Planning Models for Component–based Software Offerings under Uncertain Operational Profiles

Mary Helander
Department of Mathematical Sciences
IBM Watson Research Center
Cambridge, MA 02142
helandm@us.ibm.com

Bonnie Ray
Department of Mathematical Sciences
IBM Watson Research Center
Yorktown Heights, NY 10598
bonnier@us.ibm.com

Abstract

We present a modeling framework for allocating development or support effort among software components to meet a specified cost or reliability objective. The approach is based on knowledge of the linkage between the operational profile and the architecture of the system, but allows for uncertainty through the use of probability distributions to characterize usage. Stochastic optimization methods are presented to obtain efficient solutions to the allocation problem.

1 Introduction

Guidelines for achieving a specified reliability target under resource constraints play a key role in the software development planning process. In particular, methods that determine how to allocate resources among components of a software system to facilitate cost–efficient progress toward a quantified system reliability goal are essential. Here, a component is defined as a set of operations, a subsystem, a module, an object, or any other distinguishable software entity that can be assigned a failure intensity representing its reliability. If the intended usage of the system is given, then system reliability can be computed as a function of component failure intensities, the expected component utilizations, and the specified usage. Specification of usage through the assignment of occurrence probabilities to operations form the quantification known as the *operational profile*. (See Musa, 1993). However for many systems, in particular commercial software systems, the usage of the system in a production setting may vary considerably from an expected usage, or the usage of the system during the test phases of system development. This paper describes a modeling framework for cost and reliability planning in software system development that allows for random variation in the operational profile.

Previous authors have addressed the problem of software reliability allocation in the context of a specified operational profile. For instance, Poore, Mills and Mutchler (1993) used a spreadsheet approach to consider various strategies for allocating reliability to software modules. The paper of Helander, Zhang, and Ohlsson (1998) provides an analytical solution to the optimal allocation problem based on standard nonlinear optimization methods. A component utilization matrix is used to link system structure to *operations*, which are partitionings of the software system from a user's perspective. We build directly on the framework of Helander et al. (1998), by allowing for uncertainty in the operational profile through the use of probability distributions on the operation occurrences. This is especially important as the high-level componentization of software systems becomes more common. By high-level, we mean that what were once considered as individual products are now bundled together in different ways and sold as software systems. An example is the componentization of IBM software offerings, which bundles key elements of products such as DB2, Tivoli, and Websphere Application Server together so that they are quickly deployable in a customized setting. The usage of the components may vary widely from setting to setting, but reliability targets for the individual components need to be allocated so that overall reliability targets are achieved across a range of settings.

In Section 2, we provide the general model formulation for specification of component failure intensities to plan for a software system reliability target while minimizing costs. Section 3 gives the derivation of the stochastic optimization problem under a particular distributional assumption about usage. Section 4 presents an example.

2 Mathematical Formulation as an Optimal Planning Problem

Let n denote the number of software components. In assembling components to form a software offering, we want the components to be reliable enough so that the probability of failure-free execution for a configured system has probability of at least ρ $(0 < \rho < 1)$ of achieving failure–free execution with respect to an execution time interval of length τ. The problem of interest is to determine what the failure intensities of the individual components should be to achieve this target in the most cost effective manner. This assumes a cost associated with achieving a failure intensity value, a common concept in software engineering economics.

Denote $f(\lambda_1, \lambda_2, \ldots, \lambda_n)$ as the total cost of achieving the failure intensities $\lambda_1, \lambda_2, \ldots, \lambda_n$, which we assume to be a pseudoconvex, non–increasing function of the λ_j's. The function $\mathcal{R}(\lambda_1, \lambda_2, \ldots, \lambda_n; \tau)$ measures system reliability in terms of $\lambda_1, \lambda_2, \ldots, \lambda_n$, which will be used as the main decision variables in the reliability allocation cost–optimization problem. The following model simultaneously finds $\lambda_1, \lambda_2, \ldots, \lambda_n$:

$$\text{Minimize} \quad f(\lambda_1, \lambda_2, \ldots, \lambda_n) \tag{1}$$

$$\text{Subject to:} \quad \mathcal{R}(\lambda_1, \lambda_2, \ldots, \lambda_n; \tau) \;=\; \exp\left\{ -\sum_{j=1}^{n} \sum_{i=1}^{m} p_i \mu_{ij} \lambda_j \tau \right\} \;\geq\; \rho \tag{2}$$

$$\lambda_j \;\geq\; 0 \quad \text{for } j = 1, \ldots, n \tag{3}$$

where $0 < p_i < 1$, $\sum_{i=1}^{m} p_i = 1$, are the operational profile parameters and μ_{ij} $(0 \leq \mu_{ij} \leq 1)$ are the component usage parameters for $i = 1, \ldots, m$ operations and $j = 1, \ldots, n$ components with $\sum_{j=1}^{n} u_{ij} = 1$ for each i. Note that the reliability function in (2) is derived under the assumption that failure events are statistically independent and the system architecture is such that a failure in one component causes failure of the entire system, *i.e.* the system is not fault-tolerant with respect to the components as identified in the planning model. Other forms for $\mathcal{R}(\cdot)$ could be derived under different assumptions concerning the system architecture.

This model specified by (1), (2), and (3) is the multivariate, nonlinear, constrained optimization model introduced in Helander et al. (1998). It can be equivalently restated as:

$$\text{Minimize} \quad f(\lambda_1, \lambda_2, \ldots, \lambda_n) \tag{4}$$

$$\text{Subject to:} \quad g_0(\lambda_1, \lambda_2, \ldots, \lambda_n) \;=\; \rho \ln(\rho) \;+\; \rho\tau \sum_{i=1}^{m} p_i \sum_{j=1}^{n} \mu_{ij} \lambda_j \;\leq\; 0 \tag{5}$$

$$g_j(\lambda_1, \lambda_2, \ldots, \lambda_n) \;=\; -\lambda_j \;\leq\; 0 \quad \text{for } j = 1, \ldots, n \tag{6}$$

which follows a standard nonlinear programming model form involving minimization of a nonlinear objective function constrained by "≤ 0" inequalities. Conditions and solutions for this form are given in Helander et al. (1998) for some common cost functions in software economics (see, e.g., Boehm, 1981).

In this framework, uncertainty in the operational profile can be introduced by treating the operational profile parameters as random variables, giving a *stochastic programming* formulation. This formulation and general solutions are discussed in the next section.

3 Stochastic Optimal Reliability Allocation

Let $\tilde{\xi} = \{\xi_i, \ldots, \xi_m\}$ denote the random vector representing random variables replacements for the deterministic operational profile parameters in constraint (5), p_i, $i = 1, \ldots, m$. The decision model stated by (4), (5) and (6) in the context of a random operation profile leads to a problem statement

$$\text{“Minimize”} \quad f(\lambda_1, \lambda_2, \ldots, \lambda_n; \tilde{\xi}) \tag{7}$$

$$\text{Subject to:} \quad G_0(\lambda_1, \lambda_2, \ldots, \lambda_n; \tilde{\xi}) \;=\; \rho \ln(\rho) \;+\; \rho\tau \sum_{i=1}^{m} \xi_i \sum_{j=1}^{n} \mu_{ij}\lambda_j \;\leq\; 0 \tag{8}$$

$$g_j(\lambda_1, \lambda_2, \ldots, \lambda_n) \;=\; -\lambda_j \;\leq\; 0 \quad \text{for } j = 1, \ldots, n \tag{9}$$

In this formulation, we have replaced g_0 by G_0 to indicate that g_0 is now a random variable. As noted in Kall and Wallace (1994), this problem as a whole, and the constraint (8), are not well defined when trying to make a decision for setting the $\lambda_1, \lambda_2, \ldots, \lambda_n$ values, prior to knowing a realization of $\tilde{\xi}$. To address this, we need a deterministic equivalent of the model specified by (7), (8) and (9). One alternative is to solve a problem such as

$$P[G_0(\lambda_1, \lambda_2, \ldots, \lambda_n; \tilde{\xi}) \leq 0] \geq p, \tag{10}$$

where $0 < p < 1$ represents the chance that the overall reliability target is met.

Note that when a realization of $\tilde{\xi}$ is specified, e.g., as $p_1, \ldots, p_m$, then the values of p_i should sum to one, and each value must be between zero and one. A convenient probability distribution that achieves these properties is a Dirichlet distribution. A random variable $\tilde{\xi}$ is said to follow a Dirichlet distribution if its probability distribution function has the form

$$p(\tilde{\xi}) = Dirichlet(\tilde{\xi};\ \mathbf{u}) = \frac{1}{Z(\mathbf{u})} \prod_{i=1}^{m} \xi_i^{u_i - 1}, \tag{11}$$

when $\xi_1, \ldots, \xi_m \geq 0, \sum_{i=1}^{m} \xi_i = 1$ and $u_1, \ldots, u_m > 0$. The normalization constant is

$$Z(\mathbf{u}) = \frac{\prod_{i=1}^{m} \Gamma(u_i)}{\Gamma(\sum_{i=1}^{m} u_i)} \quad . \tag{12}$$

A univariate Dirichlet distribution reduces to a standard Beta distribution. The values $\mathbf{u} = (u_1, \ldots, u_m)$ determine the shape of the distribution. See Kotz, Balakrishnan, and Johnson (2000) for additional details on Dirichlet distributions.

Standard results in stochastic optimization, for example Kall and Wallace (1994), show that a solution involving (10) is obtained by applying the same non–linear programming techniques used to solve the deterministic formulations. For example, when (10) is quasiconvex and differentiable with respect to $\lambda_1, \lambda_2, \ldots, \lambda_n$, then the Kuhn-Tucker conditions may be used to characterize and identify an optimal solution. Application of such techniques, including validation of properties such as quasiconvexity for (10), requires derivation of the distribution of G_0. The next subsection explores this derivation.

3.1 Derivation of the Distribution for G_0

Upon examination of expression (10), we see that G_0 is basically a shifted and scaled linear combination of Dirichlet random variables. Recent results by Provost and Cheong (2000) provide an expression for the cumulative distribution function of a linear combination of Dirichlet random variables in integral form. Letting $a = \rho \log(\rho)$, $b = \rho\tau$ and $c_i = \sum_{j=1}^{n} \mu_{ij}\lambda_j$ in (8), the distribution function is given by

$$F_{G_0}(z) = \frac{1}{2} - \frac{1}{\pi} \int_0^{\infty} \frac{sin[\frac{1}{2} \sum_{i=1}^{m} 2u_i \tan^{-1}\{(c_i - \frac{(z-a)}{b})u\}]}{u \prod_{i=1}^{m} \{1 + (c_i - \frac{(z-a)}{b})^2 u^2\}^{(2u_i)/4}}\, du, \tag{13}$$

for $a + b\min(c_i) < z < a + b\max(c_i)$.

As stated above, the Kuhn-Tucker sufficient conditions can be used to find the solution to the stochastic optimization problem provided (10) is quasiconvex and differentiable with respect to the λ_i. For other cases, numerical approaches will be explored. Properties based on the assumption of the distribution function (13) will be reported in the full version of the paper, along with numerical results for a number of examples.

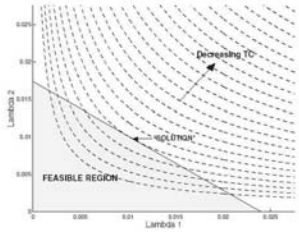

Figure 1: Solution for a deterministic version of the model specified by (1), (2), and (3)

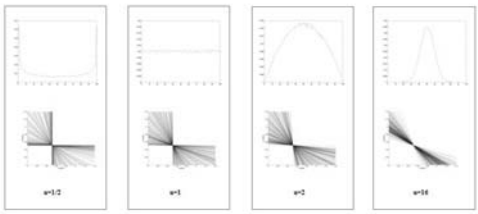

Figure 2: Solution under different parameterizations of ξ_i

4 Example

In this section, we graphically illustrate the effect of uncertainty on the allocation problem by extending the $n = 2$ component example from Helander et al. (1998). Figure 1 shows the solution for a deterministic version of the model specified by (1), (2), and (3). If we consider uncertainty in the usage profile by assuming the p_i's follow a Dirichlet distribution with a common shape parameter u, the effect is observed by the shifting of the main reliability constraint. This is illustrated graphically in Figure 2 for $u = 0.5, 1, 2, 16$.

Acknowledgements

The authors wish to thank Samer Takriti for extensive discussions concerning stochastic optimization methods and tools.

References

Boehm, B. W.(1981). *Software Engineering Economics.* Prentice–Hall: New Jersey.

Helander, M. E., Zhao, M., and Ohlsson, N. (1998) Planning Models for Software Reliability and Cost. *IEEE Transactions on Software Engineering*, vol. 24, no. 6, 420–434.

Kall, P. and Wallace, S. W. (1994). *Stochastic Programming.* John Wiley and Sons: New York.

Kotz, S., Balakrishnan, N., and Johnson, N. (2000). Continuous Multivarite Distributions, Volume 1: Models and Applications, John Wiley and Sons: New York.

Musa, J. D. (1993). Operational Profiles in Software Reliability Engineering, *IEEE Software*, vol. 10, no. 2, 14–32.

Provost, S.B. and Cheong, Y.-H. (2000). On the distribution of linear combinations of the components of a Dirichlet random vector, *The Canadian Journal of Statistics*, vol. 28, 417–425.