

A market-driven approach based on Markov decision theory for optimal use of resources in software development

Joost Noppen

TRESE Software Engineering group
Dept. of Computer Science
University of Twente
P.O. Box 217, 7500 AE, Enschede
The Netherlands
noppen@cs.utwente.nl

Victor Nicola

Telematics Systems and Services
Dept. of Electrical Engineering
University of Twente
P.O. Box 217, 7500 AE, Enschede
The Netherlands
vfn@cs.utwente.nl

Mehmet Aksit

TRESE Software Engineering group
Dept. of Computer Science
University of Twente
P.O. Box 217, 7500 AE, Enschede
The Netherlands
aksit@cs.utwente.nl

Bedir Tekinerdogan¹

TRESE Software Engineering group
Dept. of Computer Science
University of Twente
P.O. Box 217, 7500 AE, Enschede
The Netherlands
bedir@cs.utwente.nl

¹ At the time of the writing of the paper this author was visiting the Bilkent University in Ankara

Abstract: Changes in requirements may have a severe impact on development processes. For example, if requirements change during the course of a software development activity, it may be necessary to reschedule development activities so that the new requirements can be addressed in a timely manner. Unfortunately, current software development methods do not provide explicit means to adapt development processes with respect to changes in requirements. This article proposes a method based on Markov Decision Theory, which determines the estimated optimal development schedule with respect to probabilistic product demands and resource constraints. This method is supported by a tool and applied to an industrial case.

1. Introduction

In general, requirements hardly remain constant during the lifetime of a software system. Software systems, should therefore, cope with changes in requirements in cost effective manner. Unfortunately, several case studies have shown that more than 80% of the total cost of software systems is devoted to software maintenance [1]. Adapting software to changing requirements covers a major part of these costs. Changes in requirements do not only have impact on software systems, but also may cause rescheduling of software development processes. For example, if requirements change during the course of a software development activity, it may be necessary to reschedule development activities so that the new requirements can be addressed in a timely manner.

To reduce the maintenance costs often the potential changes in requirements are anticipated as much as possible. Various languages and systems provide techniques to cope with changes within a given context. Each technique assumes that certain aspects of software should be evolvable, while other aspects should be fixed. For example, separation of interface and implementation techniques, such as the Model-View-Controller pattern [2], assumes that implementation should be evolvable but not the interface. Meta programming techniques assume that *base-level* should be evolvable but not the meta interface [3]. Of course, to

support a larger degree of evolution, multiple techniques can be combined with each other. For example, separation of interface and implementation techniques can be combined with meta programming to make interfaces more adaptable, meta programming languages can be made adaptable by introducing meta meta languages, etc.

Changes in requirements do not only have impact on software systems, but also may cause rescheduling of software development processes. For example, it may be necessary to redistribute human resources so that software systems can be delivered in time. This article proposes an approach for optimum distribution of human resources when facing non-deterministic changes in requirements. The approach is based on the Markov Decision Model, which constructs probabilistic state-transition-diagrams that represent market change scenarios. By estimating changes in market demands, available human resources can be distributed optimally at certain points in time.

The remainder of this article is organised as follows. In the next section, motivations for coping with unanticipated software evolution are presented. Section 3 introduces an example case, which is used to describe the problems that are addressed in this article. Section 4 explains the proposed method by applying it to the example case. Section 5 describes the tool that supports this method. Section 6 evaluates the method. In section 7, related-work is presented. Finally, section 8 concludes the article.

2. Unanticipated software evolution

Software evolution usually implies the adaptation of the software system due to certain changes. These changes could include both changes that were expected and that were not. Software evolution can be supported by planning the expected changes as much as possible. Unfortunately, not all changes can be planned for and remain unanticipated. Nevertheless unanticipated changes form an important role in software evolution. To support better software evolution it is therefore required that we provide also techniques to cope with the unanticipated changes.

The required modifications that were unanticipated might be caused by different factors such as insufficient information, lack of diligence, lack of required skill, or simply because it was humanly impossible to predict due to the complex and dynamic nature of the problem being solved. By definition, unanticipated software evolution (USE) is not something for which we can prepare during the design of a software system [1]. The adaptations might require implementing the changed requirements in the application or sometimes it could even include adaptation to the changing environment so that the application keeps satisfying the earlier requirements. Although the extent of software evolution can be very broad, in this paper we primarily focus on software evolution due to unanticipated changes in the requirements.

We categorize unanticipated requirements as follows:

Uncertainty in time of change

In this category, the actual change to the requirements is known at the moment of system design. However, changes in requirements can occur at a different moment in time than was expected (all be it earlier or later). For example major legislative adaptations on insurance laws may have passed the parliament but the actual enactment of this change in the law can still be unclear. These changes will possibly have a considerable impact on the software systems that the insurance companies use.

Uncertainty in content of change

This kind of unanticipation means that the actual occurrence of a certain change in requirements is not known, even though the change itself is known. For example, a new insurance law may have been proposed in the parliament without any clarity on whether this law will come to pass or not. In this case, a possible impact of this change on the related software systems is known but it is not clear whether this change will be finally implemented.

Uncertainty both in time and content of change

In this category, there is no clarity on what will happen. Every change is possible at any given time.

The above classification implies that unanticipation is related both to the content as well as the occurrence of the change to the requirements. As such even if requirements are known in advance the time of occurrence for the requirement change can be fully unanticipated. Including both types of unanticipation will support the software evolution activities. It should be noted that a practical evolution case might be classified under several of these evolution categories.

A number of different approaches already have been proposed that address specific issues related to unanticipated software evolution. These approaches range from extensions of programming languages [4] [5] to adaptable architectures [6] and calculi for describing and reasoning about language properties [7]. The main focus of these approaches is on solving technical issues related to changes that could occur at model, artefact and/or code level. As such, most of these approaches provide solutions to unanticipated software evolution in case of uncertainty in content of changes. To cope with uncertainty in time of changes to requirements we think that unanticipated software evolution must also be considered at the process level. This is because changes in requirements can invalidate process configurations, person hour distributions, etc.. As a complementary approach to the existing software evolution techniques we will provide an approach that is specifically targeted at uncertainty in time of occurrence of the requirements.

3. Problem statement

In the following sub-section, first we will present an industrial example case. Second, by using this example, we will explain some problems that are caused by unanticipated changes. These problems will be addressed in section 4.

3.1 An example: application framework for insurance systems

Consider a software company that produces and maintains software systems to be used by insurance companies. It is assumed that the software company delivers a dedicated software system for each insurance product. In addition, the software company aims to market the same software system to multiple insurance companies worldwide. Insurance companies can tailor these systems according to their specific policies.

Generally, insurance companies have a large customer base, with a variety of insurance products and policies. Of course, the insurance product characteristics are not static but they evolve in accordance with the needs of society. For example, a growing trend in the insurance market is the demand for tailored and personalized insurances.

For the software company, to a certain degree it is possible to deal with the complexity of this evolution by adopting application frameworks, object-oriented and component-oriented techniques. Application frameworks are programs that capture the generic parts of several software systems and encapsulate them in reusable components. Actual software systems can be produced by composition of several of these components². The components are identified by discovering the variabilities and commonalities in the problem domain. For example, Figure 1 shows a feature diagram of insurance products³. A feature diagram represents both the commonality and variability of a product [8]. In this article, to avoid confusion, products that are sold by software and insurance companies are referred to as insurance software systems and products, respectively.

In figure 1, four different parts of insurance products are shown: *Insured Object*, *Coverage*, *Payment* and *Payee*. Each of these parts can be represented by several entities. For example, an insured object can be a *Person* or a *Corporation*. The symbols that are depicted in the legend are used as restrictions on the variability. For instance, an insured object cannot be a

² Application frameworks differ from product-line architectures in that frameworks are defined at programming language level whereas product-line architectures are specified more abstractly. Application frameworks can be implemented as reusable class hierarchies or as pluggable components.

corporation and a person at the same time. Using this specification, a large variety of insurance products can be defined, ranging from bicycle insurances to tailored insurances for large corporations. In addition, new insurance products may be introduced in the future, possibly using advanced personal computing facilities and Internet.

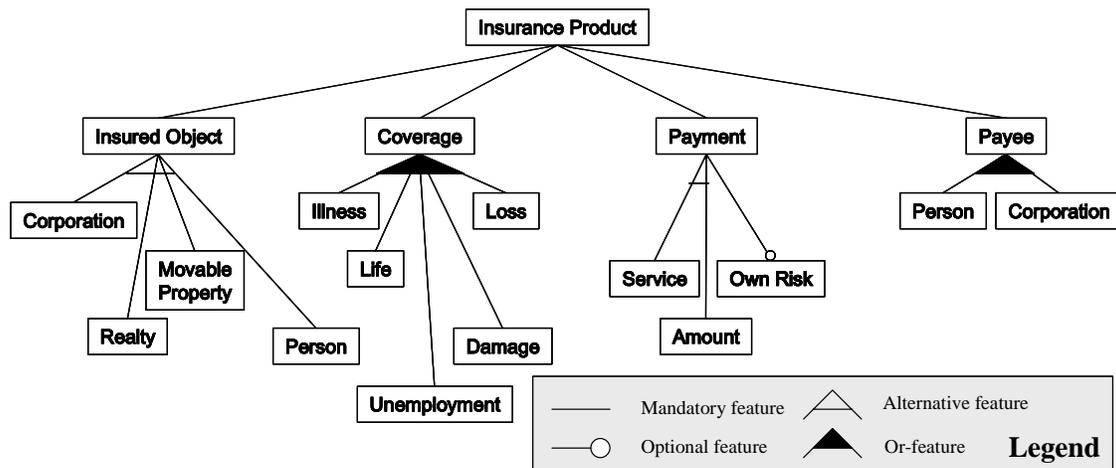


Figure 1. Feature diagram of insurance products

3.2 Resource-scheduling problems due to unanticipated changes

Although framework based approaches provide means to deal with problems that may occur when variability, process-scheduling problems may still exist in case of unanticipated changes. In the following, we will present two problems when dealing with evolution of the insurance software systems demands.

Difficulty in prioritising requirements

Assume that a large set of components have to be implemented in a time span of 2 to 4 years. These components will be used to create several different insurance software systems whenever they are demanded. The components and systems share dependency relations among each other. Therefore, it is crucial to deliver the essential components first, before the

³ This project has been carried out together with Utopics B.V.

insurance companies demand systems that incorporate these components. This requires giving higher priorities to the components that are demanded first.

A number of software development methods suggest prioritisation of requirements [9]. However, these methods do not provide explicit means for determining the priority values with respect to the changes in demands.

Several publications and tools are proposed for supporting decision-making processes. Most of the approaches are based on the Analytic Hierarchy Process (AHP) method [10] [11] [12]. In the AHP method, a pair wise comparison is made between two entities that should be ranked with respect to each other. Subsequently the pair wise ordering is normalized. The AHP method is applied to various kinds of problems such as resource allocation, risk assessment, etc. However, the proposed approaches assume a fixed set of requirements and therefore they are not directly suitable for dealing with unanticipated software evolution.

Configuring software development processes for delivering products in time

After prioritising the components, the project managers should be able to configure the implementation trajectory accordingly. Since there may be many options for the implementation effort, preferably, ideally the project managers should be able to compare all the relevant options and select the best configuration that fulfils the timing and resource constraints. For the example case, this means that the expected optimal planning of resources for a given time span should be determined to maximise profit in selling insurance software systems.

Configuring software processes with respect to available resources is not new [13]. However, in contrast to existing methods, also the demands for future systems and changes in requirements must be considered.

4. The approach

4.1 Possible Approaches

In the following sub-section a short description of a number of alternatives for decision making optimizations.

Real options

In [14] the approach of real options is described as a means of supporting decision making in software design. In real options theory each decision that can be made is considered an option and has its own *added value*. With this added value the actual gain of choosing this option is signified. In addition to the options that are identified also scenarios are identified and their change of occurrence. From these scenarios an event tree with which it is possible to determine the expected added value for choosing one of the available options. The underlying mathematical theory for the real options approach is based on dynamic programming, also known as Markov Decision Theory.

Quality Function Deployment

Quality Function Deployment [15] is a well-known tool for analysis of importance of decisions in software design. One of the particular uses for this approach can for instance be found in resolving conflicting requirements. In Quality Function Deployment a distinction is made between the requirements from the customer side (called the voice of the Customer) and the parameters of design (called the voice of the Engineer). The relationship between these two voices is described in the Relationship Matrix. QFD can balance and prioritize the customer needs and design parameters based on the relationship matrix.

Markov Decision Theory

Markov Decision theory is a mathematical theory for determination of optimal decisions at a given point in time. This is done by determination of the decisions that can be made and the definition of scenarios that can occur. For each of the scenarios the probability of occurrence is also described. These two inputs are then combined into a tree structure, which is commonly termed a Markov Decision Problem. The optimality of the decisions is then calculated based on expectation values of rewards for each decision. Markov Decision Theory is also known as Dynamic Programming.

In our approach we have elected Markov Decision Theory as the underlying mathematical theory since it provides a sound mathematical framework for solving the type of problems described in section 3. Real options theory is related to our approach since the same mathematical basis is used, and globally the approaches share similarities. However, our approach offers a more clear-cut solution for resource allocation problems since our application of the mathematical theory explicitly addresses the typical inputs for these problems.

Quality Function Deployment has a less formal mathematical basis than Dynamic Programming and as such is easier to understand in its application. However, with the large amount of inputs and the probabilistic character of the problem QFD does not provide the proper tools for efficient determination of optimal development trajectories.

4.2 The Method

The options in scheduling software development processes depend on the possible software system demands. The following three steps are used to determine the most profitable option:

1. Determine probabilistic estimations for market, cost and profit changes;
2. Specify the possible market scenarios that can occur;

3. Evaluate all possible schedules with respect to the scenarios and select the most cost effective schedule.

This method should be flexible enough to cope with changes in the estimations at any given moment in time.

4.3 Markov Decision Model

To achieve the objectives stated in the previous section, the problem needs to be formalized mathematically. In our approach the optimisation problem is formalized using Markov Decision Theory [16]. Below the formalization is defined for this specific problem.

- Consider a time horizon of length I time intervals (steps). Let i ($1 \leq i \leq I$) be the index of the time interval with i steps to go. That is, the time horizon starts at the beginning of interval I and ends at the end of interval 1.
- Let J be the total number of artefact types that may be produced to satisfy market demands during the entire time horizon. Market orders for artefacts are placed at the beginning of each time interval, and are lost if not satisfied by the end of the same time interval. At the beginning of interval i ($1 \leq i \leq I$), the demand for artefact j is a random number D_{ji} characterized by the (discrete) distribution F_{ji} with a finite mean given by d_{ji} . The price of one unit of artefact j at the end of interval i is given by r_{ji} . Denote by \mathbf{d}_i and \mathbf{r}_i the *demand* and *price* vectors with elements d_{ji} and r_{ji} ($1 \leq j \leq J$), respectively.
- At the beginning of interval i ($1 \leq i \leq I$), artefact j ($1 \leq j \leq J$) requires a total of w_{ji} person-hours to be completed (i.e., to make it ready for marketing). Denote by \mathbf{w}_i the remaining work vector with elements w_{ji} ($1 \leq j \leq J$). Note that $w_{ji} = 0$ if artefact j is completed at some interval before the beginning of interval i .

- At the beginning of interval i ($1 \leq i \leq I$) we have M_i available person-hours that can be distributed arbitrarily among the set of (uncompleted) artefacts. Let a_{ji} be the fraction of available person-hours allocated to artefact j during interval i , and denote by \mathbf{a}_i the allocation vector with elements a_{ji} ($1 \leq j \leq J$). (Note that for any i , $\sum_{j=1}^J a_{ji} \leq 1$.) Then the amount of person-hours allocated to artefact j during interval i is given by $a_{ji} M_i$, with $a_{ji} \leq w_{ji} / M_i$. Note that if artefact j is completed at some interval before i , then $a_{ji} = 0$. It follows that the total work allocation (cost) during interval i is given by $M_i \mathbf{a}_i^T \mathbf{e} \leq M_i$, and that the remaining work on artefact j at the end of interval i (or at the beginning of interval $i - 1$) is given by

$$w_{j,i-1} = w_{ji} - a_{ji} M_i = w_{ji} - \sum_{k=1}^i a_{jk} M_k, \quad 2 \leq i \leq I,$$

where w_{ji} is the total work required to complete artefact j at the beginning of the time horizon.

- Let the vectors \mathbf{d}_i and \mathbf{w}_i define the state of the Markov decision process at the beginning of interval i . Then, the expected immediate return R_i at the end of interval i depends on the state vectors $(\mathbf{d}_i, \mathbf{w}_i)$ and the allocation vector \mathbf{a}_i ; specifically

$$R_i(\mathbf{d}_i, \mathbf{w}_i, \mathbf{a}_i) = \sum_{j=1}^J \mathfrak{S}(\{w_{ji} - a_{ji} M_i = 0\}) d_{ji} r_{ji}, \quad 1 \leq i \leq I,$$

where the indicator function $\mathfrak{S}(\cdot)$ equals 1 if its argument is true (i.e., artefact j is completed at or before interval i), otherwise it is equal to 0.

- Let $Y_i(\mathbf{d}_i, \mathbf{w}_i)$ be the maximum total expected reward at the end of interval i , then the optimality equations can be written as follows:

$$Y_1(\mathbf{d}_1, \mathbf{w}_1) = \max_{\mathbf{a}_1} [(R_1(\mathbf{d}_1, \mathbf{w}_1, \mathbf{a}_1) - M_1 \mathbf{a}_1^T \mathbf{e})],$$

$$Y_i(\mathbf{d}_i, \mathbf{w}_i) = \max_{\mathbf{a}_i} [(R_i(\mathbf{d}_i, \mathbf{w}_i, \mathbf{a}_i) - M_i \mathbf{a}_i^T \mathbf{e}) + Y_{i-1}(\mathbf{d}_{i-1}, \mathbf{w}_{i-1})], \quad 2 \leq i \leq I.$$

The above recursion can be applied to determine the optimum allocation vectors \mathbf{a}_i for $i = 1, 2, \dots, I$, respectively.

4.4 Scheduling the implementation of the framework

In this section, we will apply our method to the example case. To this aim, first insurance software system demands are estimated. Second, an insurance software system demand scenario is presented. Finally, based on this scenario, the optimal alternative in scheduling the implementation of the framework is determined.

Step 1: Estimations

As explained in the following sections, at this stage, three kinds of data are defined. First, expected products and sub-products are specified. Second, the insurance software system demands are estimated. Third, the available human resources and costs are specified.

Specification of the product demands

Assume that according to the marketing study there will be a substantial demand for *Unemployment Insurance* systems, *Illness Insurance* systems and *Illness Insurance with Own Risk* systems. These should be assembled through the composition of the necessary framework parts. In Figure 5, the artefacts *framework parts* and *systems* are represented as rectangles. The directed relations among the rectangles indicate assembly dependence relations. For example, *Unemployment Insurance* systems are assembled from the framework parts *Insured Person* and *Unemployment Coverage*. The framework parts *Insured Person* and *Unemployment Coverage* are, in turn, assembled from *Insured Object* and *Coverage*, respectively.

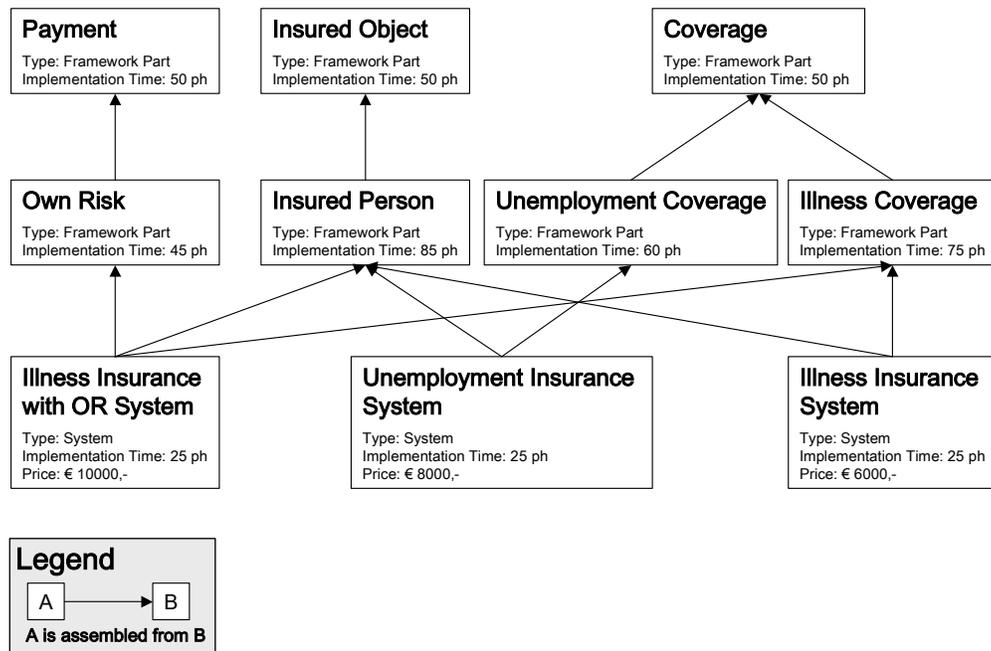


Figure 2. Dependencies among insurance software systems and framework parts

As shown in Figure 2, every artefact has the attribute *Implementation Time*, which indicates the amount of time necessary to implement that artefact. The products have an additional attribute, *Selling Price*, which indicates the estimated market price of the system.

Market Demand Scenarios

In general, market demand scenarios cannot be determined with certainty. Therefore probabilistic estimations need to be made. Typically, standard normal distributions are used for this purpose. Consider, for example, the following demand estimations:

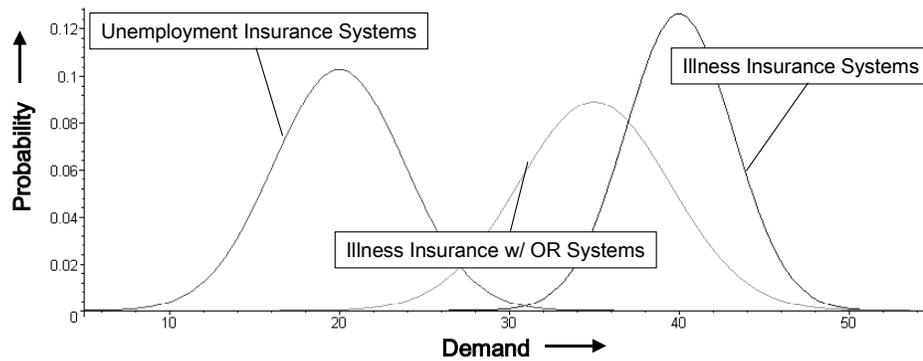


Figure 3. Market volume estimations for three kinds of insurance software systems

In Figure 3, curves represent from left to right the estimations for market demand for *Unemployment Insurance* systems, *Illness Insurance with Own Risk* systems and *Illness Insurance* systems, respectively. The curve for *Unemployment Insurance* systems varies between approximately 16 and 24 units and concentrates around 20 units. The second distribution, which depicts the demand estimation for *Illness Insurance with Own Risk* systems, varies in demand between approximately 30 and 40 units. For this software system, a demand of 35 units will occur with the highest probability. Finally, the third distribution varies between approximately 38 and 43 units and concentrates around 40 units. For this example we will take the mean values for ease of computation. However, the computations can be made more accurate by taking the entire deviation interval.

Human Resource Definitions

Assume that two persons will be available in every week for up to 50 person hours each. It is also required that a person cannot be assigned to multiple tasks in one week. The cost per person hour is fixed to € 100.

Step 2: Modelling the product demand scenarios

To compute the optimum schedule, the Markov Decision Model as described in section 4.2 has to be discretized. For the distributions this is done by taking the mean values. The mean values of the distributions of three insurance products are listed in the following:

- *Illness Insurance systems* 40 units
- *Unemployment Insurance systems* 20 units
- *Illness Insurance with Own Risksystem* 35 units

In addition, for this example, we have discretized the amount of person hours that can be allocated as per week and per employee. Since a person can only work on one task per week, the smallest amount of person hours that can be assigned is 50.

Step 3: Selecting the most cost effective schedule

Finally, all possible schedules are evaluated for every single scenario. For this purpose, schedules are modelled as a *sequence of decisions*, where each decision is taken at the beginning of a week. For instance, in the first week, it is possible to assign 100 person hours to any possible artefact. When a system is completed and there is a demand for it, it can be sold with a profit. Finding the best schedule is an optimisation problem as defined below:

Given

- *System demand scenarios with probabilistic change dependencies;*
- *A set of available person hours with a given cost price and minimal contract duration;*
- *A set of framework parts and/or systems with expected implementation time;*
- *Estimated profit of a system when it is brought to the market,*

Find a schedule, which delivers the maximum profit.

The method is supported by a tool, which is presented in section 5. Finding the optimal schedule is computed automatically and is based on Markov Decision Theory, which was presented in section 4.2.

Results

The result of the optimisation algorithm can be displayed as a table. In our example case, the tool provides several equivalent schedules that will result in the same expected profit. Table 1 shows one of the possible schedules.

	Coverage	Insured Object	Payment	Illness Coverage	Unemployment Coverage	Insured Person	Own Risk	Illness Insurance System	Unemployment Insurance System	Illness Insurance /w OR System	Max. Expected Profit
Week 1	50	50	-	-	-	-	-	-	-	-	€ 14000,-
Week 2	-	-	-	-	100	-	-	-	-	-	
Week 3	-	-	-	-	-	100	-	-	-	-	
Week 4	-	-	-	-	-	-	-	100	-	-	
Week 5	-	-	-	-	-	-	-	100	-	-	

Table 1. An optimum allocation of resources for maximum profit

Each row in Table 1 corresponds to a week in the project time span. The elements of a row contain the allocation of person hours for that week. For instance, in the first week, 50 person hours are assigned to the *Coverage* artefact and 50 person hours to the *Insured Object* artefact. In addition, the model gives an indication of the *maximal expected* profit. Based on the given product demand scenarios and constraints, *Unemployment Insurance System* should be delivered at the end of the fourth and fifth week. Therefore, the relevant artefacts for these insurance software systems should be implemented first.

4.5 A more accurate market model

The example presented in the previous section can be extended in various ways. For example, in a more realistic scenario, the probabilities and available resources may change in time due to illness or holidays. In following, as an illustrative example we will assume time dependent mean probability values and resource availability functions.

Market Demand Scenarios

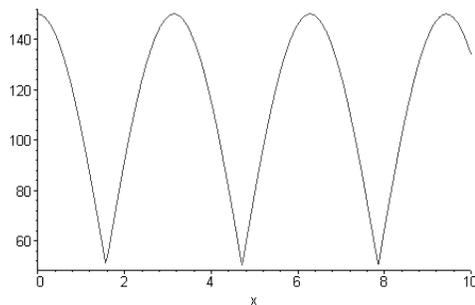
The demands for each insurance software system may change in time. To simulate this, the market demand scenarios should include changing mean demands for every relevant project time step. For example, in the table below, we make an assumption by modelling changes per week:

	Illness Insurance System	Illness Insurance with OR System	Unemployment Insurance System
Week 1	40	35	20
Week 2	30	42	31
Week 3	40	35	20
Week 4	30	42	31
Week 5	40	35	20

Table 2. Mean values of system demands per week

Human Resource Definitions

It is also possible to change the available of human resources in time. Consider for example Figure 8, which assumes a periodic change in the available human resources.



This graph is approximated in the actual calculations by the following function⁴

$$50 + 50 * \text{Round}(\text{Abs}(2 * \cos(t)))$$

Figure 4. Periodic changes in human resources

In this figure, it is assumed that the availability of human resources shows a periodic character. The presented changes in product demands and human resources should be considered as examples only. It is possible to use different functions and/or any meaningful data obtained for example from the past experiences.

Results

Based on the time varying data, Table 3 displays one of the expected optimal schedules:

	Coverage	Insured Object	Payment	Illness Coverage	Unemployment Coverage	Insured Person	Own Risk	Illness Insurance System	Unemployment Insurance System	Illness Insurance /w OR System	Max. Expected Profit
Week 1	50	50	-	-	-	-	-	-	-	-	€25000,-
Week 2	-	-	-	-	100	-	-	-	-	-	
Week 3	-	-	-	-	-	100	-	-	-	-	
Week 4	-	-	-	-	-	-	-	-	150	-	
Week 5	-	-	-	-	-	-	-	-	100	-	

Table 3. An optimum allocation of resources in changing system demands and resources

This table shows that based on the assumptions, *Unemployment Insurance Product* is expected to deliver most profit. Therefore, based on the table, the relevant artefacts for this product should be implemented first. Note that in the first step only 100 person hours of the available 150 are distributed. Apparently within the current situation it is not wise distribute the additional 50 person hours as it will lower the expected profit.

5. Tool support

The optimisation process is very labour intensive and therefore automatic support is necessary. In addition, tool support is necessary for modelling artefacts, market, and available

⁴ *Round* represents the function that rounds real numbers to the nearest natural number, and *Abs* the function that returns the absolute value of a given number.

resources. The architecture of our tool is shown in Figure 5. Here, the models and processes are represented as rectangles and ellipses, respectively.

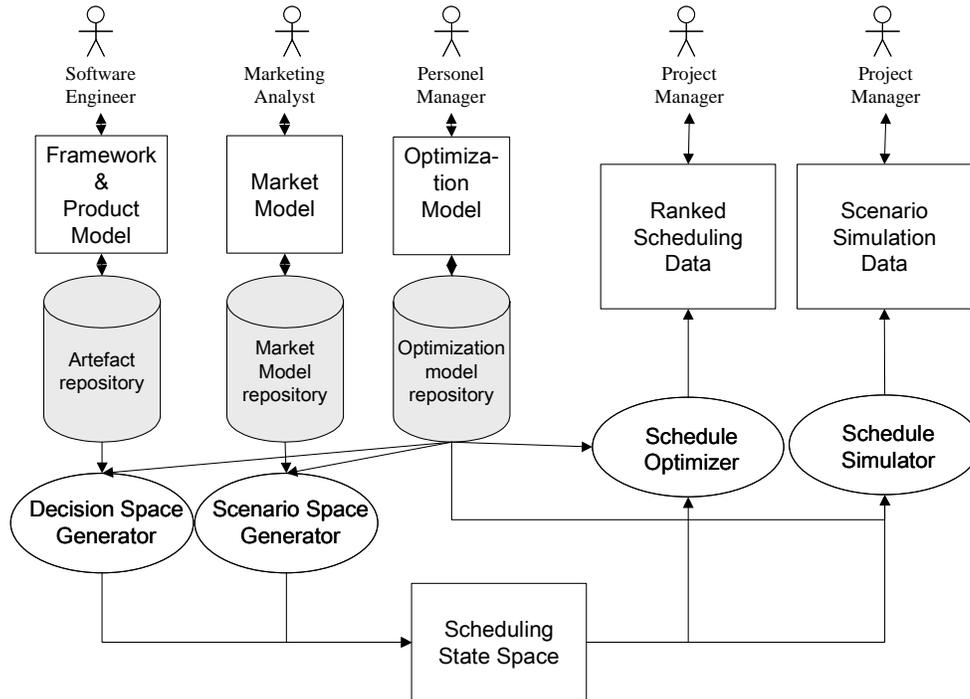


Figure 5. The tool architecture

Artefacts (framework parts and software systems), insurance software system demand scenarios and resources are modelled by the software engineer, marketing analyst and personnel manager, respectively. The Personnel Manager also defines the goals of the optimization problem, such as minimization of cost, maximization of profit, etc.

The process *Decision Space Generator* retrieves the necessary information from the artefact repository and the process *Scenario Space Generator* retrieves the necessary information from the Market Model Repository. Together they generate the data *Scheduling State Space*. The next step is the determination of scheduling advice. The process *Schedule Optimizer* takes *Scheduling State space* as input and generates the data *Ranked Scheduling Advice* as output. This data can be presented to the project manager in various formats, for example, as a table shown in Table 3. In addition the project manager can run simulations by using the *Schedule Simulator*, which results in the data *Scenario Simulation Data*.

First of all the *FrameWork & Product Model* can be defined in the *Parameter Modeller*. In this tool all the parts can be modelled, which the system engineer can directly influence. In figure 6 the component definition of section 4.3 is modelled using the *Parameter Modeller*.

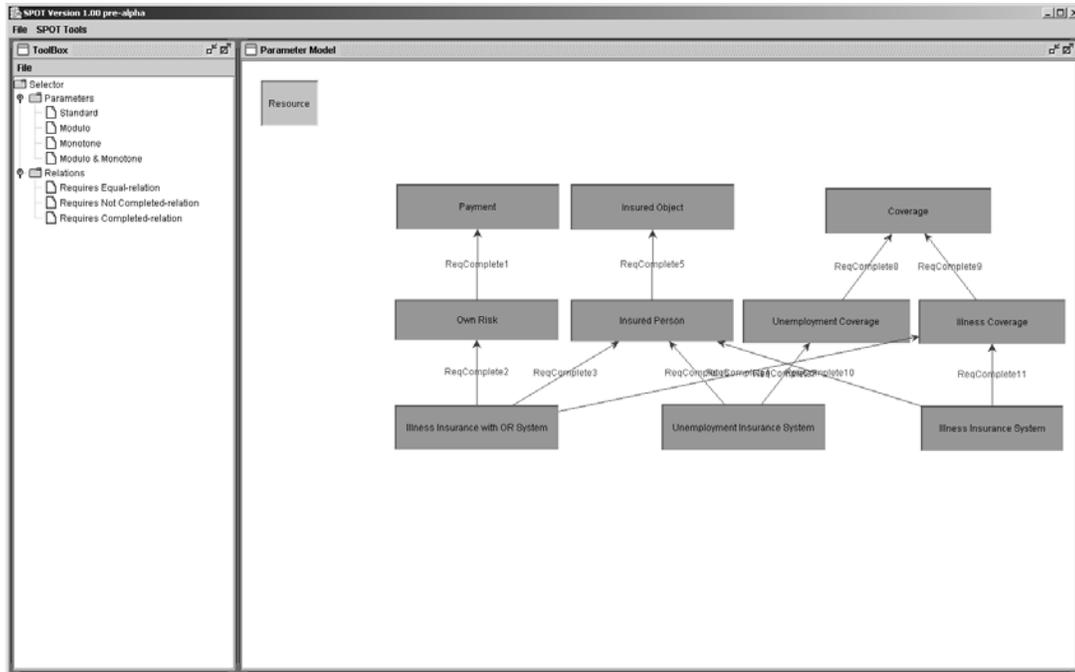


Figure 6. Modelling the insurance framework components

As shown in Figure 6, the sub-window on the left side shows the modelling tools. It is possible to define different types of parameters, corresponding to framework parts and actual products.

During the course of a project, a number of different situations may occur. To increase flexibility, if desired, the scheduling advice can be given on a per situation basis. The tool can simulate the situations that can occur for gaining a better insight into these situations. In figure 7 the *Simulation Tool* is displayed.

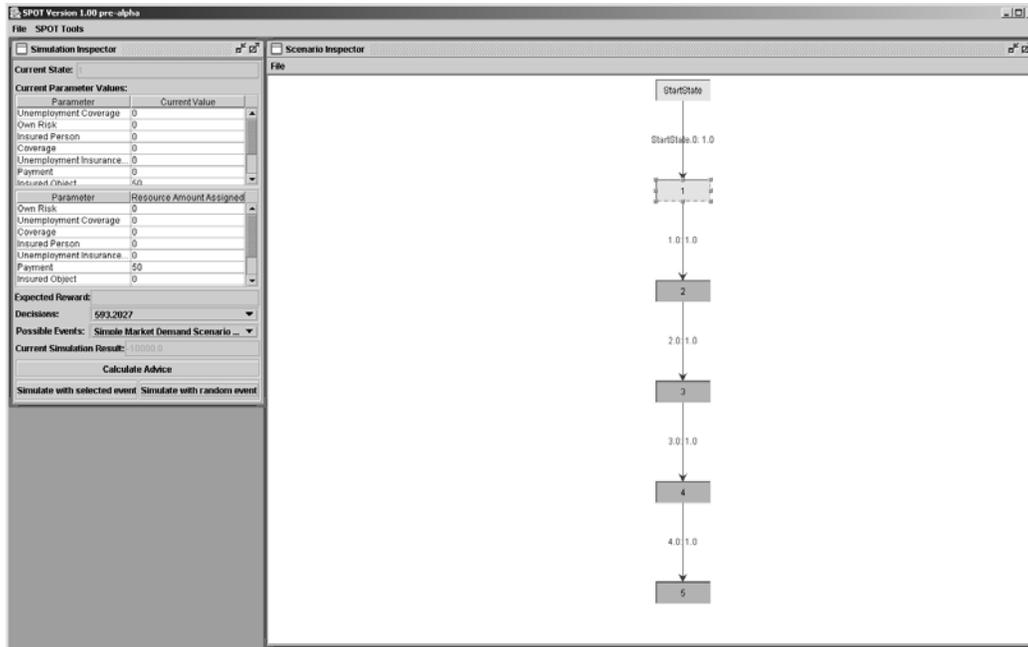


Figure 7. Running a simulation

As shown in Figure 7, the Simulator Tool consists of two sub-windows. The left sub-window is the simulation inspector and displays information on the current state of the simulation, such as profit result, amount of work done, etc.

The right sub-window displays the states that can be reached during the simulation. This graph is the scenario graph that models all possible market situations. The graph in figure 7 depicts a graph for which each state has only one possible next state and therefore represents only one scenario. While running the simulation the tool will color the visited nodes yellow, so in figure 7 already two nodes have been visited.

The tool has been implemented on the Java Platform as a part of the SP²OT Tool Suite, and is still under development. More information on the tool suite and the tool itself can be found at:

<http://janus.cs.utwente.nl/twiki/bin/view/Spot/WebHome/>

6. Evaluation

In section 3.2, prioritising changing requirements and configuring software development process accordingly were identified as two important problems. Given probabilistic product

demand scenarios and resources, we have presented a method and a set of tools that can help project managers in determining the estimated optimum process development schedule that delivers the best profit. In the following, we will evaluate our approach with respect to the following concerns:

- ***Support for unanticipated evolution***

In general, unanticipated software evolution focuses on changes in software requirements that were not foreseen during the design of a software system. In our specific case we assume that changes in software requirements are known, but their occurrence is not certain. That is to say that we do not deal explicitly with uncertainty in both content and time of change as described in section 2. However, our approach can be adopted for phenomenal potential evolution and temporal potential evolution. Our stochastic model explicitly reasons about the occurrence of requirements over a time span, as such we can deal with uncertainty in time of change. In addition, every requirement is tagged with a probability and the model computes the occurrence and the change of these probabilities, and therefore we can deal with phenomenal potential evolution. Of course, more experimentation is necessary to determine the practical use of our approach.

- ***Support for planning the software development process***

The presented approach provides a valuable input to the resource planning process in which the scarce resources have to be planned in a time span. The tool can be used at various levels of details. In addition, the tool provides flexibility in considering new data, for example, which may be available due to better market estimates. The tool will recalculate the optimal schedule from the time that new data is considered relevant. The probabilistic estimations cannot be always guaranteed to be accurate. Nevertheless, our tool may force the managers to explicitly consider various scenarios and experiment with “what-if” conditions over the scenarios. This may help the managers to gain better insight in planning the software development processes.

- ***Tool support***

Optimum resource allocation is a very difficult task to carry out manually because (a) various different input data are required; (b) requires evaluation of a large space of alternatives, and (c) input data may change continuously and therefore evaluation have to be repeated frequently. We have developed a set of tools for: (a) defining models, (b) specifying probabilistic estimations, (c) computing optimal schedules at various levels of details, and (d) displaying and inspecting the results in various formats. We have already started experimenting with the tool in different applications. By providing the tool for public use, we hope to obtain the necessary feedback necessary for improving the tool where necessary.

- ***Support for vague market characteristics***

In general, short-term market estimations can be made with a reasonable accuracy. For a time span longer than 2 years, simple probabilistic estimations will be less reliable. This problem can be partially addressed by using adaptive learning techniques so that estimations can be improved continuously and by defining more accurate market models, for example, based on fuzzy logic techniques.

- ***Support for simulation and educational purposes***

One important perspective for our tool is to use it for simulation and educational purposes. Because the tool essentially generates all the possible and relevant situations during a project, the state space can be used to examine specific project situations. In addition, by fixing market expectations at specific time points, scenarios can be simulated, for instance for determining worst-case scenarios. We are currently cooperating with a large software company in adopting the tool for consulting activities and in management courses, etc.

- ***Human Resource Model***

Our approach is sensitive to the accuracy of the human resource model. In addition, it is assumed that every employee can perform any task, if the time permitting. A more realistic approach would be to categorise human resources with respect to their skills. This adds an extra level of dependencies that needs to be considered when scheduling human resources.

We are currently working on a more realistic human resource model, which improves the model properties based on the experiences obtained in the past.

- ***Artefact Model***

In the current approach, artefacts can only be produced when all the required sub-artefacts have been delivered. By introducing a more fine-grained model, we would like to minimize the dependencies among the artefacts and thereby improve parallel allocation of resources.

- ***Scalability of the approach through managing state explosions***

To calculate the optimal result, Markov Decision Models require all relevant states. This can cause very large state spaces and lengthy computations. This is known as the *Curse of Dimensionality* [17]. To address this problem, within the field of Markov Decision Theory, various research activities have been carried out. Examples are state space minimization techniques by assessing the relevancy of states and by making trade off between accuracy of the result and computation size, or optimisation for parallel computing [18][19][20]. We are currently experimenting on various state-space reduction mechanisms.

7. Background and Related Work

7.1 Software Process Configuration Management

To cope with the constantly changing customer requirements, software products and software process must be re-configured frequently. Software configuration management aims to manage the evolution of a product. Our approach can be seen as complementary to software configuration management techniques.

Software configuration management is a broad domain and covers the entire life cycle. Further, software configuration management processes can be defined for specific domains as well. For example, in [13] an approach is presented for resource management for distributed multi agent systems. This approach takes into account a wide range of resource and entity types. By providing different types of relations between entities and resources, the actual

situation can be described and optimised accurately. The approach in [13], however, is not aimed to optimise with respect to future changes that are uncertain. Our approach aims to schedule the available resources while considering probabilistic changes in requirements and project context.

In addition to software configuration management, several researchers have focused on the configuration of processes. The basic issues are to choose the right development process and to align process configuration whenever necessary. In [21] and [22] it is argued that software development processes should be developed in much the same way as software. The common assumption is that processes should be considered as products and be configured with respect to product quality goals. Our approach can be considered complementary to these approaches. Our focus is, however, to cope with changes in requirements and project context, which is generally missing in the conventional software process configuration approaches.

7.2 Requirements Engineering

Prioritisation of requirements with respect to implementation is not new. There have been various research activities on this topic. However, research activities generally focus on fixed requirements; prioritisation is realized after the requirements are determined.

The work described in [23] focuses on prioritisation of software requirements with respect to the quality of decision-making. The basic assumption here is that new requirements should either be accepted or rejected. By analysing and simulating the acceptance/rejectance rate, the decision quality can be improved and only relevant requirements are selected for further consideration.

The difference with our approach lies in the fact that the approach in [23] acts as a filter and results in a set of requirements that should be implemented. The remaining relevant requirements are also considered equally important. However, these requirements still need to be prioritised with respect to implementation and resources, since there still might be dependencies among them. Also the approach focuses on the best requirements set for the

next release, while our approach is more aimed at finding optimal schedules for implementation with respect to changes along the time span of a project.

In [10] a methodology is proposed, which prioritises requirements based on a cost-value analysis. For each requirement, the relative cost and value are determined using the Analytic Hierarchy Process. This model then allows comparing the requirements based on these properties, so that one can be ranked over the other with respect to these two aspects. By considering the requirements of engineers, customers, users and software engineers an accurate requirements model can be made.

The proposed methodology in [10] focuses on prioritisation with respect to value and cost, and with that it aims to satisfy market demands in general. The methodology does not explicitly address the delivery constraints that are imposed during software development process. Also the model does not support analysis of requirements that share dependencies. In our approach this is specifically addressed by the framework dependencies.

7.3 Optimisation Techniques

Many different optimisation techniques have been defined, most of them having a mathematical origin. Generally speaking an optimisation model consists of four parts: a subject to be optimised, the options to be considered, a comparison criterion, and an ideal situation description (or goal). The possible options are evaluated with respect to their goal using the criterion. Dynamic programming [16, 17] is one of these techniques but many others exist. For instance, there are learning-based optimisation models such as neural networks [24] or genetic algorithms.

Most optimisation techniques are based on evaluation of the available options. Brute force approaches tend to evaluate every single option and then select the optimal one. However, this soon leads to incomputable problems. Therefore most optimisation techniques aim at “smart” evaluation of options, thus reducing computation complexity. A vital step in the

optimisation process is to map the problem to a mathematical formalism. By choosing a smart mapping, the complexity of the problem can be kept at an acceptable level.

8. Conclusions

It should be noted that the management decision problem is indeed a very challenging problem and it is difficult to claim that we have provided the complete and final solution. To fully prepare for unanticipated software evolution seems to be impossible by definition.

Nevertheless we have shown that even unanticipated software evolution can be managed to a certain degree. For this, we have made an explicit distinction between unanticipated software evolution due to uncertainty in content and uncertainty in the time of occurrence of changes to requirements. Conventional software evolution approaches have basically focused on evolution in content of the requirements and several useful techniques have been proposed. As a complementary approach we have focused on coping with evolution of changes to requirements in time.

The presented approach represents the possible development activities as a state space using Markov decision theory. By applying dynamic programming techniques the optimal schedules of software development activities are determined. The required changes are modelled using probabilistic estimations on product demands, profits and available resources. Due to the large state space it is very hard to manually perform this approach. Therefore, we have implemented a tool, which supports the modelling, the selection and the evaluation of the changes to requirements and the allocation of human resources.

The approach and the related tool explicitly forces to think about priorities and risks of decision making with respect to changes in requirements. To improve the scalability of the approach our future work will include the refinement of the human resources model and the integration of more refined state-space reduction techniques.

9. Acknowledgements

We would like to thank the anonymous reviewers for their valuable feedback on the initial version of this article.

10. References

- [1] Kniesel, G. Noppen, J. Mens, T. Buckley J. *The first international workshop on unanticipated software evolution* Workshop report for the first int. workshop on unanticipated software evolution at ECOOP 2002
- [2] Lethbridge, T. Laganière, R. *Object-Oriented Software Engineering* McGrawHill, 2001, ISBN: 0-07-709761-0
- [3] Yonezawa, A. Matsuoka, S. *Metalevel Architectures and Separation of Crosscutting Concerns* Third International Conference, REFLECTION 2001 Kyoto, Japan, September 25-28, 2001 Proceedings, Springer Verlag, ISBN 3-540-42618-3
- [4] Duggan, D. Wu, Z. *Adaptable Objects for Dynamic Updating of Software Libraries* Workshop paper for the 1st USE workshop at ECOOP 2002
- [5] Dmitriev, M. *Hotswap technology application for advanced profiling* Workshop paper for the 1st USE workshop at ECOOP 2002
- [6] McGurren, F. Conroy, D. *X-ADAPT: an architecture for dynamic systems* Workshop paper for the 1st USE workshop at ECOOP 2002
- [7] Anderson, C. Drossopoulou, S. *An imperative object based calculus with delegation* Workshop paper for the 1st USE workshop at ECOOP 2002

- [8] Kang, K. Cohen, S. Hess, J. Novak, W. Peterson, A. *Feature Oriented Domain Analysis (FODA) Feasibility Study* CMU/SEI-90-TR-21 ESD-90-TR-222, Software Engineering Institute, Carnegie Mellon University, Pittsburgh, November 1990
- [9] Jacobson, I. Booch, G. Rumbaugh, J. *The Unified Software Development Process* Addison Wesley, 1999, ISBN: 0-201-57169-2
- [10] Karlsson, J. Ryan, K. *Prioritizing Requirements Using a Cost-Value Approach* IEEE Software September/October issue, pp. 67-74.
- [11] Salo, A. Hämäläinen, R. *On the measurement of preferences in the analytic hierarchy process* Journal of multi-criteria decision analysis, vol6, 309-319, 1997
- [12] Forman, E. Sally, M. *Decision by Objectives* World Scientific, 2001, ISBN: 981-02-4142-9
- [13] Podorozhny, R. Staudt Lerner, B. Osterweil, L. *A rigorous approach to resource management in activity coordination* University of Massachusetts, Amherst MA 01003, USA, 1999
- [14] Sullivan, K. Chalasani, P. Jha, S. Sazawal, V. *Software Design as an Investment Activity: A Real Options Perspective*. In L. Trigeorgis (ed.) Real Options and Business Strategy, Risk Books. ISBN: 1899332472
- [15] Tran, T. Sheriff, J. *Quality Function Deployment: an Effective Technique for Requirements Acquisition and Reuse 2nd* IEEE Software Engineering Standards Symposium, Aug. 21-25 1995, Montreal, Canada, p. 191 ISBN 0-8186-7137-8

- [16] Puterman, Martin, *Markov decision processes, discrete stochastic dynamic programming*, 1994, Wiley-Interscience, ISBN: 0-471-61977-9
- [17] Bellman, R. *Adaptive Control Processes: A Guided Tour*, Princeton University Press, Princeton, NJ, 1961.
- [18] Chung, S. Hanson, F. Xu, H. *Parallel stochastic dynamic programming: finite elements methods* Linear Algebra and Applications 172, pp. 197-218, 1992
- [19] Larson, R. *Dynamic programming with reduced computational requirements* IEEE Transactions on Automatic Control AC-14, pp. 135-143 (1965)
- [20] Larson, R. *State Increment Dynamic Programming*, American Elsevier, New York, NY, 1968
- [21] Osterweil, L. *Software Processes are software too, revisited: An invited talk on the most influential paper of ICSE 9* Proc. 19th Intl. Conference on Software Engineering, Monterey, CA, IEEE Press, 2-13
- [22] Osterweil, L. *Architecting Processes are key to software quality* International Workshop on the Role of Software Architecture in Testing and Analysis, Marsala, Italy 1998
- [23] Regnell, B. Karlsson, L. Höst, M. *An analytical model of requirements selection quality in software product development* Published at SERPS'02
- [24] Haykin, S. *Neural Networks: A comprehensive foundation* Prentice Hall 1998

11. List of Tables

	Insured Object Coverage	Insured Object Payment	Illness Coverage	Unemployment Coverage	Insured Person	Own Risk	Illness Insurance System	Unemployment Insurance System	Illness Insurance /w OR System	Max. Expected Profit
Week 1	50	50	-	-	-	-	-	-	-	€ 14000,-
Week 2	-	-	-	100	-	-	-	-	-	
Week 3	-	-	-	-	100	-	-	-	-	
Week 4	-	-	-	-	-	-	-	100	-	
Week 5	-	-	-	-	-	-	-	100	-	

Table 1. An optimum allocation of resources for maximum profit

	Illness Insurance System	Illness Insurance with OR System	Unemployment Insurance System
Week 1	40	35	20
Week 2	30	42	31
Week 3	40	35	20
Week 4	30	42	31
Week 5	40	35	20

Table 2. Mean values of system demands per week

	Insured Object Coverage	Insured Object Payment	Illness Coverage	Unemployment Coverage	Insured Person	Own Risk	Illness Insurance System	Unemployment Insurance System	Illness Insurance /w OR System	Max. Expected Profit
Week 1	50	50	-	-	-	-	-	-	-	€25000,-
Week 2	-	-	-	100	-	-	-	-	-	
Week 3	-	-	-	-	100	-	-	-	-	
Week 4	-	-	-	-	-	-	-	150	-	
Week 5	-	-	-	-	-	-	-	100	-	

Table 3. An optimum allocation of resources in changing system demands and resources

12. List of Illustrations

Figure 1. Feature diagram of insurance products

Figure 2. Dependencies among insurance software systems and framework parts

Figure 3. Market volume estimations for three kinds of insurance software systems

Figure 4. Periodic changes in human resources

Figure 5. The tool architecture

Figure 6. Modelling the insurance framework components

Figure 7. Running a simulation