Improving Solution Architecting Practices

Poort, E.R.

2012

3

Dealing with Non-Functional Requirements across the Contractual Divide

Agreement on Non-Functional Requirements between customer and supplier is crucial to a successful IT solution delivery project. In an ideal world, stakeholders and designers cooperate to achieve their common goals in a win-win situation. In a commercial setting, however, one dominant feature often introduces powerful forces from outside the technical realm. That feature is the customer/supplier relationship, usually formalized in bidding rules or as a delivery contract. Formal customer/supplier relationships often place severe limitations on information exchange between stakeholders and designers. In this chapter, we explore the effect of limitations on the process of optimal quantification of Non-Functional Requirements, and explore a number of avenues to deal with them.

3.1 Introduction

In the commercial setting of bespoke system development and integration projects, customers and suppliers embark on a quest to converge to a point where requirements can be agreed between them. Starting points on this quest are the customer's needs and the supplier's capabilities to meet those needs. Agreeing on functional requirements is usually the first step. The harder part is agreeing on non-functional requirements, which are more tightly tied to the architecture. One of the aspects that get in the way of a proper integrated approach to develop NFRs and Architecture is the formal character of the client/supplier relationship. This formal character is expressed in bidding and tendering rules in the pre-contract stage, and in the contract itself after that. The formal representation of NFRs in these situations is often a number, especially when they refer to solution quality attributes: a quantified NFR. The process of getting to these numbers is called NFR quantification.

3.1.1 Requirements and architecture in client/supplier situations

Traditionally, designing a solution to fit stakeholders' needs is done in two phases:

RE : *Requirements Engineering* expresses the needs of the stakeholders in a set of Requirements (FRs and NFRs).

AD : *Architectural Design* finds the optimal solution to address the Requirements, and expresses the solution in a Solution Architecture.

In the commercial setting of fixed price IT projects, Requirements Engineering is done by the customer, and Architectural Design by the supplier. After RE, the customer invites a number of potential suppliers to bid for the privilege of supplying a solution that fulfills the Requirements. This invitation is usually called Request for Proposal (RfP) or Invitation to Tender (ItT); we will use the term RfP. After receiving the RfP, the candidate suppliers will perform enough of the Architectural Design to be able to calculate the cost and time needed to deliver the Solution within a reasonable margin of error.

As stated in Chapter 1, NFRs are widely seen as the driving force for shaping IT systems' architectures [Mylopoulos, 2006, Chung et al., 1999, Paech et al., 2002, Bass et al., 2003]. In other words: of all the Requirements in an RfP, the NFRs have the biggest role in the Architectural Design. In Chapter 2, we have discussed various existing approaches to derive Architectural Design from NFRs, and we presented our own approach to do the same. However, as already noted by [Boehm and Bose, 1994], the notion that an architecture can be derived from requirements in one go is an oversimplification. Architecture and requirements are so closely related, that many aspects of requirements engineering can only be addressed properly if the architecture is developed at the same time. This point is made particularly eloquently in [Paech et al., 2002], which pleads for a tightly integrated approach for Functional Requirements, Non-Functional Requirements and Architecture. In our experience such an integrated approach is indeed necessary, but it is particularly difficult to achieve in the type of fixed price tendering situation described above. This is due to the traditional strict separation of roles in the tendering process, a separation that is mandated by law [US

Government, 2005, European Commission, 2004] for many government related organizations.

In this chapter, we will first present some real-life examples of the issues related to NFR quantification in a commercial setting. We will then look at the problem of NFR quantification from a number of perspectives. We will see how the tendering process interferes with proper NFR quantification, and discuss ways of dealing with this interference.

3.2 Real-life Issues Dealing with NFRs

Sometimes NFRs are explicitly tied to RfPs or delivery contracts; sometimes they are only implicitly mentioned, or ignored altogether. Either way, their impact is significant in all but the smallest projects. Conflicts between customer and suppliers can often be traced back to NFRs. Even in conflicts that revolve around the delivery of certain functionality, the delivery is often delayed by performance or security requirements related to that functionality.

This section presents some typical anonymous real-life experiences and dilemmas. The cases are used for illustration in subsequent sections.

3.2.1 Case #1: the difficulty of communicating NFRs

A customer in the financial services market invites tenders to deliver a bespoke solution to support a funds collection process. NFRs are not stated by the customer in the tender requirements documentation, and are commensurately not addressed by the (subsequently winning) supplier. At the outset of the project, an experienced engineer emphasizes to the project manager the importance of formally agreeing performance with the customer, for fear of otherwise failing to secure acceptance. The project manger agrees and sanctions work to commence on defining a non-functional requirements specification that the customer will be asked to approve.

The engineer and his team spend a significant amount of time producing a specification of performance requirements for each of the functional transactions in the system. The specification, due to the difficulties of precisely specifying such requirements, turns out to be large, highly technical and difficult to understand. The document is presented to the customer for approval, who predictably does not understand it, and engages an independent consultant to help him. The consultant spends a considerable time digesting the specification and the value of the requirements to his client, by which time the delivery project is well advanced and the project manager is beginning to panic. He convenes a meeting with the customer and his consultant to try to force agreement of the requirements, which fails as the consultant has advised his client that the requirements as stated in the specifications are highly artificial, overly qualified and of little real value. A protracted series of subsequent meetings fails to achieve agreement and meanwhile the main development is nearing completion. The impasse continues to the point that the project starts to incur significant financial loss.

3.2.2 Case #2: the unexpected cost of high availability

A bid team is responding to an invitation to tender a solution for providing administrative support to the provision of key public services. The customer states that the solution must be highly available, stating 99.999% availability requirements with onerous penalties on breach of the operational requirement. Considering the requested services levels, the bid team embarks upon a process of crafting a solution architecture with extreme high availability qualities throughout. The tender process allows little contact with the customer to refine understanding of the customer's stated requirements.

As the bid proceeds and the solution unfolds, it becomes evident that the apparently necessary hardware and software infrastructure costs are very significant. However, the bid team is convinced that this is justified and that all competitors in the tender will be responding similarly. Costs continue to spiral as the impact of the highly available solution on envisaged transactional performance becomes apparent and ever more capable infrastructure is included to compensate.

In the later stages of the bid process and pre-contract award, the bid team is invited to present the solution to the customer: the customer is aghast at the elaborate nature of the solution and the likely price to deliver it. The customer decides to suspend the tender having not received any affordable solution. The customer and all bidding suppliers have spent considerable resources and cost, which could have been avoided by more communication about impact of the availability requirement throughout the tender process.

3.2.3 Case #3: the danger of ignoring unspecified NFRs

A very large project is underway to deliver a transaction processing system for a property services company. Within the contract, the specification did not include any requirements for performance, availability or any other NFR. In hindsight, the specification did not accurately define the quantity of information to be processed or accurately express the complexity of the business process the client actually needed. The project begins work and some members of the team begin to voice concerns amongst themselves about the growing complexity of the functional requirements and growing data requirements, and in particular the potential for this combination of factors to adversely impact transactional performance experienced when the functions are implemented, fearing that the customer will find the solution to be unacceptable. These concerns are not raised with the customer. The team attempts to gain acceptance of the solution on functional grounds only, strictly fulfilling the contractually agreed specifications.

At the stage of formal acceptance testing, the customer states that he is very unhappy with the performance of the application and says that it will not viably satisfy his needs. The project team points out that no contractual obligation exists to require the supplier to deliver in-line any particular non-functional requirements. The customer steadfastly refuses to accept the system and stalemate is reached - a situation that could have been avoided by recognizing and communicating about the implicit performance requirement at an early stage.

3.2.4 Dilemmas for suppliers

In our experience, there are two flavors of dealing with NFRs in RfPs: they are either not mentioned at all (Cases #1 and #3), or they are present as hard, quantified requirements (Case #2), often poorly and ambiguously stated. Both flavors lead to dilemmas when writing a proposal in response to the RfP.

If an RfP contains hard-quantified quality attribute requirements for systems that are newly to be designed, the dilemma is caused by the uncertainty in the cost of fulfilling them. The level of uncertainty is often much larger than many stakeholders realize. New architectural combinations may have to be tried out with highly unpredictable effects on a number of interacting quality attributes. In a tendering situation, there often is no time to reduce the uncertainty by executing e.g. a proof of concept. This leaves the supplier with essentially two options: either going along with the requirement and taking on the full risk of the uncertainty, or offering a non-compliant solution - thereby risking losing the job. This is essentially a regular risk management issue - except that, as stated before, the risk in NFRs is that they can very considerably stress projects.

For quality attributes that are ignored in RfPs, the dilemma to the supplier is of a different kind. Their professionalism leads them to take into account these attributes even if no quantified NFRs are present, but how far should they go with this? Too much attention may inflate the price, potentially causing the bid to be lost due to quality attributes that are not even formally required by the customer. Too little may lead to severe problems later on, because customers have expectations about quality attributes, even if they are not explicitly quantified in the RfP (see Case #3). There are well-documented court cases [RACV Insurance Pty Ltd v. Unisys Australia Ltd, 2001] that show that suppliers have a duty of care in this area that can go beyond the contractually explicit requirements.

NFRs, whether documented in the RfP or not, are a regular source of dilemmas for suppliers responding to RfPs. One sometimes gets the impression that the tendering rules force customers to contract the supplier that has the lowest level of understanding of the NFRs. A supplier that is insufficiently aware of the impact of NFRs will generally submit a lower priced offer, because other suppliers will calculate the proper cost of addressing the NFRs, or the proper contingency needed to deal with NFR-related risks. Of course, by the same token, the NFR-unaware supplier that wins the bid will subsequently perform poorly in terms of quality attributes, and probably overrun delivery time and budget once the NFR trouble has come to light. This impression is confirmed by the example of the Dutch highway tunnel safety systems [Gram and Keulen, 2010]. The project was plagued by quality issues so severe that they caused years of delay. The government committee that investigated the trouble reports that "at the time of awarding the bid, it was known that the winning bidder scored quite badly on quality [. . .], but the quality criterium weighed insufficiently to compensate for the low price. The winning party, when asked, confirmed that, in their opinion, they could realize the project." [Gram and Keulen, 2010]. Due to European tendering rules [European Commission, 2004], in this situation the customer would not even have been *allowed* to award the bid to another supplier. This is a clear example of a supplier that won a bid due to a lack of NFR-awareness combined with tendering rules.

3.3 NFR Quantification as an Economic Problem

NFRs in RfPs can be expressed to various degrees of (un)certainty. They can be documented as vague goals that still need to be clarified and disambiguated, like the softgoals of [Chung et al., 1999]. They can also be expressed in quantified values. A lot of literature is available on the benefits of quantifying NFRs. Crisply quantified NFRs give architects a basis for their design decisions [Bass et al., 2003, Gilb, 2005], allow architectures to be validated [Clements et al., 2002], and give testers and customers a firm basis for acceptance testing [Pinkster et al., 2004]. There is no dispute that the most important quality attributes for a system should, at some point in time, be quantified in terms of objectives, targets and eventually (acceptance) test criteria. The ISO-25000 standard [ISO/IEC 25000, 2005] provides a model and metrics to do so, and several approaches exist [Barbacci et al., 2003, Gilb, 2005] as alternatives or supporting processes for these standards.

From an economic perspective, NFR quantification can be seen as an exercise in optimizing the value/cost ratio. Quantified NFRs have to be related to two economic entities: the business value of the realized NFRs (quality attribute) to the customer, and the cost to the supplier to realize the NFR (which in turn translates to price to the customer). These relationships have been extensively explored by [Kazman et al., 2002] and later by [Regnell et al., 2008] and [Berntsson Svensson, 2009], and will be explained briefly below.

3.3.1 Cost

Quality requirements tend to be very cost sensitive. This is because NFRs are fulfilled through architectural strategies and choices, such as technology selection or layering styles [Bass et al., 2003], which usually affect more than one Quality Attribute. These architectural decisions are usually discrete choices between alternatives, each carrying their own cost. These discrete choices cause discontinuous jumps in the relationship between quality attributes and cost [Regnell et al., 2008], as illustrated in the "Cost vs. Max Response Time" graph in Fig. 3.1(a). This relationship is called the "cost function", and it is determined by the architectural decisions influencing the NFR. At the time of writing the Request for Proposal, the cost function is actually unknown by the customer because the architectural choices have not been explored in depth (this is the job of the supplier). At bidding time, even the supplier usually does not have the time to sufficiently explore the cost and time needed to fulfill NFRs. For newly designed systems, figuring out the true cost function of NFRs often requires extensive model calculations or architectural prototyping, for which the deadline of tender submission is usually far too short. From the supplier's point of view, the affordability factor is acutely felt at bid stage, where the matter of competitive positioning is uppermost in their mind.

3.3.2 Value

The business value of quality is often a highly intractable entity [Garvin, 1984], which is illustrated in Case #1 above. Who can calculate the difference in value between a system with 99.99% availability and one with 99.999% availability? The difference in cost between the two may be prohibitive, as is illustrated in Case #2. The relationship between quality attribute and value is called the "value function". The value function is typically stakeholder-dependent: e.g. improving performance by 50% at night-time may be worth a lot to a particular department, while another department in the same company derives more value from increased security resilience or maintainability. In Fig. 3.1(b), an example of a relationship between business value and maximum response time for a function is depicted. At bid time, this relationship is unknown to the supplier: even if the response time requirement is quantified in the RfP, it is a single number or a statistical spread of numbers, but rarely explicitly related to the business value.

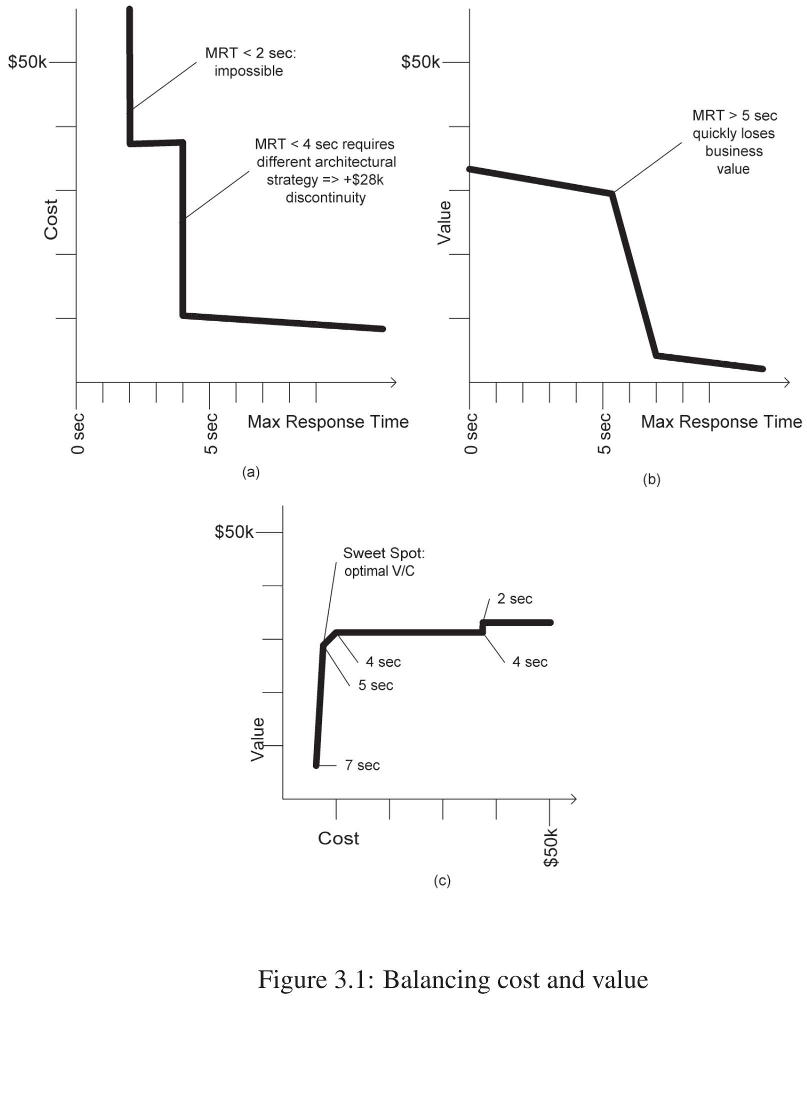

Figure 3.1: Balancing cost and value

3.3.3 Balancing cost and value

In Fig. 3.1(c), a Cost versus Value graph is derived from two underlying relationships for a particular quality attribute: the Maximum Response Time (MRT) for a function. Assuming that we want to maximize the Value/Cost ratio, finding the "sweet spot" in this graph is easy: the point on the graph that has the steepest straight line to the origin represents the optimal quantified NFR from an economic point of view. We have to keep in mind, though, that the optimization we have performed here concerns only one NFR. In reality, quality attributes are not orthogonal, so a full cost/benefit analysis would require a more complex, multidimensional calculation involving cost and value functions of all relevant quality attributes.

We now have a simplified method to quantify NFRs from an economic perspective. Three things are essential to this method:

- Supplier knowledge of the NFR's cost function.
- Customer knowledge of the NFR's value function.
- Communication of said knowledge between customer and supplier.

Even though these essentials were derived from a simplified method, it is easy to see that they are needed for any realistic approach to quantifying NFRs in a way that makes economic sense.

3.4 NFR Quantification as a Negotiation Problem

All three of the essentials mentioned in the previous section are usually low at tendering time, and significantly increase only after the contract has been signed:

Cost function knowledge increases by the research and experience of the supplier's delivery team.

Value function knowledge increases as end-users and business managers of the customer organization get more involved in the execution of the project and see more and more of the solution at work.

Communication between customer and supplier is severely restricted at tendering time, and gradually opens up after contract signing, as mutual trust grows.

So the reasonable thing to do is to postpone the quantification of NFRs until after the contract signing, when there is a relationship between customer and supplier that allows free exchange of information, and sufficient time to elaborate architectural alternatives and establish their costs. However, uncertainty in NFRs implies significant risk. It is natural for customers to seek as much certainty as possible that the system's quality attributes will fulfill their needs. The natural tendency therefore is to demand a supplier's commitment to fixed and quantified NFRs. This puts the supplier in a difficult position: should they refuse to commit, or convert the inherent risk into a contingency premium on top of the price? Either way may lead to not winning the bid, due to either non-compliance or overpricing.

The risk and cost of NFRs often become objects of contract negotiations. This does not help the three essentials mentioned above, as customer and supplier now have to deal with negotiation tactics such as risk avoidance, divide and conquer, good guy/bad guy, salami nibbling and slicing, on top of the technical difficulties of the engineering process. Especially communication of cost and value aspects between customer and supplier falls victim to the commercial necessity of playing ones cards close to the chest.

We thus come to the core of the issue: from an engineering and economic perspective, NFRs should not be quantified until cost and value knowledge and customer/supplier communication have been sufficiently established, which usually occurs well after contract signing; on the other hand, commercial reality often demands quantified NFRs committed to in the contract. In the next section, we will explore some possible solutions to this issue.

3.5 Towards Solutions

In this section, we will present two approaches that can help alleviate the issues around NFR quantification in a commercial setting: Requirements Convergence Planning and Competitive Dialogue.

3.5.1 Requirements convergence planning

As stated in §3.2.4, when responding to an RfP containing hard-quantified NFRs, suppliers with insufficient assurance that the requirements can be met basically have two options:

Scenario A Respond "compliant" and deal with the resulting risk.

Scenario B Respond "non-compliant", and offer an alternative for addressing the underlying stakeholder needs.

We asked the Logica Architecture Community of Practice for their views on these scenarios in an open e-mail question, and received a dozen responses. The anecdotal evidence in these responses led to the following:

> In *Scenario A*, the risk is usually dealt with by increasing the contingency budget, and mitigated by adding assumptions about the interpretation and measurements of the quantified NFR. This leads to a higher price for the customer, and a remaining chance that the supplier cannot achieve the required number (usually performance or availability). Failure sometimes means penalties, sometimes budget and delivery time overruns. As one respondent writes, "we usually get away with it".
>
> Most respondents prefer *Scenario B*, but indicate that only "mature" customers will agree to it - customers who are aware of the intricacies of NFR quantification. Scenario B requires room for discussion in which the supplier can highlight to the customer that particular requirements can be very expensive, and in which the hard requirement can be moved to a "target value". Instead of committing to the NFR value quantified by the customer, the supplier will commit to a *process* to find a proper balance of affordability, i.e. a number that is acceptable to the customer and achievable at reasonable cost. This process is sometimes called "calibration" or "clarification".
>
> The choice between Scenarios A and B is usually based on the following aspects:
>
> - Possibility to respond "non-compliant" without automatic disqualification.
> - Contractual conditions in case of not making the NFR (e.g. penalties).
> - Customer's openness, awareness of NFR criticality and (assessed) willingness to compromise on the NFR if the project runs into trouble.

In recent years, we have started to call the process referred to in Scenario B *Requirements Convergence Planning*. A Requirements Convergence Plan (RCP) is a plan to quantify specific quality attributes that cannot be committed to at contract signing time. This plan sets out a process of discovery and ultimately convergence on quantification of performance or other NFRs with open collaboration from a customer. The plan seeks to identify the most favorable balance of value and cost for performance attributes, whilst implicitly reducing risk.

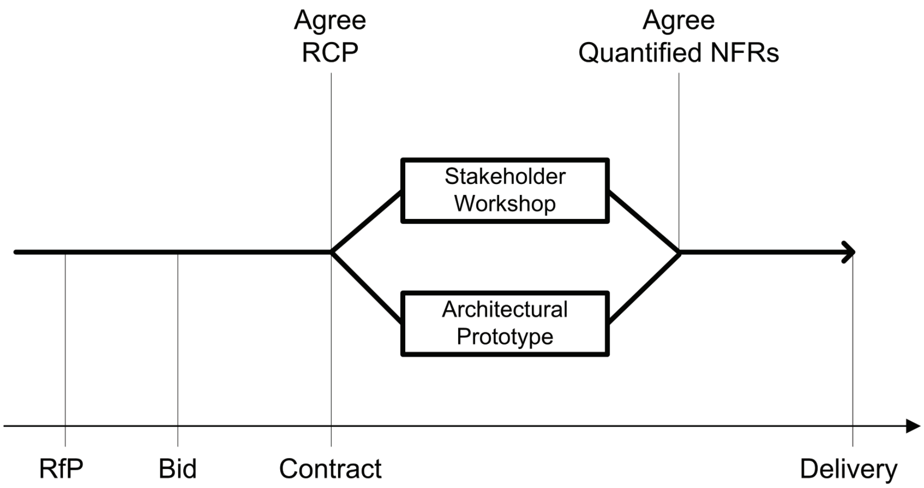

Figure 3.2: Requirements Convergence Plan

One way of planning requirements convergence is by defining a process of incrementally benchmarking architectural strategies and sharing the results of the benchmarks with the customer. In this way the customer becomes fully aware of what can feasibly be achieved within the cost constraints of the project, what risks this entails, and the impact of delivery timescales of striving for a different balance.

The RCP concept is visualized in an example in Fig. 3.2. At contract signing time, customer and supplier agree to the execution of the RCP, which in this case contains two activities: a series of stakeholder workshops (e.g. Quality Attribute Workshops [Barbacci et al., 2003]) to increase knowledge about the business value of the quality attributes, and an architectural prototype to research what quality attribute level can be achieved at what cost. Both customer and supplier are involved in both activities, stimulating the flow of information needed to make the trade-offs. At the end of the RCP period, the results of the stakeholder workshops and prototype evaluation are put together, and result in a firm quantified NFR. The supplier then commits to delivering the solution fulfilling the NFR.

Requirements convergence planning can be called a *two-stage commitment* approach for NFR quantification: at contract signing, the supplier does not commit to a quantified NFR, but to the execution of the RCP. At the end of the RCP, once a mutually agreed balanced NFR is achieved, it is signed off and committed to. The details of what happens in the RCP can be worked out on a case-by-case basis, as long as the three essentials mentioned in §3.4 above are sufficiently addressed: cost function knowledge, value function knowledge and communication between customer and supplier about them.

Apart from getting close to economically optimized NFR quantification, a benefit of this approach is that it allows the supplier to more keenly price its offer, as the performance evaluation exercise is openly effort-boxed and no explicit commitment is made to meeting specific NFRs at time of tender. The additional advantage to the customer is that he is not paying for the possibly large contingency that a supplier would otherwise have to load his offer with if this process were not to be followed.

This approach can not always be applied, since it requires an RfP that allows it to be proposed. Also, the customer must be willing to give up the certainty of a committed and quantified NFR, in exchange for the probability of better value for money spent on achieving NFRs. We have had some success with customers that welcome the openness, feel that they are more likely to get what they need and feel that they will not necessarily pay overly for it. Suppliers feel better in control of the risks, and feel as though they are in a better position to satisfy the customer's needs and make a profit.

3.5.2 Competitive dialogue

In 2004, the European Council added a new tendering procedure for the public sector, called "Competitive dialogue": *a procedure in which any economic operator may request to participate and whereby the contracting authority conducts a dialogue with the candidates admitted to that procedure, with the aim of developing one or more suitable alternatives capable of meeting its requirements, and on the basis of which the candidates chosen are invited to tender*. [European Commission, 2004] The Competitive Dialogue is meant for "particularly complex contracts". The aim of the dialogue is to "identify and define the means best suited to satisfying their needs. They may discuss all aspects of the contract with the chosen candidates during this dialogue." The Competitive Dialogue tendering procedure contains significantly less restrictions in the communication between customer and supplier at tendering stage. [US Government, 2005] contains a form of tendering called "Contracting by Negotiation", which was introduced in the regulations in 1997; like Competitive Dialogue, it has less communication restrictions than its counterpart, Sealed Bidding.

The fact that Competitive Dialogue allows a freer exchange of information between customer and suppliers makes it more suitable than the previously existing procedures for an integrated RE/AD approach in the IT solution domain. In practice, we see more and more use of the competitive dialogue tendering procedure, but it is still only applied in a minority of tenders: in 2011, only about 5% of the IT tenders Logica is interested in follows the competitive dialogue procedure. The following anonymous example shows how such a dialogue is typically conducted:

> A government ministry is looking to outsource the operation and maintenance of its Enterprise Resource Planning (ERP) application landscape to an IT service provider. The contract will be for 5 years, with a total contract value in the order of magnitude of 50M€. A Request for Information (RfI) goes out, after which a shortlist of suppliers is selected. An RfP based on the competitive dialogue model follows 5 months after the RfI. Candidate suppliers have two months to register for the bid. A six month competitive dialogue phase then starts. During the competitive dialogue phase, there are 8 workshops with each supplier. The objective of the workshops consists of two business goals: cost reduction and flexibility enhancement. Suppliers are asked to use the workshops to bring forward ideas so that the ERP landscape can be operated and maintained in a less costly and more flexible manner. Even though the original RfP contains quite detailed requirements (including many NFRs), suppliers are actively encouraged to think outside of the boundaries of these requirements during the competitive dialogue phase. The requirements may be adjusted as a result of the workshop outcome, in order to obtain the business goals. Bidders taking part in the competitive dialogue get part of their costs reimbursed.

In this example, we see that the stakeholder workshop part of the requirements convergence plan (§3.5.1) is in place. In similar cases, suppliers are also asked to provide a proof of concept - analogous to the architectural prototype in the requirements convergence plan. In other words, the four central ideas that make up requirements convergence planning are used in practice in tendering situations: two-stage commitment, stakeholder workshops, architectural prototyping and cusomter/supplier dialogue.

3.6 Discussion and Conclusions

In this chapter, we have presented some key issues related to NFR quantification in customer/supplier relationships. Critical NFRs should be quantified, but we should *beware of premature quantification*: as our real-life examples illustrate, prematurely quantified NFRs can cripple projects and lead to diverging points of view in customer/supplier relationships that are very hard to resolve.

We have concluded that, in most cases, it is impossible to find the optimal (best value/cost ratio) quantification for important NFRs at tender time. Optimal quantification requires sharing of information between customer and supplier, and it requires time to establish at least a reasonably proven estimate for the cost and value relationships. One possible way to create better NFR quantification circumstances for customers and suppliers is by means of a requirements convergence plan, which we will encounter again in Chapter 9 as a practice in our solution architecting approach RCDA. The European Union has a new tendering procedure that can be used for requirements convergence, “Competitive Dialogue”.

With the ever growing complexity of IT systems and projects, predicting system quality attributes becomes increasingly harder. Academia and industry are researching ways to improve this predictability, but they cannot win this race while the complexity of IT systems and projects increases at its current frantic rate. In the mean time, we have to deal with an imperfect world. There is no unambiguous recipe for balancing cost and value of quality attributes. Performing the balancing act while negotiating a contract is fraught with uncertainty and danger, and can even lead to failure of IT projects. The industry could benefit from a change in attitude that reflects this state of affairs. Transparency and awareness between customers and suppliers about NFRs is one part of that attitude; willingness to share the risk of unquantified NFRs is another. Both transparency and risk sharing require a basis of trust to exist between customers and suppliers in the IT industry. Without this trust, formal requirements documents or contracts with precisely quantified NFRs will not help to guarantee success.

3.6.1 Related work

Negotiating and risk balancing

Viewing NFRs as a negotiation problem was first introduced in the WinWin Spiral model of Barry Boehm et al [Boehm et al., 1995, Boehm and In, 1996]. The WinWin Spiral model is an iterative process of negotiating requirements between stakeholders, based on win-conditions. More recently, [Fricker et al., 2010] introduces the use of Implementation Proposals to facilitate the negotiation and understanding between stakeholders and architects.

The need for iterating between stakeholders to resolve requirements conflicts and reach agreement is also described in the elaboration phase of the Unified Process [Kruchten, 1998]. The issues that such elaboration iterations raise in relationship to a tendering situation have been recognized by others, and Pitette proposes a solution involving Progressive Acquisition [Pitette, 2001]. Another discussion of the difficulties of requirements specification in RfPs can be found in [Paech et al.], which reports on the experiences of a supplier in a tender process, identifies challenges and presents some possible solutions for the supplier.

An extensive treatment of balancing the forces of risk and timing can be found in [Karolak, 1995].

(Un)certainty in requirements

For dealing with uncertainty in requirements, two approaches appear in literature: modeling the uncertainty [Laplante and Neill, 2005, Noppen, 2007], and tuning the development process to better deal with change, which is one of the basic premises of the Agile movement [Agile Alliance, 2001]. The gradual increase of certainty during IT projects (§3.4) is often visualized as a "cone of uncertainty" [Mcconnell, 1997]. It was first described in [Boehm, 1981] as the "funnel curve".

[Davis, 2005] gives lots of practical advice on how to prevent overspecification of requirements. [Glinz, 2008] presents another economic perspective on NFR quantification, focusing on the risk-based need to quantify versus the cost of the quantification activities.

NFR trade-off approaches

We have seen in §3.3 that economic reasoning about NFR quantification requires knowledge of the various architectural strategies that influence the NFRs. In other words, we see that economic justification of NFR quantification requires knowledge of the solution architecture. This confirms the need for an integrated approach for requirements engineering and solution architecture as identified previously by [Paech et al., 2002]. As stated before, the one-dimensional value/cost trade-off method presented in §3.3 is a simplification: we use it here because it allows simple reasoning about quantifying an NFR, and clearly demonstrates the need for the "three essentials" for proper quantification: cost knowledge, value knowledge and communication. Almost the same method is used in QUPER [Regnell et al., 2008], who apply it in the context of product roadmapping. Several approaches for *multi*-dimensional trade-off also exist, such as CBAM [Kazman et al., 2002] and the NFR Framework [Chung et al., 1999, Lamsweerde, 2009]. [Supakkul et al., 2010] classifies such approaches as "selection patterns" and compares a number of them. All of these are more sophisticated than the method presented in §3.3, but in the end all require the same "three essentials".

[Fricker and Glinz, 2010] presents a case study analyzing and monitoring the hand-over and negotiation process between stakeholder and architect. The case study reports that a substantial part of the requirements change after the solution is presented to the stakeholder: the intended solution changed the stakeholder's position and triggered substantial requirements modifications "to exploit strengths and account for weaknesses of a possible solution". They report that "requirements understanding was perceived good-enough only after negotiation". Even though that study was not specifically targeted at NFRs, the results strongly confirm our position that proper NFR determination requires knowledge of the solution architecture.

[Gilb, 2005] appears to be a strong opponent to this position, advocating *'How Good' and 'How Much' before 'How'* as a matter of principle: "All performance requirements and resource requirements must be stated before any design idea can be fully and properly evaluated." However, in the same list of principles Gilb also states: "You cannot have correct knowledge of all the interesting requirements levels for a large and complex system in advance," indicating at least partial agreement with our position. The Design Engineering process presented in [Gilb, 2005] requires the same "three essentials" of cost knowledge, value knowledge and communication to work, indicating that it too would suffer from the communication limitations often encountered in formal client/supplier relationships.

Summarizing, we have found no existing method in industry and literature that allows proper quantification of NFRs in a situation with severe communication constraints between customer and supplier. The only way to address the issues highlighted in §3.2 is to use contract negotiation models with less constraints, taking into account the characteristics of NFR quantification.