

# Simulations for Very Early Lifecycle Quality Evaluations



## Research Section

Eliza Chiang<sup>1\*,†</sup> and Tim Menzies<sup>2</sup>

<sup>1</sup> Department of Electrical & Computer Engineering, University of British Columbia, 2356 Main Mall, Vancouver BC, Canada

<sup>2</sup> Lane Department of Computer Science, West Virginia University, PO Box 6109, Morgantown, WV, 26506-6109, USA

Chung *et al.* have proposed a graphical model that captures the interdependencies between design alternatives in terms of synergy and trade-offs. This model can assist in identifying quality/risk trade-offs early in the lifecycle of software development, such as architectural design and testing process choices. The Chung *et al.* method is an analysis framework only: their technique does not include an execution or analysis module. This paper presents a simulation tool developed to analyze such a model, and techniques to facilitate decision making by reducing the space of options worth considering. Our techniques combine Monte Carlo simulations to generate options with a machine learner to determine which option yields the most/least favorable outcome. Experiments based on the above methodology were performed on two case studies, and the results showed that treatment learning successfully pinpointed the key attributes among uncertainties in our test domains. Copyright © 2003 John Wiley & Sons, Ltd.

KEY WORDS: modeling methodology; software process models; requirement engineering; software quality assurance; Monte Carlo simulation

## 1. INTRODUCTION

Software system must meet all the functional requirements in order to provide desired functionalities to users. In addition, it must exhibit extra *non-functional* software quality attributes such as

accuracy, security, performance and other business goals. As there are no clear-cut criteria to determine whether these goals are satisfied, Chung *et al.* (1999) used the notion of *softgoals* to represent such goals. Chung *et al.* also define an entire softgoal modeling framework, featuring trade-offs and interdependencies between system quality attributes and design alternatives, but their framework is a paper design only; if an analyst wants to simulate a softgoal system, they face the problem of simulations across a space of uncertainties intrinsic to softgoals. For example, an analyst can connect two softgoals and say (for example) '*softgoal helps softgoal2*' where '*helps*' is the second strongest of the four qualitative influences defined in a softgoal

\* Correspondence to: Eliza Chiang, Department of Electrical & Computer Engineering, University of British Columbia, 2356 Main Mall, Vancouver BC, Canada

†E-mail: echiang@interchange.ubc.ca

Contract/grant sponsor: NASA; contract/grant number: NCC2-0979

Contract/grant sponsor: NASA Office of Safety and Mission Assurance under the Software Assurance Research Program led by the NASA IV&V Facility



framework.<sup>1</sup> Analysts find it intuitive to specify their connection in such a simple qualitative format. However, qualitative influences (such as 'helps') are subjected to individuals' beliefs, and are thus prone to be inconsistent. Our goal, therefore, is to develop a simulation tool that finds stable conclusions across inconsistencies within a softgoal framework.

Aside from inconsistent beliefs, another problem with drawing conclusions from a softgoal framework is the lack of supportive data. In the current software engineering practice, there are not many data available on which to perform statistical analysis (Menzies 2001). This is especially true during the early lifecycle of software development, when decisions are made based on uncertain and subjective knowledge. In the case of advanced technologies and systems, there is little past experience to learn from. Without supportive data, the relevance of any conclusion drawn from a softgoal framework is questionable. In spite of this, estimations on the potential risks and benefits of design decisions during the earlier requirement phase is essential, because these early decisions have the most leverage to influence the development to follow. The *softgoal simulation tool* presented in this paper, therefore, is designed to aid decision making in times such as early software development lifecycle, a time when domain knowledge is incomplete and inconsistent.

The premise of our methodology is that within a large space of uncertainties generated from a model, there often exist emergent stable properties (Menzies *et al.* 2002). If isolated, these properties can be used to drive a system towards the more/less preferred direction. In order to find such consistent behaviors, we apply 'bounded' randomness (i.e. guesses that fall within some defined range) to handle imprecise knowledge, and utilize *Monte Carlo* simulation (Kalos and Whitlock 1986) to explore a wide range of system behaviors. This generates a large range of behaviors which must be analyzed. TAR2 treatment learner, an analytic tool developed by Hu and Menzies (2001), is employed to automatically summarize these behaviors and return recommendations that can drive the system to some preferred mode. For example, Feather and Menzies (2002) describe one application that used formal requirements models written at the *NASA Jet Propulsion Laboratory (JPL)* (NASA website, see

references) for deep space satellite design. The formal model could generate a cost and a benefits figure for each possible configuration of the satellite (some  $10^{30}$  options in all). The black dots at the top figure in Figure 1 show what happens after 30 000 Monte Carlo simulations of that model; note the very wide range of cost and benefits. After treatment learning, a small number of constraints on the satellite configurations were found that, after 30 000 more Monte Carlo simulations, yielded the black dots at the bottom figure of Figure 1. There are two important features of these black dots at the bottom figure. First, compared to the initial black dots (shown in the top figure), the variance in the costs and benefits is greatly reduced. Secondly, the mean value of the costs and benefits are improved, i.e. reduced cost and higher benefits. The success of the Feather and Menzies application led to the speculation that one might understand the space of options within softgoals via Monte Carlo simulation and treatment learning.

As for the implementation of the softgoal simulation tool, it is designed to be lightweight and highly customizable to different business goals. Our approach is somewhat different from the standard simulation methods in the software engineering or process modeling community. Standard methods include distributed agent-based simulations (Clancey *et al.* 1996), discrete-event simulation (Harrell *et al.* 2000, Kelton *et al.* 2002, Law and Kelton 2000),<sup>2</sup> continuous simulation (also called system dynamics) (Abdel-Hamid and Madnick 1991, Stermar 2000) and state-based simulation (which includes Petri net and data flow approaches) (Akhavi and Wilson 1993, Harel 1990, Martin and Raffo 2000); our methods are closer to logic-based (Bratko 2001, chapter 20), or rule-based simulations (Mi and Scacchi 1990). In our approach:

- A model is defined that is a set of logical constraints between variables.
- A solution is generated from that model that satisfies those constraints.

In the ideal case, all model constraints can be satisfied. However, in the case of models generated from competing stakeholders, this may not be possible. Hence, our approach offers a range of operators which tries to satisfying all, many, or one of a set of

<sup>1</sup> The qualitative influences defined by Chung *et al.* (1999) are 'MAKE', 'HELP', 'HURT' and 'BREAK'.

<sup>2</sup> See also the <http://imaginethatinc.com> web site.

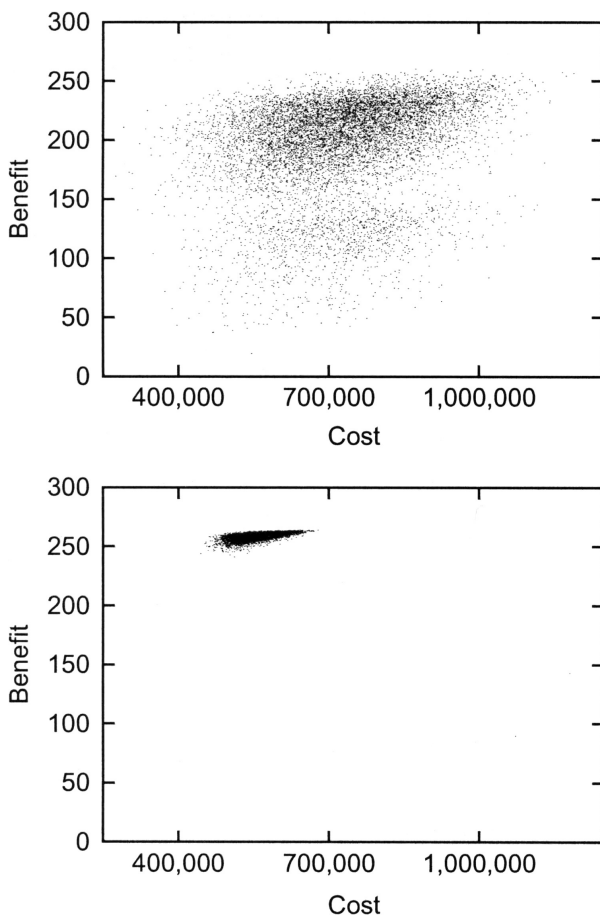


Figure 1. An example of treatment learning results

constraints. The appropriate selection of operators depends on the business at hand. In the case of software quality assurance, for example, one might combine all softgoals within the framework with logic ANDs to model the strictest quality assurance scheme,<sup>3</sup> or with logic ORs for the loosest. Users may also use a combination of the available operators to create a framework that most resembles the actual system. Examples on how these operators can be configured to suit individual business needs are presented in the case studies sections (Sections 3 and 4).

The rest of this paper is organized as follows: using the *Keyword in Context (KWIC)* framework

<sup>3</sup> Some frameworks may not yield any relevant conclusion when inferred under the strictest constraint. It is because frameworks may consist of softgoals that benefit some while harm other softgoals, and this may prevent the inference engine from drawing any conclusion if a majority of softgoals need to be satisfied.

(Shaw and Garlan 1996) as an instructional example, we first introduce the softgoal framework model proposed by Chung *et al.* Second, we present the inference process adopted by our softgoal simulation tool to execute such model. Third, we explain how Monte Carlo simulation and TAR2 treatment learning are coupled to pinpoint consistent properties within the softgoal framework, properties that can drive the system toward some preferred state. Two case studies are presented later in this paper: the first one is the analysis of the KWIC framework, a small example which is discussed comprehensively to illustrate our proposed technique; then, we introduce an advance satellite design project (SR-1) taken from *NASA IV&V Facility* as our second example. This example demonstrates how our method can be scaled up to real-world business use. Experimental results of both case studies and their implications are also presented. Finally, we summarize the proposed simulation technique and points to directions for future research.

## 2. SOFTGOAL MODELING AND SIMULATION

In this section we outline a novel approach for system modeling and simulation. This approach involves the softgoal framework model, Monte Carlo simulation, and treatment learning (Hu and Menzies 2001) to analyse and identify influential properties in the modeled system.

### 2.1. Softgoal Framework: An Overview

Softgoal framework consists of three types of softgoals: the *Non-Functional-Requirement (NFR) softgoals*; the *operationalizing softgoals*; and the *claim softgoals*. NFR softgoals represent quality requirements such as 'time-performance'. Operationalizing softgoals comprise possible solutions or design alternatives to achieve the NFR softgoals (e.g. 'incorporate javascript in online storefront'). Claim softgoals argue the rationale and explain the context for a softgoal or interdependency link (e.g. a claim may argue that 'client-side scripting loads faster'). As there are no clear-cut criteria for success, NFR softgoals may not be absolutely achieved, yet they can be



sufficiently satisfied<sup>4</sup> (Simon 1957). NFR softgoals can have an associated *priority*. *Priority softgoals* are shown in a softgoal framework with exclamation marks ('!', '!!'), or *critical/veryCritical* textually. *Priority* specifies how important a softgoal is to be fulfilled for the success of the system. *Contribution* represents the interdependency between softgoals, as well as the influence a (*claim*) softgoal has on an interdependency link.<sup>5</sup> Listed in an increasingly positive magnitude, these contributions include BREAK ('--'), HURT ('-'), UNKNOWN ('?'), HELP ('+') and MAKE ('++'). Contributions can also be combined among multiple softgoals and/or interdependency links through logic operations such as AND and OR.

The following presents the Keyword in Context (KWIC) system, a well-known example on software architectural design, to illustrate how the softgoal framework models design alternatives and quality attributes.<sup>6</sup>

The framework shown in Figure 2 defines the trade-offs among NFRs and the architectural design alternatives within the KWIC domain. The top-level *NFR softgoals* – *Comprehensibility*, *Modifiability*, *Performance*, *Reusability* – are the quality requirements to be satisfied. The design alternatives – *Shared Data*, *Abstract Data Type*, *Implicit Invocation*, *Pipes and Filters* – populate the bottom-level as operationalizing softgoals. The sub-softgoals at the middle-level of the framework are obtained by decomposing the top-level NFR softgoals. In Figure 2, for example, *Modifiability* considerations for a system are decomposed into concerns for *data representation*, *processes* and *functions*. The links from the operationalizing softgoals to NFR softgoals indicate the positive/negative impacts for each design alternatives had among quality factors. For instance, the *implicit invocation* regime makes an architectural design more extensible but requires more space, thus contributing to the

corresponding softgoals (*Extensibility[function]* and *SpacePerformance[system]*), respectively, as illustrated in Figure 2. Arguments such as '*expected size of data is huge*' is used to justify the statement: '*Pipe & Filter[Target System]* BREAKS (--) '*Space Performance[System]*', and is represented by *Claim softgoal* ('*Claim [c4]*').

In addition, the '!!' symbol associated with the NFR softgoal '*Modifiability[Data Rep]*' indicates that it is a high priority quality attribute to be satisfied. Several other attributes, such as '*TimePerformance[System]*', '*Deletability[Function]*', '*Updatability[Function]*', also serve as critical factors for overall system quality.

The KWIC framework discussed above sets an example of how softgoal modeling technique can be applied to other systems to capture the trade-offs/synergy between quality attributes and design alternatives. After a softgoal framework is constructed for the target system, it can be fed to the softgoal simulation tool for automatic inference and simulation.

## 2.2. Inference

Once the softgoal framework is defined, it is then encoded into text format for automatic inference.<sup>7</sup> The methodology for inferring the softgoal framework structure is described in this section.

Each softgoal framework requires a top-level softgoal node that represents an abstraction of overall quality. Each search performed on the framework generates a consistent 'world' – a scenario where the top-level softgoal is satisfied/denied. This 'world' can be different for each search, depending on the topology and randomness embedded in the framework definition. After a 'world' is generated, its 'goodness' is rated by computing the 'benefit' and 'cost' of this particular 'world' based on various user-configured parameters.

When uncertain qualitative inferences (e.g. HELP, MAKE) meet, combination logic is required to sum these influences. Logic operators supported by the softgoal simulation tool include AND, OR and ANY. Chaining softgoals with AND imposes the strictest constraint towards satisficing their parent softgoals, whereas chaining with OR requires

<sup>4</sup> Coined by H. A. Simon (United States social scientist and economist), 'satisfied' is defined as to be satisfied with a minimum or merely satisfactory level of performance, profitability etc., rather than a maximum or optimum level. In the context of the softgoal framework, a softgoal is said to be satisfied when it is achieved not absolutely but within acceptable limits.

<sup>5</sup> In graphical terms, they are the labels of the arrows between softgoals.

<sup>6</sup> The KWIC framework, taken from Chung *et al*'s book (1999), is a graphical expression of the architectural assessment knowledge from Shaw and Garlan (1996).

<sup>7</sup> See site <http://www.ece.ubc.ca/~elizac/vio/papers.html> for framework encoding scheme and keyword definitions.

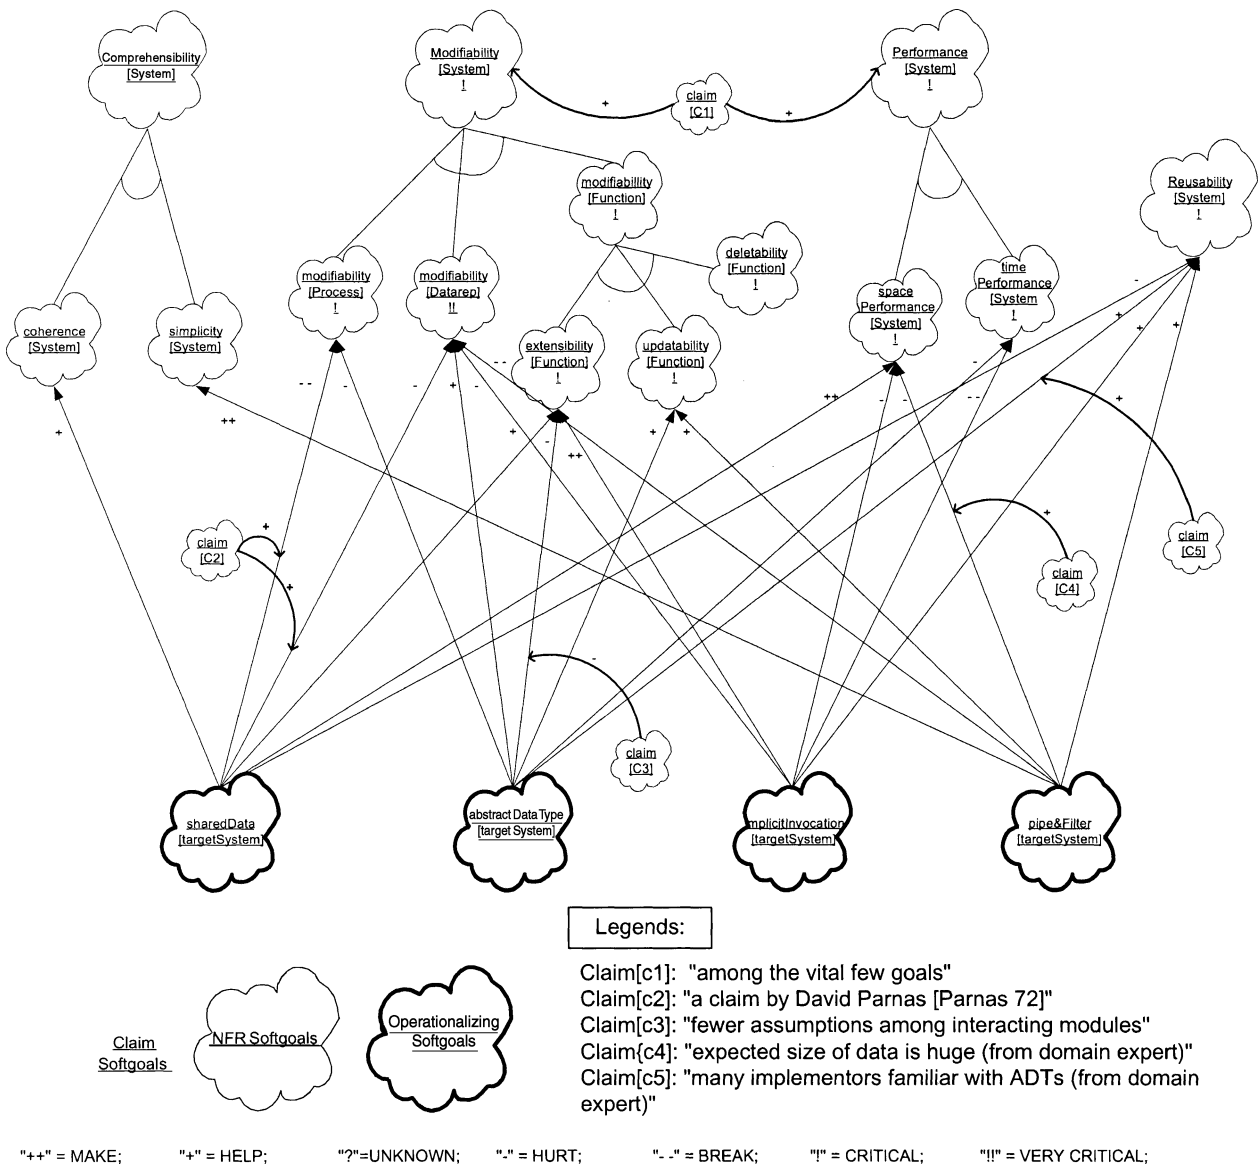


Figure 2. KWIC framework

only one satisfied softgoal to satisfy its parent. Logic ANY is similar to OR in its satisficing criteria, except that the inference engine would try to prove more than one of the chained softgoals.<sup>8</sup> Different business concerns on a domain can be addressed by using different combinations of logical operators in the analysis of softgoal frameworks. We

constructed two sample applications, the Rigorous Quality Assurance and the Weak Quality Assurance Scheme, in our study of the KWIC framework to demonstrate such capability. Details and experimental results are presented in Section 3.

### 2.3. Calculation of Cost and Benefit

As mentioned in the previous section, each inference on the framework results in a cost and benefit score. To compute these values, qualitative factors such

<sup>8</sup> For implementation details, please see site: [http://www.ece.ubc.ca/~elizac/vio/softgoal/logic\\_operations\\_used\\_in\\_framework\\_inference.html](http://www.ece.ubc.ca/~elizac/vio/softgoal/logic_operations_used_in_framework_inference.html).



as softgoal priorities ('!', '!!') and contributions ('-', '+') are involved. Numerical values are required to represent these factors during automatic inference, yet there is no definition available for quantification. This section outlines the approach taken by our simulation tool to handle the calculations of benefit and cost under such limitation.

2.3.1. Handling Source of Uncertainty within Softgoals

As there is no conventional basis for quantifying subjective knowledge (e.g. HELP, MAKE), a quantification rule is thus created to state the rankings of various qualitative strengths. Under this rule, the mean of all quantified parameters must satisfy some numerical constraints. A typical quantification rule is stated below:

0 <= score ("--") <= score ("-") <= 1
<= score ("+") <= score ("++") <= 2

A score less than 1 reflects a weakening effect. Having fixed a (0 . . . 1) range for weakening, a range between 1 and 2 is used to reflect strengthening of the influence.

Each score is expressed as a Gaussian with user-defined mean and variance values. Variances characterize the score distributions and determine how much the Gaussians overlap each other. Users can adjust the variances according to the business model and the degree of inconsistency of domain knowledge.9

Cost of each softgoal can be configured as either a static or random value. We implemented the costs of all the design alternatives in the KWIC framework (Section 3) to be static. Randomized cost calculation is applied to the framework of the SR-1 Project (Section 4).

As mentioned before, the cost and benefit computed in each inference produce a 'rating' on the 'desirability' of a particular 'world'. These ratings, once obtained by performing the Monte Carlo process, are used by a treatment learner for classification. Details on treatment learning are given in the next section.

9 For more details on benefit/cost implementation, please see site: http://www.ece.ubc.ca/~elizac/vio/softgoal/benefit.and.cost.calculation.html.

2.4. Treatment Learning with TAR2

After the ratings and the corresponding behaviors of the 'worlds' are recorded, we apply treatment learning to summarize this data. To allow TAR2 treatment learner to classify each 'world' according to its cost and benefit score, a classification and ranking scheme is required to map ranges of costs/benefits to appropriate categories. For the case study of the KWIC framework, each 'world' is rated based on the following classification scheme. The ranges of benefit were sub-divided into six bands - vvolow, vlow, low, high, vhigh and vvhhigh (in increasing magnitude) - whereas cost is sub-divided by its discrete values (from 0 to 5). Each band has roughly the same number of samples. Combining each band of the cost and benefit yields 36 classes (see Table 1 for the ranking function described). This classification scheme is applied to both rigorous (Figure 3) and loose (Figure 4) quality assurance on KWIC framework. This scheme takes account of both benefit and cost with slight preference towards lower cost. For example, Cost=zero, Benefit=vhigh has a higher ranking than Cost=one, Benefit=vvhhigh. Different preference scheme can be configured for specific business concerns.

For our softgoal simulation tool, TAR2 treatment learner (Hu and Menzies 2001) is used to perform data analysis. Base on the class ranking (Table 1), TAR2 searches the datasets for the candidate attribute ranges, that is, ranges that are more common in the highly ranked classes than the other classes. In the KWIC domain, such a candidate is a range of design approaches that drive the system into high quality/low cost, or low quality/high

Table with 3 columns: <combine\_logic>, <op>, <arithmetic[op]>. Rows include settings for all\_of, any\_of, any\_one\_of, contribution, priority, softgoal, and operationalizing softgoal.

Figure 3. Settings for KWIC framework




Figure 4. Logic configuration for rigorous quality assurance scheme

Table 1. Class rankings for KWIC framework

Cost	Benefit					
	vvlow	vlow	low	high	vhigh	vvhigh
Zero	26	17	10	5	2	1
One	28	19	12	7	4	3
Two	30	21	14	9	8	6
Three	32	23	16	15	13	11
Four	34	25	24	22	20	18
Five	36	35	33	31	29	27

cost if the range of undesirable design options is of concern. Knowing this range of attributes can greatly assist in making design decisions, as the space of considerations is narrowed down to only the attributes that would assert positive/negative impacts towards the system.

Prior to the discussion on the case studies in the next section, we introduce the incremental treatment learning (Menzies *et al.* 2002) strategy that is used in our study of the KWIC framework. To apply incremental treatment learning, a Monte Carlo simulator executes and generates datasets on the softgoal framework model. TAR2 condenses this dataset to a set of proposed treatments. After some discussions, users add the approved treatments as constraints for another round of Monte Carlo simulation. This cycle repeats until users see no further improvement.

The next section presents two case studies where our proposed simulation technique is applied.

### 3. EXPERIMENTS AND RESULT ON THE KWIC SYSTEM

As detailed in Section 2, the KWIC framework models the architectural design alternatives and quality attributes that are of business concern. The objective of this experiment is to look for design alternatives that would significantly impact the KWIC system among inconsistent knowledge such as trade-offs and beliefs.

In this section, we present two different logical interpretations applied onto the framework topology with respect to different business concerns. The rationale behind these interpretations and our observations are also discussed.

#### 3.1. Experiment 1: Rigorous Quality Assurance

This experiment is intended for system designs where strict quality assurance is mandatory. The goal is to find out what design alternatives would optimize system quality attributes. As shown in Figure 4, the NFR softgoals are combined with logic AND, meaning that all the softgoals have to be satisfied in order to satisfy their parent softgoal. To satisfy the NFR softgoals immediately above, the inference would try to satisfy the operationalizing softgoals as often as possible, and hence they are combined with logic ANY. The top-level NFR softgoals below the overall goal 'goodness of system' are combined with logic ANY.<sup>10</sup>

This above schematic is applied to the KWIC framework, and its implications is explained as follows: the operationalizing softgoals, namely the *sharedData[targetSystem]*, *abstractDatatype[targetSystem]*, *implicitInvocation[targetSystem]* and *pipe&Filter[targetSystem]*, are attached to *modifiability[DataRep]* with ANY, meaning that the inference engine will try to prove as many of the operationalizing softgoals as it can to satisfy the *modifiability[DataRep]*. Satisficing *modifiability[System]* means all its precondition softgoals – *modifiability[Process]*, *modifiability[DataRep]*, and *modifiability[Function]* – are satisfied, for they are combined with an AND. As it is impossible to find a 'world' where the *comprehensibility[System]*,

<sup>10</sup> ANY is used instead of AND to allow proper inference on this particular framework, as it is impossible to satisfy all the desired system quality attributes represented by the top-level NFR softgoals.



*modifiability[System]*, *performance[System]* and *reusability[System]* are satisfied at the same time, they are chained with ANY (instead of AND) so that the top-level goal '*goodness[System]*' can be satisfied.

Calculations on benefits and costs, as well as other parameters, are summarized in Figure 3. Notice that *deletability[System]* in Figure 2 does not associate with any operationalizing softgoal. Thus, it is assumed that certain operation is performed for its fulfillment, and the cost of this unknown operation is equal to 1. The class ranking function is described in Table 1.

Results from incremental treatment learning of the KWIC framework using the rigorous quality assurance settings are summarized in Tables 2, 3, 4 and 5. In order to clearly show how TAR2 condenses data ranges to improve the mean of the more preferred class, the results for each of the incremental process are presented as a percentile matrix. Each cell is colored on a scale ranging from white (0%) to black (100%).

As this experiment represents the case where business users put more focus on software quality

Table 2. Round 1: Percentage distributions of *benefits* and *costs* seen in 10 000 runs of Figure 4; no treatment

Cost	Benefit					
	<5.5	<11	<16.5	<22	<27.5	<32
0						
1	3.86					3.86
2	32.9	0.27				33.17
3	10.01	3.49	0.02			13.52
4	6.66	22.68	5.88	0.45	0.02	35.69
5	1.06	6.74	4.94	0.95	0.07	13.76
Total	54.49	33.18	10.84	1.4	0.09	100

Table 3. Round 2: Constraint: *sharedData of targetSystem=yes* (rigorous quality assurance)

Cost	Benefit					
	<5.5	<11	<16.5	<22	<27.5	<32
0						
1						
2	27.84					27.84
3	10.71	7.05	0.28			18.04
4	4.26	8.89	1.7	0.11	0.01	14.97
5	2.83	19.2	14.68	2.24	0.19	39.15
Total	45.64	35.14	16.66	2.35	0.2	100

Table 4. Round 3: Constraint: *implicitInvocation of targetSystem=yes, sharedData of targetSystem=yes* (rigorous quality assurance)

Cost	Benefit					
	<5.5	<11	<16.5	<22	<27.5	<32
0						
1						
2						
3	25.21	2.7				27.91
4	8.86	18.51	2.13	0.04		29.54
5	3.15	21.13	15.5	2.45	0.32	42.55
Total	37.22	42.34	17.63	2.49	0.32	100

Table 5. Round 4: Constraints: *abstractDataType of targetSystem=yes, c3=yes, implicitInvocation of targetSystem=yes, sharedData of targetSystem=yes* (rigorous quality assurance)

Cost	Benefit					
	<5.5	<11	<16.5	<22	<27.5	<32
0						
1						
2						
3						
4	10.34	24.86	3.52	0.08		38.8
5	4.68	31.01	21.51	3.64	0.34	61.2
Total	15.02	55.87	25.03	3.72	0.34	100

(i.e. benefit scores) versus costs, the benefit improvement is emphasized in the following discussion on treatment results.

Table 2 shows resulting data ranges when no constraint was imposed on the architectural design options and claims. Tables 3, 4 and 5 show the result of applying incremental treatments to Figure 4. Note that as the key decisions accumulate, the variance in behavior decreased and the mean benefit scores improved. The mean benefit drifted from <5.5 before treatment (Table 2) to <11 at treatment round 4 (Table 5). Moreover, the number of samples fell into the high benefit ranges (<27.5 and <32) increased after treatment. Based on this result, developers may focus on key issues that would greatly impact overall software quality, such as whether or not to implement *shared Data* for the system. Alternatively, if in some dispute situation, an analyst could use *c2;c4;c5* as bargaining chips. Since these claims have little overall impact, our analyst could offer them in any configuration as part of some compromise





deal in exchange for the other key decisions being endorsed.

As a last note on the results shown in Tables 3, 4 and 5, we observed that cost increases as treatments accumulate. In other words, rigorous quality assurance costs the most but doubles the average benefit. With this new information, users are now informed enough to intelligently debate the merits of rigorous quality assurance over its cost.

### 3.2. Experiment 2: Weak Quality Assurance

Often, when outside consultants are called in to offer a rapid assessment on how to improve a problematic project, they seek the fewest actions that offer the most benefit. To handle this situation, we defined a variation of the KWIC framework to simulate a weaker form of quality assurance. This assurance scheme is a simple modification in terms of its logical operations (i.e. swapping logical operators between softgoals). As shown in Figure 5, the NFR softgoals are combined with logic OR, meaning that the parent softgoal is satisfied when one of its contributing softgoals is satisfied. Similarly, only one of the operationalizing softgoals is needed to fulfill the satisficing criteria of the NFR softgoal immediately above. The parameter configuration for inference is the same as of rigorous quality assurance. In order to find the least preferred behavior, the class rankings are reversed as opposed to that of the rigorous quality assurance scheme.

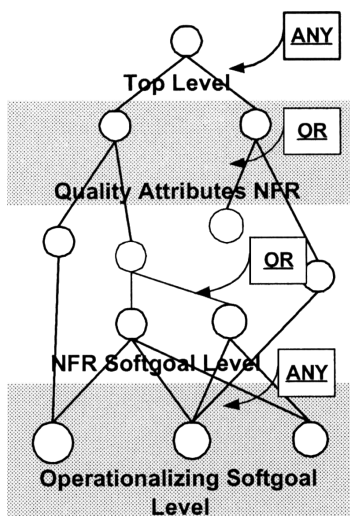


Figure 5. Logic configuration for weak quality assurance scheme

Incremental treatment learning is performed on Figure 5, and the results are shown in Tables 6 to 9.

The goal of this experiment is to determine what would negatively impact software quality in the most liberal quality assurance scheme. Comparing Table 9 (after treatments) with Table 6 (before treatments), the number of samples that fell into the lowest benefit range (<14.7) increased, which showed that benefit suffered as treatments

Table 6. Round 1: Percentage distributions of *benefits* and *costs* seen in 10000 runs of Figure 5; (weak quality assurance) no treatment

Cost	Benefit						
	<14.67	<29.33	<44	<58.67	<73.33	<87	
0							
1	3.06					3.06	
2	6.99	10.36	0.62	0.02		17.99	
3	4.18	25.26	11.92	1.41	0.1	42.87	
4	1.72	12.88	13.26	3.61	0.31	0.04	
5	0.27	1.57	1.77	0.59	0.06	4.26	
Total	16.22	50.07	27.57	5.63	0.47	0.04	100

Table 7. Round 2: Constraints: *c4=yes, pipeAndFilter of target System=no* (weak quality assurance)

Cost	Benefit						
	<14.67	<29.33	<44	<58.67	<73.33	<87	
0							
1	12.34	0.02				12.36	
2	9.18	18.66	1.77	0.05		29.66	
3	4.71	26.32	16.1	3.24	0.24	0.02	
4	0.5	3.29	2.8	0.7	0.06	7.35	
5							
Total	26.73	48.29	20.67	3.99	0.3	0.02	100

Table 8. Round 3: Constraints: *c3=yes, c4=yes, pipeAndfilter of targetSystem=no* (weak quality assurance)

Cost	Benefit						
	<14.67	<29.33	<44	<58.67	<73.33	<87	
0							
1	11.89	0.01				11.9	
2	9.03	18.6	1.95	0.04		29.62	
3	4.77	26.17	17.05	3.47	0.29	0.01	
4	0.38	2.77	2.91	0.56	0.1	6.72	
5							
Total	26.07	47.55	21.91	4.07	0.39	0.01	100



Table 9. Round 4: Constraints: c2=yes, c3=yes, c4=yes, pipeAnd-Filter of targetSystem=no (weak quality assurance)

Cost	Benefit						
	<14.67	<29.33	<44	<58.67	<73.33	<87	
0							
1	20.34	0.05				20.39	
2	8.38	28.29	4.62	0.18		41.47	
3	1.84	15.66	15.84	4.27	0.48	0.05	38.14
4							
5							
Total	30.56	44	20.46	4.45	0.48	0.05	100

accumulated. The results also suggest that, using the weaker form of quality assurance scheme, the overall software quality of the KWIC system suffers if Pipe & Filter is not implemented. Hence, users may center their discussions on the possibilities of implementing the Pipe & Filter option. Most importantly, high cost solutions can be avoided (note those results 35% over cost=3 in Table 9) without degrading overall benefits.

We have presented the treatment learning results of the KWIC framework experiments on two distinct settings, and how these settings address different business concerns. Even though their implications are different, these experimental results demonstrated how TAR2 discovers a range of consistent behavior among the space of inconsistent information. Also, as incremental treatment is applied, variance is reduced and mean values of preferred classes improved.

#### 4. CASE STUDY: NASA IV&V ACTIVITY PRIORITIZATION – A STUDY ON THE SR-1 PROJECT

The case study discussed in Section 3 demonstrated our simulation technique working on a small example, albeit one often cited in the literature. The SR-1 Project presented in this section shows how our softgoal simulation tool scales up to modern real-world software.

##### 4.1. The SR-1 Project

This case study demonstrates the capability of the softgoal simulation tool to ‘learn’ treatments from incomplete data. The overview and the

goal of our study are outlined first, followed by details of the investigations and analysis process performed in this study. The treatment learning strategy is then defined, and the results of our experimentations presented.

##### 4.1.1. Introduction

This section begins with some basic facts and terminologies being used throughout this case study. Then the objective of our study is defined, and the issues on data unavailability addressed.

The NASA SR-1 (T.S. Corporation 2002) program refers to the technologies involved in advance satellite design. It is under a contract from NASA’s Space Launch Initiative (SLI). In this study, we focus on the software components of this technology, specifically on a list of Catastrophic/Critical/High Risk (CCHR) functions and the standards used for evaluating their risk and criticality.

The NASA IV&V Facility is one of the organizations performing V&V on software projects such as SR-1. Verification and validation (V&V) is a system engineering process employing a variety of software engineering methods, techniques and tools for evaluating the correctness and quality of a software product throughout its lifecycle (Nasa website, see References). Independent V&V (IV&V) is performed by organization that are technically, managerially and financially independent of the development organization.

The Criticality Analysis and Risk Assessment (CARA) (T.S. Corporation 2002) process is a quantitative analysis used by the NASA IV&V personnel to determine the appropriate scope of V&V on a project. CARA is based on the notion that a function that has high criticality and high risk requires more extensive inspections than a function of lower criticality/risk. The CARA analysis evaluates and rates the criticality and risk of software functions based on factors such as size and complexity of program code. These ratings are then used to calculate the CARA score for each of these functions. Appropriate IV&V resources are assigned based on these scores.

##### 4.1.2. Overview and Objective of this Study

Like many other companies, project management at NASA IV&V Facility has to deal with business issues such as delivery deadlines and resource allocations. It is every manager’s goal to optimize



resource usage and reduce project costs while meeting deadline dates. On the other hand, each IV&V analysis activity consumes a different degree of resources, and some of these activities perform better in the V&V process than the others. Finding out which of these V&V activities are more powerful, and less costly at the same time, would be helpful for project resource management and task prioritization. The objective of our study, therefore, is to look for the analysis activities that are more cost-effective than others.

In our study of the SR-1 project, we applied the softgoal framework idea to sketch out the interdependencies between IV&V analysis activities and the criticality/risk assessments on SR-1 functions, which are summarized as follows:

- Criticality and risk criteria, such as *Performance and Operation*, are viewed as the quality attributes which each validated SR-1 functions are trying to satisfy. They are the NFR softgoals in the SR-1 framework.
- SR-1 software functions (e.g. *vehicle management*) are also viewed as NFR softgoals, as no software validation process can guarantee these functions to be absolutely flawless. None the less, their correctness can be sufficiently satisfied by applying IV&V analysis activities.
- IV&V analysis activities serve as the operationalizing softgoals in the framework.
- CARA ratings (*catastrophic, critical, high, moderate, low*) define the impacts of each function towards SR-1's overall criticality and risk factors upon failure. Thus, they become the interdependencies between SR-1 software functions and the criticality and risk criteria.
- Effectiveness of analysis activities relates the IV&V analysis activities to the applicable SR-1 software functions.
- Significance of SR-1 functions defines its *priority*.

To illustrate the above idea, a sample segment of the SR-1 framework is shown in Figure 6. Figure 7 shows the criticality and risk ratings of the SR-1 functions, as well as their analysis levels resulting from the CARA process. The analysis activities these levels provide by NASA IV&V for requirements, design, code and test are listed in Figure 9. Each of these SR-1 functions maps to an analysis level, which is mapped to a set of activities assigned to analyze the function. For example, the function 'Target State Filter' (*f(tFilter)*) is assessed to be level

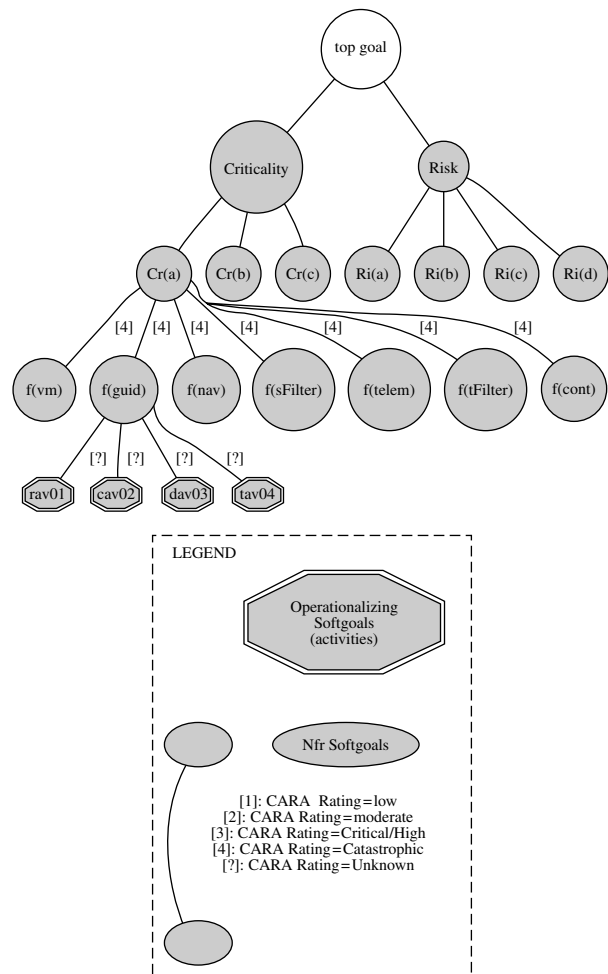


Figure 6. Segment of the SR-1 framework

'Limited' (L), thus the analysis activities *rav01–rav09*, *dav01–dav09*, *cav01–cav06*, and *tav01–tav07* are performed to validate the integrity of this function. Figure 8 shows the cost of these analysis activities in qualitative terms (*low, high* and *veryHigh*). The term 'cost' is used to generalize factors such as time, manpower, and other IV&V resources.

As we proceeded with our analysis, we found that the SR-1 softgoal framework displayed typical features of real-world systems – lack of domain knowledge and supporting data. First of all, we were unable to obtain any expert opinions regarding to the effectiveness of each analysis activities. Similarly, we have no information on which of the SR-1 functions are more important than the others. Moreover, there is discrepancy in the scaling factors for cost calculations. These factors would



Function	Criticality			Risk				Level
	Cr[a]	Cr[b]	Cr[c]	Ri[a]	Ri[b]	Ri[c]	Ri[d]	
f[vm]	4	4	2	3	2	3	3	F
f[guid]	4	4	2	3	2	2	3	F
f[nav]	4	4	2	3	2	2	3	F
f[sFilter]	4	4	2	3	1	2	3	F
f[telem]	4	1	3	3	1	3	2	F
f[tFilter]	4	1	2	3	2	2	3	L
f[cont]	4	4	1	2	1	2	1	L
f[cam]	2	1	1	3	3	2	3	B
f[exInf]	3	2	1	1	1	2	1	N
f[oReqm]	3	3	3	1	2	2	3	F
f[sMode]	3	1	3	3	2	2	3	L
f[aMode]	3	1	3	3	2	2	3	L
f[tMode]	3	1	3	3	2	2	3	L
f[dMode]	3	1	2	3	2	2	3	L
f[sbMode]	1	3	1	1	1	2	2	L
f[pMode]	3	1	1	1	1	2	2	N
f[aexInf]	3	1	1	2	1	2	1	N
f[comm]	3	1	1	1	1	3	1	N

Figure 7. CARA ratings on SR-1 software functions and corresponding analysis level

defeat the traditional quantitative approach in requirement analysis. With our proposed technique, however, we were able to perform inference and draw useful conclusions in spite of the lack of domain knowledge. The next session details how we handled the incomplete information in the SR-1 framework.

#### 4.2. Softgoal Framework Construction

The list below describes the uncertainty factors in the SR-1 framework, translated into softgoal framework-specific terminologies:

- the contributions of operationalizing softgoals (analysis activities) to NFR softgoals (SR-1 functions);
- the priorities of all NFR softgoals;
- the uncertainties intrinsic to the use of qualitative representations (e.g. 'catastrophic CARA' ratings, 'veryHigh' cost).

To allow inference while accounting for the above factors, we have made some assumptions and corresponding adjustments to the inference process.<sup>11</sup> They are listed as follows:

- Analysis activities will always contribute positively to the integrity of SR-1 functions, i.e. all the operationalizing softgoals will either HELP

<sup>11</sup> For further details, please see <http://www.ece.ubc.ca/~elizac/vio/papers.html>.

Code	Level	Cost
Requirements Analysis Activities		
rav01	B,L,F,C	notHigh
rav02	B,L,F,C	notHigh
rav03	B,L,F,C	notHigh
rav04	B,L,F,C	notHigh
rav05	B,L,F,C	notHigh
rav06	B,L,F,C	notHigh
rav07	B,L,F,C	notHigh
rav08	B,L,F,C	notHigh
rav09	B,L,F,C	notHigh
rav10	F,C	veryHigh
rav11	F,C	veryHigh
rav12	F,C	veryHigh
rav13	F,C	veryHigh
rav14	C	extremelyHigh
Design Analysis Activities		
dav01	L,F,C	high
dav02	L,F,C	high
dav03	L,F,C	high
dav04	L,F,C	high
dav05	L,F,C	high
dav06	L,F,C	high
dav07	L,F,C	high
dav08	L,F,C	high
dav09	L,F,C	high
dav10	F,C	veryHigh
dav11	F,C	veryHigh
dav12	F,C	veryHigh
dav13	F,C	veryHigh
dav14	C	extremelyHigh
Code Analysis Activities		
cav01	L,F,C	high
cav02	L,F,C	high
cav03	L,F,C	high
cav04	L,F,C	high
cav05	L,F,C	high
cav06	L,F,C	high
cav07	F,C	veryHigh
cav08	F,C	veryHigh
cav09	F,C	veryHigh
cav10	F,C	veryHigh
cav11	F,C	veryHigh
cav12	C	extremelyHigh
cav13	C	extremelyHigh
cav14	C	extremelyHigh
Test Analysis Activities		
tav01	B,L,F,C	high
tav02	B,L,F,C	high
tav03	B,L,F,C	high
tav04	B,L,F,C	high
tav05	B,L,F,C	high
tav06	B,L,F,C	high
tav07	B,L,F,C	high
tav08	L,F,C	veryHigh
tav09	L,F,C	veryHigh
tav10	F,C	veryHigh
tav11	F,C	veryHigh
tav12	C	extremelyHigh

Figure 8. Analysis activities, applicable analysis levels for SR-1's functions; and cost

or MAKE their parent NFR softgoals. To comply with this assumption, each positive contribution is randomly chosen to be either HELP or MAKE by our softgoal simulation tool during inference.

- Performing V&V on either catastrophically\_rated or critically/highly\_



Code	Requirements Analysis Activity
Rav01	Verify documentation meets intended purpose, has appropriate detail and all necessary elements.
Rav02	Validate ability of requirements to meet system needs
Rav03	Verify Traceability to and from parent requirements
Rav04	Analyze data/adaptation requirement
Rav05	Analyze Testability, Qualification requirements
Rav06	Analyze Data FnotHigh, Control FnotHoigh, moding and sequencing
Rav07	Assess development metrics
Rav08	Analyze development risks/mitigation plans
Rav09	Analyze Timing and Sizing requirements
Rav10	Review developer timing/sizing, loading engineering analysis
Rav11	Perform engineering analysis of key algorithms
Rav12	Review/use developer prototypes or dynamic models
Rav13	Develop alternative static representations (diagrams, tables)
Code	Design Analysis Activity
Dal01	Verify documentation meets intended purpose, has appropriate detail and all necessary elements
Dal02	Validate ability of design to meet system needs
Dal03	Verify Traceability to and from requirements
Dal04	Analyze database design
Dal05	Analyze design Testability, Qualification requirements
Dal06	Analyze design Data FnotHigh, Control FnotHigh, moding, sequencing
Dal07	Analyze control logic, error/exception handling design
Dal08	Assess design development metrics
Dal09	Analyze development risks/mitigation plans
Dal10	Review developer timing/sizing, loading engineering analysis
Dal11	Perform design analysis of select critical algorithms
Dal12	Review/use developer prototypes or dynamic models
Dal13	Develop alternative static representations (diagrams, tables)
Code	Code Analysis Activity
Cal01	Verify documentation meets intended purpose, has appropriate detail and all necessary elements
Cal02	Verify Traceability to and from design
Cal03	Verify Architectural design compliance (structure, external I/O, and CSCI executive moding, sequencing and control)
Cal04	Verify supportability and maintainability
Cal05	Access code static metrics
Cal06	Verify CSU and CSC level logical structure and control fnotHigh
Cal07	Verify internal data structures and data fnotHigh/usage
Cal08	Verify error and exception handling
Cal09	Verify code and external I/O data consistency
Cal10	Verify correct adaptation data and ability to reconfigure
Cal11	Verify correct operating system and run time libraries
Code	Test Analysis Activity
Tal01	Analyze System level verification requirements to verify that test definition, objectives, plans and acceptance criteria are sufficient to validate system requirements and operational needs associated with CCHR functions
Tal02	Verify Software Test Plan qualification testing methods and plans are sufficient to validate software requirements and operational needs
Tal03	Verify test cases traceability and coverage of software requirements, operational needs and capabilities
Tal04	Verify software STD test case definition inputs, expected results, and evaluation criteria comply with STP plans and testing objectives
Tal05	Analyze correct dispositioning of software test anomalies
Tal06	Validate software test results compliance with test acceptance criteria
Tal07	Verify trace and successful completion of all software test case objectives
Tal08	Verify ability of software test environment plans and designs to meet software testing objectives
Tal09	Verify regression tests are sufficient to determine that the software is not adversely affected by changes
Tal10	Analyze STD procedures for test setup, execution, and data collection; confirm procedures completely and correctly test referenced requirements, confirm inspection and analysis completely verifies referenced requirements
Tal11	Monitor execution of software testing as needed

Figure 9. Analysis activities keys to Figure 8

rated functions is assumed to be always beneficial towards the overall safety of SR-1 software. Also, we assumed that doing V&V to lowly-rated SR-1 functions has negative impacts on

the framework. The rationale of such assumption is that the additional workload may hinder job performance of V&V specialists, and hence out-weigh the gains.



- For the *moderately\_rated* functions, the effect is assumed to be either positive or negative. Therefore, such rating is transformed to be either *lowly\_rated* or *critically\_rated* during inference.
- All the NFR softgoals have the same priorities.
- The numeric values of the qualitative terms fall within pre-defined ranges, as shown in Figure 11.

Regarding the cost discrepancy, two versions of cost functions were presented to us. Figure 10 describes these functions.

Other settings used in inferring SR-1 framework are shown in Figure 11. The resulting SR-1 framework consists of 48 operationalizing softgoal nodes, 28 NFR softgoal nodes, and hundreds of edges representing softgoal contributions. As this framework was too large to be legible, we were unable to present it in this paper.

After the SR-1 softgoal framework was constructed, we carried out our investigations using the softgoal simulation tool. Details on the experimental settings (class ranking functions, logic configurations etc.) and treatment learning results are given in the next section.

### 4.3. Experiments and Results

The class ranking function used for the SR-1 framework is similar to that of the KWIC framework. The

version	1	2
<notHigh> (X)	$mean(X) = 1$	$mean(X) = mean(Y) * 0.7$
<high> (Y)	$Y = mean(X) * 10$	$mean(Y) = 2, 0 < Y < 10$
<veryHigh> (Z)	$Z = mean(Y) * 10$	$mean(Z) = Y * F, 1.1 \leq F \leq 1.7$

Figure 10. Two versions of cost function

<combine_logic>	<op>	<arithmetic[op]>
<i>all_of</i>	rand	minimum
<i>any_of</i>	rany	summation
<i>any_one_of</i>	ror	maximum
<contribution>	<value> [contribution]	<arithmetic> [contribution]
<i>helped</i>	mean=1.4	multiply
<i>made</i>	mean=1.8	multiply
<i>catastrophically_rated</i>	mean=1.9	multiply
<i>critically_rated</i>	mean=1.4	multiply
<i>highly_rated</i>	mean=1.1	multiply
<i>lowly_rated</i>	mean=0.4	multiply

Figure 11. Miscellaneous settings for SR-1 framework

range of benefits and costs were sub-divided into four bands (from *vlow* to *vhigh*), and each band consists of roughly the same examples. Table 10 shows these class rankings. Two studies were conducted based on this ranking function:

- In the ‘most preferred’ study, TAR2 looks for behaviors that would contribute to the integrity of the SR-1 functions.
- Conversely, in the ‘least preferred’ study, we reversed the order of class ranks to find treatments, which would assert negative impact to the framework.

We constructed and experimented on two variations of the SR-1 framework, which differed in their logical compositions<sup>12</sup> Figure 12 shows the weakest

Table 10. Class rankings for SR-1 framework

Cost	Benefit			
	vlow	low	high	vhigh
vlow	10	5	2	1
low	12	7	4	3
high	14	9	8	6
vhigh	16	15	13	11




Figure 12. SR-1 framework: 1

<sup>12</sup> The appropriateness of these framework variants in representing the real situations has yet to be determined by NASA experts.



quality assurance scheme build to represent the realistic business situation, where analysts try to perform as many analysis activities to fulfill the integrity of a SR-1 function, and hence the overall software quality. In this framework, the analysis activities at the bottom level were chained with an ANY, and attached to their corresponding SR-1 function. The SR-1 functions were bound to each of their critical/risk criteria with an OR, which were also bound to their upper-level softgoals with OR. For instance, *rav01-rav09* and *tav01-tav07* are chained with ANY and attached under SR-1 function *f[cam]*, meaning that the associated activities would be proven as many times as it could during inference to satisfy *f[cam]*. Function *f[cam]*, together with other functions (such as *f[vm]*, *f[guid]*, *f[nav]* etc.) were chained to the risk criteria *Ri[d]* with logic OR. Therefore, satisficing one SR-1 function would be sufficient to satisfy the risk criteria. Similarly, the risk criteria *Ri[a]* to *Ri[d]* are combined with OR under the overall *risk* softgoal, which was connected to the top-level softgoal with an ANY. In other words, satisficing one *risk* criteria would be enough to satisfy the overall *risk* softgoal, which would lead to the top-level softgoal being satisfied.

For comparison with Figure 12, we proposed another SR-1 framework to represent rigorous quality assurance. For this, we derived two prototypes of such a framework, presented in Figures 13 and

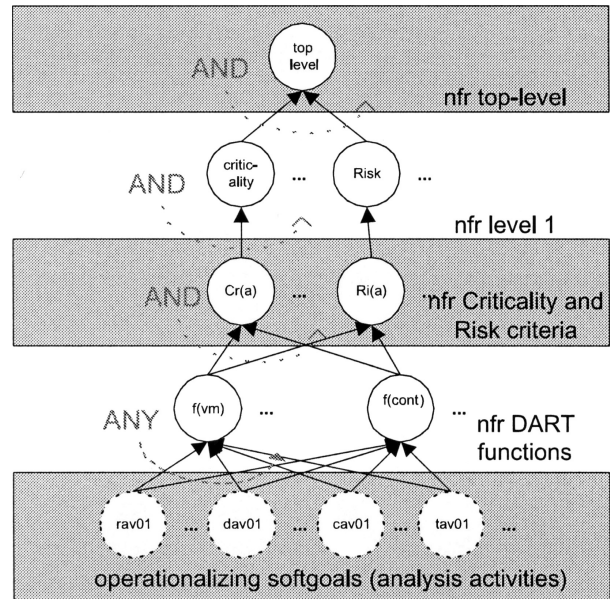


Figure 14. SR-1 framework: 3

14. Figure 14 defined the strictest form of quality assurance, in which all the NFR softgoals (i.e. SR-1 functions, risk/criticality criteria, *risk* and *criticality* softgoals, and the top-level softgoal) are combined with AND, except for the bottom level operationalizing softgoals (i.e. analysis activities), which were combined with ANY. We reviewed this configuration and found it corresponded to a 'Utopia' model of rigorous quality assurance, since it is impractical in a real-world situation to fulfill the complete set of quality requirements as implied by this model. Because of its lack of practical application, we abandoned this model and experimented on a more 'pragmatic' rigorous quality assurance model, as shown in Figure 13. In this 'pragmatic' model, ANY were used to replace AND in the 'Utopia' model as a weaker form of conjunctive logical operator. Further, OR was used to replace ANY to further relax the satisficing constraint. Under this scheme, one satisfied analysis activity at the bottom level of the framework would be sufficient to fulfill the SR-1 function softgoal at the upper level. The inference engine would attempt to satisfy these SR-1 function softgoals as many times as possible to satisfy the *criticality/risk* criteria softgoals, which would also be attempted as many times as possible for the upper level overall *criticality/risk* softgoals. Finally, the top-level softgoal would be satisfied when either one or both of the overall *criticality/risk*

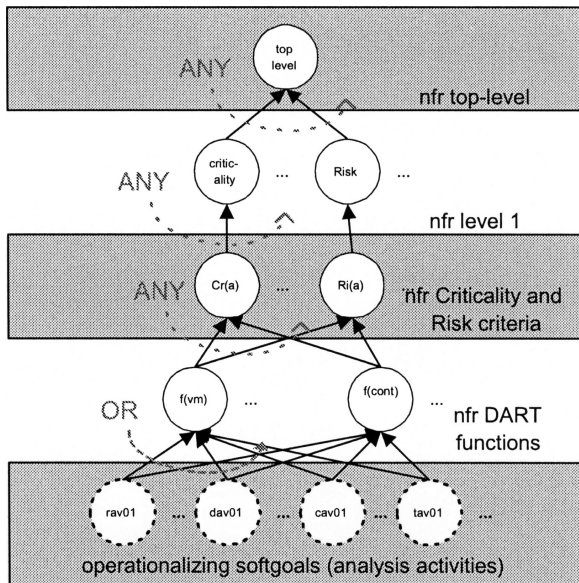


Figure 13. SR-1 framework: 2



softgoals were satisfied. We found this model to be a closer match to the real-world business case; hence it was used in performing the experiments in our studies.

After the framework variations were defined, Monte Carlo simulations were applied to each variant twice, each time with different cost functions. After that, treatment learning was applied to each set of data in finding the most/least favorable treatments. The effects of these treatments are compared with the control situations (i.e. no treatment) in terms of costs and benefits. Results from the experimentations described above are presented in two groups: Tables 11, 12 and 13 for weak quality assurance, and Tables 14, 15 and 16 for 'pragmatic' rigorous quality assurance.

Recall that the class ranking function defined in Table 10 accounted for both benefit and cost with slight preference towards lower cost (e.g.  $Cost=low, Benefit=high$  has a higher ranking than

Table 11. SR-1 framework 1: Percentage distributions of *benefits* and *costs* seen in 10000 runs of Figure 12; no treatment

Cost	Benefit				Total
	vlow	low	high	vhigh	
vlow	34.15				34.15
low		4.02	6.26	5.58	15.86
high		6.2	9.98	8.82	25
vhigh		5.64	8.76	10.59	24.99
Total	34.15	15.86	25	24.99	100

Table 12. More preferred system: SR-1 framework 1: Percentage distributions of *benefits* and *costs* Seen after applying treatments ( $tav09$  of  $tal=y$ ) for a more desirable system

Cost	Benefit				Total
	vlow	low	high	vhigh	
vlow					
low		4.70	7.74	7.30	19.7
high		9.95	16.0	14.2	40.1
vhigh		9.05	14.1	17.0	40.1
Total		23.7	37.8	38.5	100

Table 13. Less preferred system: SR-1 framework 1: Percentage distributions of *benefits* and *costs* seen after applying treatments ( $cav10$  of  $cal=y$ ) for a less Desirable System

Cost	Benefit				Total
	vlow	low	high	vhigh	
vlow					
low					
high		5.16	9.52	8.6	23.27
vhigh		17.32	26.9	32.51	76.73
Total		22.47	36.41	41.11	100

Table 14. SR-1 framework 2: Percentage distributions of *benefits* and *costs* seen in 10000 runs of Figure 13; no treatment

Cost	Benefit				Total
	vlow	low	high	vhigh	
vlow	17.63	2.21	2.67	2.5	25.01
low	3.84	8.76	6.16	6.24	25
high	2.48	8	7.12	7.4	25
vhigh	1.06	6.03	9.05	8.85	24.99
Total	25.01	25	25	24.99	100

Table 15. More preferred system – SR-1 frame work 2: Percentage distributions of *benefits* and *costs* seen after applying treatments ( $dav12$  of  $dal=n$ ) for a more desirable system

Cost	Benefit				Total
	vlow	low	high	vhigh	
vlow	25.35	3.13	3.83	3.6	35.91
low	4.88	11.58	8.4	8.66	33.52
high	2.26	7.65	6.67	7.61	24.19
vhigh		1.56	2.16	2.32	6.38
Total	32.84	23.91	21.06	22.18	100

$Cost=low, Benefit=vhigh$ ). Because of this setting, treatment learner would always recommend treatments that sacrifice a lower benefit for a lower cost. All out result sets reflected this particular class setting.





Table 16. Less preferred system – SR-1 frame work 2: Percentage distributions of *benefits* and *costs* seen after applying treatments (*cau07 of cal=y*) for a less desirable system

Cost	Benefit				Total
	vlow	low	high	vhigh	
vlow					
low	3.4	5.69	1.65	2.11	12.86
high	5.51	15.52	7.71	6.98	35.72
vhigh	4.13	12.95	16.71	17.17	50.96
Total	13.22	34.44	26.08	26.26	100

Several features of these results deserve comment. Consider the results of the experiment on SR-1 framework 1 (weak quality assurance). First, treatment eliminated samples within the *Cost=vlow*, *Benefit=vlow* range, from >34% before treatment to 0% after treatment. Secondly, for the ‘more preferred’ system (where *tav09 of tal=y*), treatment learning drove the sample distributions towards a higher benefit range (*Benefit=high* and *Benefit=vhigh* occupied >76%, as opposed to <50% with no treatment). Third, the distributions of total benefit received after treatments were roughly the same for both the ‘most preferred’ and ‘least preferred’ system. However, the ‘least preferred’ system (where *cav10 of cal=y*) suffered from very high cost (77% of the sample was classified as *Cost=vhigh*), compared to the ‘most preferred’ system (41% of the sample). This effect could be explained by the way the class ranking function was defined. Treatment learning for the ‘more preferred’ system would give recommendations that yield lower cost over lower benefit, whereas it would identify treatments that result in higher cost for the ‘least preferred’ system.

The result for the ‘pragmatic’ rigorous quality assurance (Figure 13) scheme is presented in a similar fashion as the weak one described above. The experimental data is shown in Tables 14, 15 and 16. First of all, treatments for the ‘more preferred’ system (*dav12 of dal=n*) resulted in an increase of samples in the *Cost=vlow* and *Cost=low* range, from >50% to >69%. None the less, the samples within the range *Benefit=vlow* and *Benefit=low* also increased (from 50% to <57%), a clear indication on the proportionality of cost and benefit. On the other hand, treatment for the ‘least preferred’ system (*cav07 of cal=y*) resulted in very high cost

(>85% in the *Cost=high* and *Cost=vhigh* range) compared to that of no treatment (<50%). However, the benefit did not significant increase correspondingly, as >52% in the *Benefit=high* and *Benefit=vhigh* range after treatment, where it was <50% before treatment.

To conclude the SR-1 case study, the following points summarize the results and our observations:

- The use of logic components significantly affects the treatments that TAR2 recommends.
- Variation on cost functions was found to have no observable effect on the resulting treatment recommendations from all SR-1 framework variants studied.
- For all experimentations on the SR-1 framework, TAR2’s treatment recommendations have been ten-way cross-validated<sup>13</sup> and their trustworthiness ensured. In other words, all of TAR2’s treatment recommendation remained stable, in spite of the uncertainties and imprecise factors (as discussed in Section 4.2) within the framework.
- For the results shown in Table 15 (on SR-1 framework 2), TAR2 suggested not doing *dav12* would be beneficial. Our treatment learning method can give advice on which activities should not to be done in order to receive the most preferred outcome.

## 5. RELATED WORK

Inference diagrams (Shachter 1986) (a form of Bayes nets) have been used to sketch out subjective knowledge, then assess and tune each knowledge variable based on available data. This modeling scheme is used by Burgess *et al.* (2001) in evaluating requirements that are candidates to be included in the next release of some software. Bayesian reasoning theory adopts a quantitative, probabilistic approach to solving decision problems. Its algorithm is widely studied and well understood, and the inference tool is widely available. To apply Bayesian theory in defining the softgoal framework, one needs to translate the relationship between operational

<sup>13</sup> Cross-validation is a method of estimating the generalization error based on ‘re-sampling’. In ten-way cross-validation, the entire dataset is randomly split into ten mutually exclusive subsets of approximately equal size. Each subset is being tested on the inducer that is trained by the other nine subsets of data (Kohavi 1995).



softgoals (design alternatives) and NFR softgoals (software quality attributes) into probabilities (e.g. 0.6, 0.8). These numerical values may not be as intuitive as linguistic descriptions ('HELP', 'MAKE'). Moreover, when data required to determine the likelihoods for the calculation of posterior probabilities are not available, certain assumptions have to be made. These assumptions negate the authenticity of the answer obtained by Bayesian inference. Therefore, inference diagrams may not be the best for early lifecycle requirement modeling.

Fuzzy Petri nets, a formalism that combines fuzzy set theory and Petri net theory, is a tool for the representation of uncertain knowledge about a system state. It has been used to reason about uncertainty in robotic systems (Cao and Sanderson 1993). However, within a Petri net, the way membership functions attach to each token and certainty factors associated with transitions can be hard to understand. For business users that are involved in constructing the requirement model, a simple modeling convention is much preferred to a sophisticated one. Thus, using Petri nets may not be very practical.

NASA's DDP (Defect Detection and Prevention) model (Feather and Cornford 2002) is a risk management framework designed to aid decision-making during the earlier phases of advanced technology and system development. It utilizes quantitative analysis, which accepts only one numeric value for each required quantitative value. It is unable to represent the actual situations where uncertainty factors exist.

## 6. CONCLUSION AND FUTURE WORK

The requirement analysis technique proposed in this paper is summarized as follows. First, a model is needed to capture the trade-offs/synergy between software quality attributes and applicable design alternatives. We have adopted the softgoal framework by Chung *et al.* (1999) in modeling such knowledge for analysis. Second, an inference engine is built to automatically execute the text-encoded softgoal framework. Thirdly, we incorporate Monte Carlo simulation to explore the wide range of behaviors in the model, and summarize these behaviors with a treatment-learning tool named TAR2.

From what we have observed in our experiments on both KWIC and SR-1 frameworks, our simulation tool successfully discovered consistent

behaviors within each framework despite various uncertainty factors, and provided treatment recommendations relevant to business concerns. As this approach does not require much concrete domain knowledge, time and expenses dedicated to data collection (e.g. appointments with domain experts, gathering surveys) can be minimized. Moreover, TAR2's treatment results pinpointed the most critical decision towards the problem domain, thus users can focus on this key issue and allot less time in discussing the non-critical ones.

The requirement analysis technique presented in this paper is in its preliminary state, and research is in progress. Much remains to be done to investigate the softgoal framework behavior and refine the proposed technique with more real-world case studies. Specifically, work will be done to determine the sensitivity to benefit and cost functions in the general softgoal framework. With our connection with the NASA IV & V Facility, we are optimistic in receiving more case study materials to conduct further research activities.

## ACKNOWLEDGEMENTS

This research was conducted at the University of British Columbia and West Virginia University, partially under NASA contract NCC2-0979. In part, this work was sponsored by the NASA Office of Safety and Mission Assurance under the Software Assurance Research Program led by the NASA IV & V Facility. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not constitute or imply its endorsement by the United States Government.

## REFERENCES

- Abdel-Hamid T, Madnick S. 1991. *Software Project Dynamics: An Integrated Approach*. Prentice-Hall Software Series.
- Akhavi M, Wilson W. 1993. Dynamic simulation of software process models. In Proceedings of the 5th Software Engineering Process Group National Meeting (held at Costa Mesa, California, April 26–29), Software engineering Institute, Carnegie Mellon University.
- Bratko I. 2001. *Prolog Programming for Artificial Intelligence*. 3rd edn. Addison-Wesley.



- Burgess CJ, Dattani I, Hughes G, May JHR, Rees K. 2001. Using influence diagrams to aid the management of software change. *Requirements Engineering* 6(3): 173–182.
- Cao T, Sanderson AC. 1993. A Fuzzy Petri approach to reasoning about uncertainty in robotic systems. Proceedings of the IEEE International Conference on Robotics and Automation, p. 317–322.
- Chung YM, Nixon Yu, Mylopoulos. 1999. *Non-Functional Requirements in Software Engineering*. Kluwer Academic Publishers.
- Clancey W, Sachs P, Sierhuis M, van Hoof R. 1996. Brahms: simulating practice for work systems design. In *Proceedings PKAW '96: Pacific Knowledge Acquisition Workshop*, Compton P, Mizoguchi R, Motoda H, Menzies T, (eds). Department of Artificial Intelligence.
- Feather M, Menzies T. 2002. Converging on the optimal attainment of requirements. International Conference on Requirements Engineering.
- Feather MS, Cornford SL. 2002. Quantitative risk-based requirements reasoning. *Requirements Engineering Journal* (submitted).
- Harel D. 1990. Statemate: a working environment for the development of complex reactive systems. *IEEE Transactions on Software Engineering* 16(4): 403–414.
- Harrell H, Ghosh L, Bowden S. 2000. *Simulation Using ProModel*. McGraw-Hill.
- NASA Jet Propulsion Laboratory web site. <http://www.jpl.nasa.gov>.
- Kalos M, Whitlock P. 1986. *Monte Carlo Methods, Volume 1: Basics*. Wiley: New York.
- Kelton D, Sadowski R, Sadowski D. 2002. *Simulation with Arena*. 2nd edn. McGraw-Hill.
- Kohavi R. 1995. A study of cross-validation and bootstrap for accuracy estimation and model selection. IJCAI-95.
- Law A, Kelton B. 2000. *Simulation Modeling and Analysis*. McGraw-Hill.
- Martin R, Raffo DM. 2000. A model of the software development process using both continuous and discrete models. *International Journal of Software Process Improvement and Practice*, June/July.
- Menzies T. 2001. *Practical machine learning for software engineering and knowledge engineering; Handbook of Software Engineering and Knowledge Engineering*. World-Scientific.
- Menzies T, Chiang E, Feather M, Hu Y, Kiper JD. 2002. Condensing uncertainty via incremental treatment learning. *Annals of Software Engineering* (special issue on Computational Intelligence).
- Hu Y, Menzies T. 2001. Constraining discussions in requirements engineering via models. Q1
- Mi P, Scacchi W. 1990. A knowledge-based environment for modeling and simulation software engineering processes. *IEEE Transactions on Knowledge and Data Engineering*, September: 283–294.
- NASA IV & V web site. <http://www.ivv.nasa.gov/>.
- Parnas D. 1972. On the criteria to be used in decomposing systems into modules. *Communications of the ACM* 5(12): 1053–1058.
- SWI-Prolog web site. <http://www.swi-prolog.org/>.
- Shaw M, Garlan D. 1996. *Software Architecture: Perspectives on an Emerging Discipline*. Prentice Hall.
- Shachter RD. 1986. Evaluating influence diagram. *Operations Research* 34(6): 871–882.
- Simon HA. 1957. Rational choice and the structure of the environment. In *Models of Man*. Simon HA (ed.). Wiley: New York.
- T.S. Corporation. 2002. IV&V Catastrophic/Critical/High Risk Function List for the Demonstration of Autonomous Rendezvous Technology Project.
- Sterman H. 2000. *Business Dynamics: Systems Thinking and Modeling for a Complex World*. Irwin McGraw-Hill.

**QUERIES TO BE ANSWERED BY AUTHOR (SEE MARGINAL MARKS)**

**IMPORTANT NOTE: Please mark your corrections and answers to these queries directly onto the proof at the relevant place. Do NOT mark your corrections on this query sheet.**

---

Query No.	Query
-----------	-------

---

Q1	?
----	---

---