD model of the architecture process refines the n -dimensional optimization problem of the consensus model into a 3×3 matrix. The cells of this matrix contain strategies from each of the three strategy dimensions fulfilling each of the three types of requirements. Whereas the diagonal of the matrix contains obvious strategies (e.g. functional solutions to functional requirements), the off-diagonal cells often suggest important solutions that can help achieve requirements that would otherwise pose problems. Without being complete, we provide some examples of each of the matrix cells.

Functional strategies aimed at functional requirements: the required functions should be implemented.

Functional strategies aimed at quality-attribute requirements: these are functions

CHAPTER 2. RESOLVING REQUIREMENT CONFLICTS THROUGH NON-FUNCTIONAL DECOMPOSITION

like encryption, access control, caching and duplication, that are specifically designed to achieve a quality objective.

Functional strategies aimed at delivery requirements: delivery requirements, like outsourcing or time-to-market limitations, often exist. Their realization points to e.g. reuse of off-the-shelf components or purchasing and integrating commercial products.

Structural strategies aimed at functional requirements: examples of structures that contribute to functional requirements are database normalization, extracting of generic functionality, functional or non-functional decomposition.

Structural strategies aimed at quality-attribute requirements: lead to programming in patterns, but there are also other structural strategies contributing to quality attributes, such as the choice of programming language or correct parametrization. The NFD method itself also contributes to fulfilling quality-attribute requirements.

Structural strategies aimed at delivery requirements: delivery requirements like preferred release schedules can be realized by adapting the structure of the solution to accommodate incremental deployment. Another example is the choice of a rapid development platform (fourth generation language), which dictates a particular solution structure. NFD can also be applied here.

Process strategies aimed at functional requirements: an example is using a conventional cascade-development method, which prioritizes system functionality over time and budget limitations.

Process strategies aimed at quality-attribute requirements: examples of these are best practices from the Software Process Improvement community to improve reliability, or Common Criteria assurance packages to achieve security goals.

Process strategies aimed at delivery requirements: delivery requirements such as “user involvement” can be realized by prototyping, or “strict deployment deadline” by using a development method such as EVO [Gilb, 2005] or the Rational Unified Process^{®2} (RUP[®]) [Kruchten, 1998] that employs time-boxing techniques.

The combination of applied strategies in the process dimension results in the best *development process* to fit the solution requirements. The sum of the applied functional strategies and the realization of the primary functional requirements together compose the *solution functionality*. The *solution architecture* consists of a high-level description of the applied functional (logical view) and structural (subsystem, development, deployment views) strategies.

²RUP, Rational and Rational Unified Process are trademarks of International Business Machines Corporation.

2.3. MODEL OF REQUIREMENTS AND ARCHITECTURE

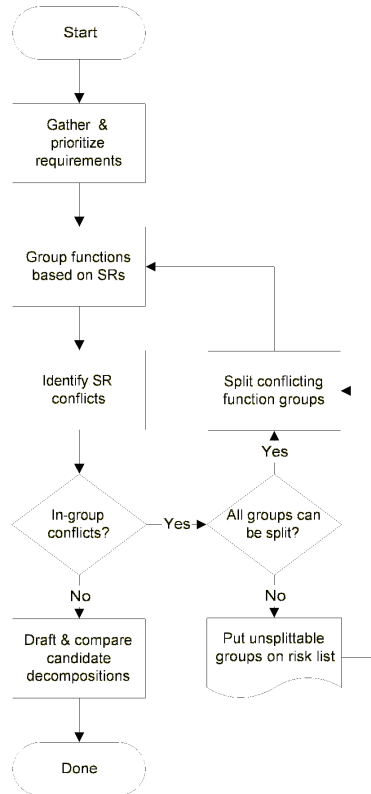


Figure 2.3: The NFD Process.

2.3.5 The role of the NFD process

Non-Functional Decomposition (NFD) is a strategy in the structural dimension of solution construction. The NFD process helps to optimize the structure of the solution for all supplementary requirements, including delivery and secondary functional requirements, which are generally associated with process or functional strategies first. It does this by adapting the solution structure to the requirement conflicts in the solution, and isolating conflicting requirements in subsystems that can then be individually optimized by applying process, structural or functional strategies. It is essentially an iterative divide-and-conquer strategy for resolving requirement conflicts.

2.4 The NFD Process

The process of NFD is depicted in Fig. 2.3 and contains the following steps.

Gather and prioritize requirements can be based on any modern requirements elicitation technique, provided that the documented requirements show how the Primary Functional Requirements (PFRs) are mapped to the Supplementary Requirements (SRs). It is important that the requirements are somehow prioritized, since prioritization of PFRs is important for project and release planning, and prioritization of SRs is important for the architecture. *Example from the Unified Process:* in the UP, FRs are generally documented as use cases, and SRs as supplementary specifications. The NFD method requires that the supplementary specifications are made specific to (groups of) use cases, e.g. by documenting them in the Use-Case Descriptions, or specifying to which use cases SRs apply in the Supplementary Specifications document. Another way of linking SRs to FRs is the use of quality attribute scenarios as described in [Bass et al., 2003].

Group functions based on supplementary requirements is the process of finding all (primary) functional requirements that share or have similar supplementary requirements, and grouping them together. This will yield a number of cross-sections of the functionality of the solution, depending on which supplementary requirement is used as a grouping criterion. In this step, the distinction between PFRs and SRs is less important: each group will have a number of functions, originating from both primary and secondary requirements, which will be treated equally during the remainder of the process. *Example:* a time-to-market priority grouping will divide functionality into groups that are candidates to be included in different release phases of the solution, while availability grouping will divide functionality into candidate groups to run on platforms with differing availability characteristics.

Identify supplementary requirement conflicts yields two types of conflicts:

1. *Grouping conflicts* are caused by differences in grouping of functions, i.e. the grouping of the functions is significantly different from one SR to the other. *Example:* there are three function groups, called WorkFlow, DataEntry and Analysis (Fig. 2.4). Security requirements for DataEntry and Analysis are similar and more restrictive than those for WorkFlow, but modifiability requirements for Analysis are more stringent than those for DataEntry and WorkFlow.
2. *In-group conflicts* are conflicting supplementary requirements within one function group. *Example:* the Analysis function group from the previous example has both critical performance requirements and high modifiability requirements.

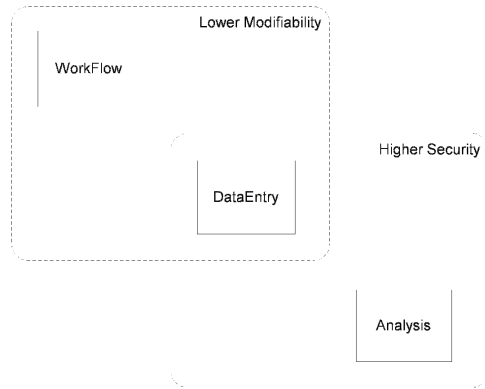


Figure 2.4: Grouping conflict example.

Split conflicting function groups deals with in-group conflicts. Most of the time, a further analysis of an in-group conflict will show that the conflicting requirements can actually be assigned to different functions. These functions are then separated, leading to a splitting of the function group. The resulting two or more new function groups may then be reconsidered for being included in other function groups, so the process re-enters the “Group functions based on SRs” stage. This loop is repeated until no in-group conflicts are left that can be split further. Function groups that cannot be split in any way are flagged as risk factors. They deserve close attention during the rest of the process and need to be dealt with prior to large-scale project implementation, e.g. in an architectural prototype.

Draft and compare candidate decompositions: After the in-group conflicts are solved, the resulting grouping conflicts will be the basis for the architectural decomposition. A number of candidate decompositions into architectural components result, each favoring the main supplementary requirement that the function grouping is based upon.

At this stage, prioritization of supplementary requirements becomes important. In our experience, the candidate decomposition that is based on the SRs having the highest stakeholder priority, yields the architecture that best fits the stakeholder requirements. This does not mean that we suggest that the n -dimensional optimization problem mentioned earlier can be reduced to a series of one-dimensional optimizations, designing for the most important requirement first and then narrowing down the design choices further for each requirement. But in the decomposition process, it turns out that the decompositions based on the SRs with the highest priority have the best chance of yield-

ing a solution the stakeholders can live with. In case of doubt, an impact analysis of the top three decompositions can be made, e.g. by using the CBAM method [Kazman et al., 2002]. The decomposition with the best fit to the stakeholder's needs (including risk and cost) is selected for implementation.

2.5 Case Study: Criminal Investigation System

2.5.1 Background

Two IT organizations affiliated with the Dutch ministry of internal affairs, the Concern Informatiemanagement Politie (CIP, "Concern Information Management Police") and the ICT-Service Coöperatie Politie, Justitie en Veiligheid (ISC, "ICT Service Co-operation Police, Justice and Safety"), were developing a product line for nationwide processing of and access to criminal investigation and intelligence data.³ The product line was called Politie suite Opsporing (PSO, "Police Suite for Investigation"). One of the authors was the lead software architect for PSO, and NFD was used to design the suite's top-level decomposition. The main challenge was to create an architecture that would allow the addition of many new products to the suite in the years to come, without compromising the strict privacy and confidentiality requirements on the system.

2.5.2 Summary of requirements

The suite's primary functionality is the support of all business processes related to criminal investigation, including management of the processes, gathering of data through multiple channels, and structuring and analysis of the data. The three most important supplementary requirements according to the stakeholders are:

- SR1** *Authorization*: access to criminal investigation data is restricted by special privacy laws. Unauthorized access to privileged data is by far the biggest threat to a criminal investigation system.
- SR2** *Reliability*: reliable application of authorization and other business rules is crucial. The system should be designed in such a way that the enforcement of especially authorization rules is reliable and stable, even after several product generations.

³The authors would like to thank the ISC and CIP organizations for granting permission to publish this case study.

SR3 *Development time*: the criminal investigation systems currently in use are based on obsolete architectures and there is an urgent need in the field to support new functionality. Exceeding the stated deadline of one year of development time is unacceptable.

SR1 is a secondary functional requirement, SR2 a quality attribute requirement, and SR3 a delivery requirement.

2.5.3 Results

NFD was applied by first mapping the most important supplementary requirements onto the functional features, and then basing the main architectural decomposition on this mapping. SR1 applies specifically to the data gathered for criminal investigation purposes. It turned out that the most reliable and best maintainable solution for the future was to create a central component for access control and storage of these data. Since the legal name for storage of such data is a police register, this central component was named the “register vault”. By using off-the-shelf components supplied by a database vendor, the register vault could be assembled and an architectural prototype evaluated within a few months time, making a good start at satisfying SR3.

In this case, RUP was used to streamline the development process. The RUP Supplementary Specifications artifact is the place to document quality and other supplementary specifications, but the standard template treats these as system wide or “general” requirements. We changed the template slightly to accommodate documenting the mapping between primary and supplementary requirements. We did this at the level of “features” as defined by RUP. This allowed us to document the trace from primary and supplementary requirements to the system decomposition design decisions.

In the end, the Non-Functional Decomposition principles turned out to be very useful in communicating to the stakeholders how our design decisions were related to their stated supplementary requirements.

2.6 Case Study: Dutch Road-Pricing System

In this section we will apply the NFD model and process to analyze a large system on roadpricing. One of the authors was involved in this project (until it was suspended for political reasons), which is described below. Although NFD was not available at the time, applying the method retrospectively to this case presents a good illustration of its principles.

2.6.1 Background

In the late 1990s, the Dutch government decided to drastically change the tax system for automobile owners and drivers. Automobile tax traditionally consisted of gasoline tax, annual motor vehicle tax and BPM (personal motor vehicle tax payed once when purchasing a vehicle). The system should be replaced by a direct tax system based on usage: gasoline tax would be reduced to the legal European minimum, and annual motor vehicle tax and BPM would be completely replaced by a roadpricing scheme called “Kilometerheffing” (“Charging by kilometer”), hereafter referred to as KMH. By differentiated pricing of road segments based on the time of day and location, the scheme could also be used to make drivers avoid congested (and thus more expensive) areas during rush hours. Feasibility of the scheme would highly depend on the use of IT systems, some of which would have to be in the vehicle. The government planned to share the cost of developing and manufacturing the in-vehicle systems (called “Mobimeters”) with the multimedia, communication and automobile industries by making generic components of those systems available to those industries.

2.6.2 System requirements

The Mobimeter requirements were shared with industry partners⁴ in order to facilitate discussion. Without going into too much detail, the original requirements can roughly be summarized as follows.

PF1 The system shall continuously measure the position and driving direction of the vehicle.

PF2 The tariff used for calculating the cost during a journey is determined on the basis of the following parameters (hereinafter referred to as Tariff Parameters):

- Date and time;
- Vehicle position;
- Direction of travel;
- Tariff table;
- Vehicle category.

PF3 A driver shall be notified of the applied tariff while driving.

PF4 The system shall determine the distance travelled by a vehicle.

⁴The original requirements were drawn up by a team led by Maarten Boasson (University of Amsterdam). They were based on the “Mobimiles” report of Roel Pieper (University of Twente), which was never formally published

2.6. CASE STUDY: DUTCH ROAD-PRICING SYSTEM

PF5 The mobility costs due is calculated as being the product of the distance travelled times the current tariff.

PF6 At least once every month in which 1000 kilometers has been driven or at least once per elapsed year, whichever comes earlier, all data shall be communicated to the tax office.

PF7 The system shall be able to receive new tariff tables.

The associated supplementary requirements (SRs) are summarized below. For brevity, we only mention the most important ones:

S1 *Privacy*: a vehicle's mobility patterns may not be deducible, either in real time (tracking) or afterwards from system data (tracing).

S2 *Verifiability*: the KMH Road-Pricing System shall enable verification that the road pricing charge has been determined correctly, without requiring more than one physical inspection per year.

S3 *Provability*: The system shall enable drivers to verify the correctness of the charges by inspecting all relevant data.

S4 *Security*: All data required for the KMH Road-Pricing System process will be protected against unauthorized modification.

S8 *Re-usability*: all in-vehicle system functions that could be useful for other applications shall be made available for re-use by third parties.

S12 *Viability*: industry partners and the relevant governmental and non-governmental organizations shall be involved as much as possible in the development of the KMH system.

S13 *Standardization*: any interfaces to be designed for in-vehicle equipment shall be developed in close cooperation with the relevant standardization bodies.

S1, S4 and S8 are quality-attribute requirements. S2 and S3 are secondary functional requirements. S12 and S13 are delivery requirements that originally did not occur in the requirements document, but in the project plan. We mention them here because in the NFD model they qualify as supplementary requirements. As will be shown, they did impact the system architecture.

Table 2.1 shows how the original supplementary requirements map onto the primary requirements. Note that some of the mappings apply to subsets of a particular

CHAPTER 2. RESOLVING REQUIREMENT CONFLICTS THROUGH NON-FUNCTIONAL DECOMPOSITION

Table 2.1: Original mapping of supplementary onto primary requirements.

	PF1	PF2	PF3	PF4	PF5	PF6	PF7
S1	X					X	
S2	X	X		X	X	X	
S3	X	X		X	X	X	
S4	X	X		X	X	X	X
S8	X		X			X	X
S12	X		X			X	X
S13						X	X

primary functions only, e.g. standardization of *only* the interface, or protection of *only* the data:⁵ a finer granularity of functionality will help us split and re-group the functions later on. The next step in the mapping process would be to split up the PFs to achieve a more exact mapping, but that exercise would be too detailed for this thesis.

A number of grouping strategies present themselves, but let us first look at the glaring in-group conflict concerning privacy versus verifiability and provability in the group PF1+PF6. According to NFD, we split the group to resolve the conflict. Clearly, if all data on which the charge is based are communicated to the tax office, the privacy requirement is violated. The objective is to split PF6 in such a way that only less privacy-sensitive data are communicated and simultaneously maintain the verifiability. After checking back with the stakeholders for the business-need behind PF6, it turns out that we can split as follows

F6a At least once every month in which 1000 kilometers has been driven or at least once per elapsed year, whichever comes earlier, the total charges and the total distance per tariff shall be communicated to the tax office.

F6b Spot checks: a travelling vehicle shall be able to answer challenges made by roadside-enforcement equipment by transferring all currently measured data and the current tariff table.

F6c On request by the driver, the system shall communicate all data used to calculate charges to him or her.

F6a is sufficient to fulfill PF6's underlying need; F6b is a functional solution to S2, and F6c is a functional solution to S3. F6b adds a short-range communication function

⁵The presence of supplementary requirements that apply to data elements or storage is quite common in our experience; these requirements often lead to special data storage components, especially if they have high priority.

2.6. CASE STUDY: DUTCH ROAD-PRICING SYSTEM

Table 2.2: Second iteration mapping of supplementary onto primary requirements.

	PF1	PF2	PF3	PF4	PF5	F6a	F6b	F6c	PF7
S1	X						X	X	
S2	X	X		X	X	X	X		
S3	X	X		X	X	X		X	
S4	X	X		X	X	X	X	X	X
S8	X		X			X	X	X	X
S12	X		X			X	X	X	X
S13						X	X	X	X

to the Mobimeter, which falls under the re-usability and standardization requirements S8, S12 and S13. This analysis leads to a new mapping table.

Table 2.2 shows the SR/FR mapping after splitting PF6. The conflict between S1 and S2 is now isolated in component F6b, the spot check function. This reflects the issue that privacy-sensitive data are present in the spot-check equipment. We put this conflict aside as a risk that will be managed by a protocol surrounding the management of these data: how long they may be stored, to what purpose, etc.

Let us now look at grouping criteria. According to the stakeholders, security, privacy and viability through re-usability are the most important supplementary requirements and in that order of priority. The security requirement S4 groups the associated data of all functions except the driver display, which basically means that the Mobimeter should contain a secure data storage component or “trusted element” that all KMH functions should have access to. Privacy requirement S1 no longer applies to F6a, since we split off the data from which the mobility pattern can be deduced. So S1 now groups PF1, F6b and F6c. S1/F6b becomes the basis for storage requirements on the roadside spot-check system, and the PF1/F6c group leads to a subsystem called the “user log”.

Finally, the viability and re-usability requirements S8 and S12 group the communications functions of F6a/b/c and PF7, the display of PF3 and the vehicle localization of PF1 into a subsystem called the “In-vehicle Telematics Services platform”. This platform is designed to have standardized interfaces (S13) both to the KMH-specific part of the in-vehicle equipment and to any additional (optional) service-related components such as a navigation system. It gives access to services like GPS localization and long-range and short-range digital communication. The final system decomposition is shown in Fig. 2.5. In the final version, the localization function was separated from the ITS platform in order to fulfill an additional reusability requirement on GPS equipment, which was already supplied in some cars.

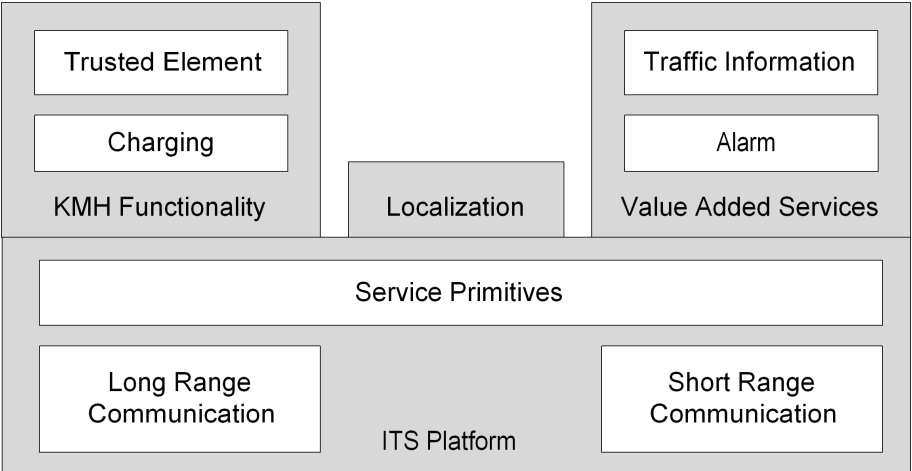


Figure 2.5: Mobimeter architecture.

2.7 Conclusions and Discussion

In this chapter, we have presented the Non-Functional Decomposition (NFD) model as a technique to bring more clarity and structure in the mapping of requirements onto a solution architecture. The key of our technique is to split the requirements into primary and supplementary requirements, and to create a mapping between those categories. The NFD process helps in optimizing the structure of the solution for all supplementary requirements, including delivery and secondary functional requirements. NFD adapts the solution structure to the requirement conflicts in the solution and isolates conflicting requirements in subsystems that can then be individually optimized by applying process, structural or functional solution strategies of which examples were presented.

The PSO product line was presented as a case study of the application of NFD. The main result here was a well documented traceability between supplementary requirements and solution decomposition design decisions. This traceability supported the project team in communicating to the stakeholders the effects of their stated requirements, and the rationale behind the main design decisions.

The KMh project was used as another case study, where we have indicated how the Mobimeter architecture as it was published, can be reconstructed with the NFD model. The examples of resolving in-group conflicts and grouping functions according to supplementary requirements were given to show the application of the method and

principles of NFD.

The validity of the observations on architecting is not only confirmed by our daily work, but can be easily verified by evaluating successful architectures like client/server or n -tier architectures. The components in these architectures all differ in their supplementary behavior, and display specific geographical accessibility, modifiability, efficiency or portability attributes.

Of the existing requirements engineering literature, a large part focuses mainly on obtaining and maintaining the right requirements [Robertson and Robertson, 2006, Wiegiers, 2003, Jackson, 2001, van Lamsweerde, 2009]. Decomposition is important in these approaches, but applies to the structure of the requirements, rather than the structure of the solution. In §2.1, we list a number of existing approaches in literature for deriving a solution's architecture from its NFRs [Boehm and In, 1996, Bosch, 2000, Chung et al., 1999, Bass et al., 2003, Gruenbacher et al., 2001]. As mentioned before, these approaches all rely on a pre-existing catalog of strategies, called by various names. NFD adds to these approaches a common rationale behind the strategies: a rationale based on the principle that conflicting NFRs can be dealt with by separating the functions that they apply to in the solution structure.

We view NFD as a framework that uses the observations made in §2.2 to improve the architecting process. These observations are not new, we believe they have always been implicit in the work of experienced designers and existing patterns and tactics. By making them explicit, NFD makes the architecting process of transforming requirements into solution design more reproducible, more transparent and more reliable. It also reveals rationale behind existing architectural patterns and tactics, and can be helpful in developing new patterns and tactics to deal with conflicting NFRs. Future work is in further application of NFD in actual technically complex projects, and in the exploration of other areas in which it could be deployed.

