Reliability Engineering and System Safety 96 (2011) 1053–1062

Contents lists available at ScienceDirect

Reliability Engineering and System Safety

journal homepage: www.elsevier.com/locate/ress

Quantification of margins and uncertainties: A probabilistic framework

Timothy C. Wallstrom

Theoretical Division, Los Alamos National Laboratory, Los Alamos, NM 87545, United States

ARTICLE INFO

Available online 6 April 2011

Keywords:
QMU
Graphical model
Bayesian network
System reliability

ABSTRACT

Quantification of margins and uncertainties (QMU) was originally introduced as a framework for assessing confidence in nuclear weapons, and has since been extended to more general complex systems. We show that when uncertainties are strictly bounded, QMU is equivalent to a graphical model, provided confidence is identified with reliability one. In the more realistic case that uncertainties have long tails, we find that QMU confidence is not always a good proxy for reliability, as computed from the graphical model. We explore the possibility of defining QMU in terms of the graphical model, rather than through the original procedures. The new formalism, which we call probabilistic QMU, or pQMU, is fully probabilistic and mathematically consistent, and shows how QMU may be interpreted within the framework of system reliability theory.

1. Introduction

Quantification of margins and uncertainties (QMU) is a methodology for assessing confidence in the performance of nuclear weapons. It arose out of joint workshops on nuclear weapon certification, held between Los Alamos National Laboratory (LANL) and Lawrence Livermore National Laboratory (LLNL) in June and December of 2001. It was described in a short paper, first drafts of which were circulated in early 2002, by the responsible Associate Directors at the two laboratories, Goodwin and Juzaitis [1]; we will refer to this paper as GJ. GJ is available online, and may be regarded as the founding document of QMU, although it should be noted that it was not intended as a definitive statement.

QMU was formally reviewed by the JASON Defense Advisory Panel in 2005 [2], and by the National Research Council (NRC) of the National Academy of Science (NAS) in 2008 [3]. Both of these reviews have embraced the methodology as scientifically sound, and have encouraged its further use and development. The National Nuclear Security Administration (NNSA) has also embraced QMU and proposed requiring its use in annual assessments [4, p. 15].[1] QMU was proposed as part of the certification process for the Reliable Replacement Weapon (RRW), which was to be deployed without underground testing [3].

In this paper, we will consider the status of QMU as a general methodology for assessing confidence in complex engineered systems. QMU is based on a simplified model of the operation of the complex system. The operation is modularized into a set of consecutive stages, each of which is essential to successful operation. The success at each stage is determined by whether a key quantity, known as a metric, falls within a certain range, known as a performance gate. Confidence at each stage is assessed in terms of a ratio between the maximum permissible variation, or margin (M), for a metric, and its maximum anticipated variation, or uncertainty (U). (Precise definitions of M and U are given below.) If this "confidence ratio" is greater than one at all stages, one has confidence in the system; otherwise one does not. QMU confidence is a binary, yes/no quantity; it is not a probability.

The purpose of this paper is to provide both a critique of QMU, in its original formulation, and a way forward for its future development. We begin by showing that when uncertainties are bounded, QMU confidence may be consistently interpreted as perfect reliability. Realistic uncertainties, however, are not strictly bounded, and frequently have long tails. We show that when QMU is nevertheless applied in such situations, QMU confidence does not always correspond very well to reliability. The main difficulty is that QMU confidence can be present when the reliability is low. The reason is that QMU effectively sets the tails of the uncertainty distribution to zero, and is insensitive to the possibility that these tail probabilities may accumulate, thereby degrading system reliability.

To address these difficulties, we develop a fully probabilistic interpretation of QMU, which we call pQMU, based on the principles of Bayesian inference. This interpretation is a generalization of QMU, as originally conceived, in that it reproduces the results of the original QMU when the uncertainties are bounded. When uncertainties are not bounded, however, it provides a different method of computing reliability, which is

E-mail address: tcw@lanl.gov

[1] The three nuclear laboratories are Lawrence Livermore National Laboratory (LLNL), Los Alamos National Laboratory (LANL), and Sandia National Laboratory (SNL). They are a part of NNSA, which in turn is part of the US Department of Energy (DOE).

0951-8320/$ - see front matter © 2011 Elsevier Ltd. All rights reserved.
doi:10.1016/j.ress.2011.01.001

based on explicit probabilistic calculations, and not on the confidence ratio.

The new interpretation, in fact, is simply the representation of the complex system as a graphical model (GM) [5]. GMs provide a flexible format for representing the simplified system model which is at the heart of QMU. The key nodes represent system components, or stages of system operation, and the connections between these nodes are modeled by conditional probability distributions. The performance gate is modeled as a binary variable, whose value depends on the comparison of two random inputs.

GMs, which are also commonly called Bayesian networks (BNs), were developed in the field of artificial intelligence [6], were introduced into reliability theory by Barlow [7,8], Almond [9], and others, and have been found to be increasingly useful, particularly as more powerful computational techniques have been developed for their solution. They provide a natural generalization of fault trees and reliability block diagrams (RBDs) [10]. In a GM, the random variables need not be binary, or even discrete, and the conditional probability distributions may be arbitrary, rather than being simple logic gates.

QMU was developed from scratch, on the basis of deep intuitions about a particular complex engineered system, and was not initially derived from or interpreted in terms of existing methodologies in probability or statistics. In GJ, for example, the word reliability does not occur, and the word probability occurs only once. In interpreting QMU as a GM, we bring QMU fully within the fold of system reliability theory. Our interpretation also relates QMU to some of the earlier methodologies used for assessing reliability in nuclear weapons, which were based on fault trees and RBDs [11,12]. Inasmuch as GMs generalize fault trees and RBDs, QMU can be seen as a natural upgrade to earlier methods, rather than as a completely new framework.

There have been a number of earlier attempts to clarify the meaning of the QMU formalism [4,13–15]. The JASON and NRC reviews, cited earlier, also contain extensive discussions of the meaning and interpretation of QMU. We have chosen the GJ framework as our starting point, rather than these more recent efforts, because these later interpretations differ from each other, both in emphasis and in technical detail. The GJ framework is explained and developed in the 2003 paper by Sharp and Wood-Schultz [16], and also in [17]. Ref. [17] also contains an early effort to provide a probabilistic framework for QMU.

To avoid confusion, we emphasize that unless otherwise stated or implied from the context, whenever we use the term QMU in this paper, we will always mean the GJ framework, and more specifically, the GJ framework as we have interpreted it in Section 2. Our comments about "QMU" may not apply to other interpretations of QMU, some of which may have already addressed, in one way or another, some of the concerns raised in this paper.

The NRC and JASON reviews both note considerable ambiguity about the definition of QMU. In 2005, JASON noted that "There is no general agreement on what QMU means to the various scientific processes commonly used in science and technology, or whether QMU is in some sense a new such process" [2, p. 25]. In 2008, the NRC notes continuing ambiguity: "Finding 3-4. The QMU framework has yet to be clearly defined by the national security laboratories collectively or individually" [3, p. 42]. I suggest that part of the difficulty in defining QMU traces to the difficulties in understanding the original formulation in terms of probability theory, difficulties which I attempt to clarify in Sections 3.2 and 4. These difficulties have surfaced, for example, in the recognition that M/U may not always be adequate for the evaluation of performance gates [4, pp. 55–57; 3, pp. 27, 28]. The problem with the M/U criterion is that it compresses all of the information in the probability distribution down to a single number, which may be too crude in some cases.

The way forward, in my view, is to identify the essential features of QMU, realize these features in a consistent formalism, and then interpret historical QMU as an *approximation* to that formalism. Using this consistent formalism, it is possible to analyze the nature of this approximation, and identify the circumstances under which it is adequate and appropriate.

We begin, in Section 2, by providing a concise description of QMU, as originally formulated by GJ. In Section 3, we show that when uncertainties are bounded, the original formulation is identical to a probabilistic formulation in terms of graphical models, provided that QMU confidence is identified with reliability one. Section 4 provides a critique of the use of QMU, as originally formulated, for unbounded uncertainties. We show that QMU confidence, for a fixed choice of U, corresponds to a wide range of reliabilities, ranging from 0.83 to 0.99997 in simple examples. In Section 5, we explore the consequences of taking the graphical model interpretation as primary. We call the resulting framework "probabilistic QMU" or pQMU. We conclude with a brief discussion.

The use of GMs in QMU was first proposed in 2008 in an earlier version of this paper [18]. Independently, Urbina realized that GMs provided a useful framework for QMU, and applied them in the analysis of a specific engineered device in his 2009 Ph.D. thesis from Vanderbilt [19]. The latter is highly recommended as an illustration of how to apply these techniques to a concrete problem.

2. QMU: background

What *is* QMU, exactly? As noted in the Introduction, the definition of QMU has been a matter of considerable debate. QMU was not initially defined as a statistical formalism, but as a specific framework for a particular engineered system. We sketch the QMU framework, following GJ. The framework contains two parts: a simplified model of the system and a means for inferring "confidence" through an evaluation of the model.

In QMU, the performance of the system is broken into a series of *critical stages*, each of which must be completed successfully in order for the system as a whole to operate successfully. Each critical stage is characterized by one or more *metrics*, which are real-valued functions of the physical state, generically denoted by the symbol m. In order to successfully pass through a stage, the associated metric must have a value in a certain range; this range is called a *performance gate*, with boundaries generically indicated by g; the interval is indicated by G. Fig. 1 illustrates a performance gate model with three critical stages.

The detailed anatomy of a performance gate is shown in Fig. 2. The device is engineered so that the metric will lie in a certain range when the device is operated under specified conditions; this range is called the *operating range* (OR). We use the symbol γ generically for parameters that characterize the operating conditions, and represent the specified conditions as the requirement that $\gamma \in \Gamma$, for some set Γ. These conditions may involve the internal configuration of the system or the external environment under which it is operated. We use m_0 to indicate the engineered value, which depends on γ: $m_0(\gamma)$. The lower boundary of the OR, $m_{0,\min}$, is the minimum of m_0 over the specified conditions; the upper boundary, $m_{0,\max}$, is the maximum.

Due to uncertainties, however, the device may not operate as intended. We write

$$m(\gamma) = m_0(\gamma) + \varepsilon, \tag{1}$$

where $m(\gamma)$ is the actual value of the metric under conditions γ, and ε is a random variable representing the difference between the actual and engineered values. In defining the key quantities of QMU, we focus on the lower gate boundary; similar quantities can