

Planning for Safety Standards Compliance: A Model-Based Tool-Supported Approach

Davide Falessi, University of Rome Tor Vergata and Simula Research Laboratory

Mehrdad Sabetzadeh, Simula Research Laboratory

Lionel Briand, University of Luxembourg

Emanuele Turella, University of Rome Tor Vergata

Thierry Coq, Det Norske Veritas, Paris

Rajwinder Kaur Panesar-Walawege, Simula Research Laboratory

// A flexible, model-based tool-supported approach assists suppliers and certifiers in developing a formal agreement about the evidence necessary to demonstrate compliance with safety standards. //



SAFETY-CRITICAL SYSTEMS THAT depend on software—such as those found in the avionics, automotive, maritime, and energy domains—usually undergo a stringent certification process to show compliance with one

or more safety standards. Although the standards provide some guidance for collecting relevant safety information for this process, the guidance is mostly textual, imprecise, and hard to specialize for context-specific needs.

An agreement about the evidence necessary to demonstrate compliance with the applicable standards is an important aspect of safety assessment practice.¹ Without such an agreement, discrepancies between the ways suppliers and certifiers interpret the standards can give rise to problems on both sides. On the supplier side, the development process might not record the information specifically necessary for certification. Recovering this information after the fact can lead to significant cost overruns and deployment delays. Indeed, given the time lapse between development and certification processes, the original developers might have moved on to a different project, department, or company. Consequently, the supplier might need to reproduce the necessary evidence from scratch, often at extremely high costs. A high-profile example of such problems occurred during the certification of the Airbus A400M computer system, when a misunderstanding in certification requirements led to substantial rework.²

On the certifier side, problems show up when supplier documentation lacks structure and direct mention of safety information. Indeed, from our experience, suppliers often provide large fragments of their existing documents with the hope that the certifier will find the required safety information in them. The certifier must then spend the time and effort to sift through the documents and, in many cases, still not find the right information.

We developed an approach and supporting tool to systematically negotiate a consistent agreement between suppliers and certifiers about the information

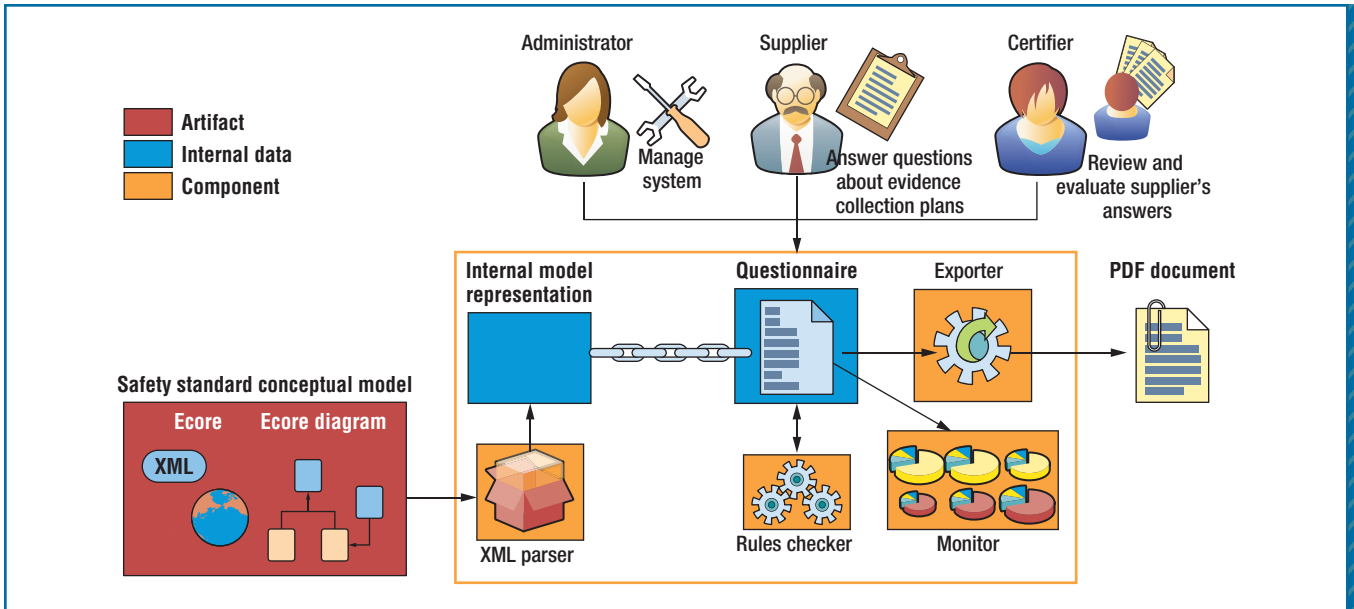


FIGURE 1. Overview of the solution for safety-evidence planning. An information model is encoded in an Eclipse-compatible format (Ecore) and provides the input to a questionnaire-based agreement process. Administrators, suppliers, and certifiers interact through the process to deliver a PDF document of the agreed-upon information requirements.

to collect prior to safety certification. Our approach and tool are standard-independent. However, for clarity in this article, we ground our discussions on IEC 61508, a widely adopted generic standard for managing the functional safety of software-dependent systems.³

A Questionnaire-Based Agreement Process

Our solution for safety-evidence planning involves a questionnaire-based agreement process, depicted in the center of Figure 1.

The process takes a safety standard's conceptual model as its input and uses a questionnaire to define details about what evidence to collect and the alternatives ways of recording and structuring it (see the "Related Work in Compliance Management" sidebar for more information). An administrator defines the questionnaire for a given safety standard. The supplier proposes answers (such as possible specializations), and a certifier accepts or rejects

the answers, providing the decision rationale via comments.

After the certifier agrees on the supplier's answers, the tool provides as output an agreement document (in a PDF) to review, print, and sign.

Model-Based Agreement

Our approach to planning safety-evidence collection is model-based. Specifically, to manage the apparent complexity of safety standards and provide an explicit and precise interpretation of the content, we use an information model that captures a given safety standard's core concepts and their relations.

In earlier work,⁴ we developed an information model for the IEC 61508 standard as a UML class model encoded in an Eclipse-compatible format (Ecore). Figure 2 shows a fragment of the Ecore information model. Briefly, an agreement concerning this fragment must specify which safety validation techniques are carried out in which phases and by which agents in relation

to the targeted safety integrity levels. We use this model fragment later to illustrate how to build a questionnaire around the concepts and relations in an information model.

The EvidenceAgreement Tool

We developed the EvidenceAgreement tool to support our approach. The tool is Web-based and lets certifiers and suppliers collaborate easily even when they're located at different geographical sites. EvidenceAgreement, its documentation, and its commented demo are publicly available at <http://modelme.simula.no/evagr/index.html>.

The Questionnaire

An administrator is in charge of creating a questionnaire that captures the information required to comply with a safety standard. In practice, the administrator role is typically played by one or several experts—usually, certifiers—who can interpret the relevant standard's details and enumerate alternative

ways of achieving compliance in different contexts.

The administrator assigns a specific questionnaire for suppliers and certifiers to use in reaching an agreement. The supplier and certifier must subsequently authenticate themselves and then choose the questionnaire to work on among the assigned ones. The EvidenceAgreement tool lets both the supplier and the certifier monitor the questionnaire's status. As Figure 3 shows, the status for each question appears in both text and color-coding. Pie charts show the status for different question types in the information model and for their aggregated "Final status."

In general, each standard has one information model, but the model can include several questionnaires. As a good

practice, we recommend using a single questionnaire per information model, encompassing all the domains to which the underlying standard applies. Our experience with IEC 61508 supports this recommendation. However, we can't be sure that one questionnaire could support any given standard in all possible domains, so the EvidenceAgreement tool allows the association of multiple questionnaires to an information model.

An administrator can assign a pre-existing questionnaire to the supplier and certifier or create a new one. A new questionnaire should provide some basic information including the questionnaire name, author, and brief description as well as an information model based on the standard for which compliance is required. The EvidenceAgreement tool

can accept any information model that can be encoded in the Ecore format.

Finally, the questionnaire must define the questions, answers, and rules as we describe in the remainder of this article.

Question Types

The questionnaire includes five types of questions.

Context. Contextual questions help suppliers better plan for evidence collection in a given context. For example, "In which domain will the product be deployed?" is a common question, and its answer would affect the safety level required. Consider a fire monitoring and control system deployed in an offshore oil platform as opposed to an on-land refinery. Each deployment would have

RELATED WORK IN COMPLIANCE MANAGEMENT

Safety certification is one facet of the more general problem of compliance management,¹ which encompasses topics such as process, medical, and environment regulations. A wide array of techniques and commercial tools exist to enable the execution and monitoring of compliance-related activities. Service-level management is a related notion,² aimed at developing a formal agreement for rendering services and ensuring their delivery accordingly.

Our work on safety certification could serve as an input to the existing compliance and service-level management tools, such as IBS's CompliantPro2 (www.ibs-us.com/en/products/compliantpro) and MetricStream's compliance management software (www.metricstream.com), which focus on the concrete collection and validation of evidence. In this context, our work's main contribution is the use of information models for formalizing the interpretation of safety standards and guiding decisions about what evidence to collect.

The research literature includes references to more systematic safety evidence collection as an important problem. In particular, Robert Lewis highlights the need for having a structured web of safety information covering not only hazards and safety requirements but also, among others, the requirements of development processes, hardware elements, human agents, and verification and validation results.³ Compliance assessment

schemes such as CASS⁴ for IEC 61508 partially address this problem by establishing guidelines for recording conformity. However, these schemes exist at a high abstraction level and must be specialized for a given domain or system. Our approach addresses this gap by helping with the specialization of safety information according to the needs of a particular context.

Our work relates most closely with questionnaire-based elicitation techniques.⁵ What differentiates our work is the use of model-driven engineering concepts to facilitate the specification of questions and possible answers and thereby ensure coverage of the underlying safety standards and consistency between the provided answers.

References

1. M. Silverman, *Compliance Management for Public, Private, or Non-Profit Organizations*, McGraw-Hill, 2008.
2. R. Sturm and W. Morris, *Foundations of Service Level Management*, Sams Publishing, 2000.
3. R. Lewis, "Safety Case Development as an Information Model," *Proc. 17th Safety-Critical Systems Symp. Safety-Critical Systems: Problems, Process and Practice*, Springer, 2009, pp. 183–193.
4. "What Is CASS? Accredited Certification for Safety Systems to IEC 61508 and Related Standards," 61508 Assoc., 2005; www.61508.org/cass.htm.
5. W. Foddy, *Constructing Questions for Interviews and Questionnaires*, Cambridge Univ. Press, 1994.

different safety concerns and might need to comply with different safety levels.

Contextual questions are associated with the whole questionnaire, with no constraints on the number of questions or answers to each question. The supplier answers contextual questions at the beginning of the process, because the context has an overarching effect on all aspects of evidence planning. These questions are the only ones that don't require an agreement on the certifier side; the context is fixed by the supplier's obligations to its customers.

For the remaining types of questions, the certifier must review and agree (or disagree) with each of the supplier's answers.

Evidence concepts. Questions about evidence concepts concern the information model's classes, and the answers are textual descriptions of the types of

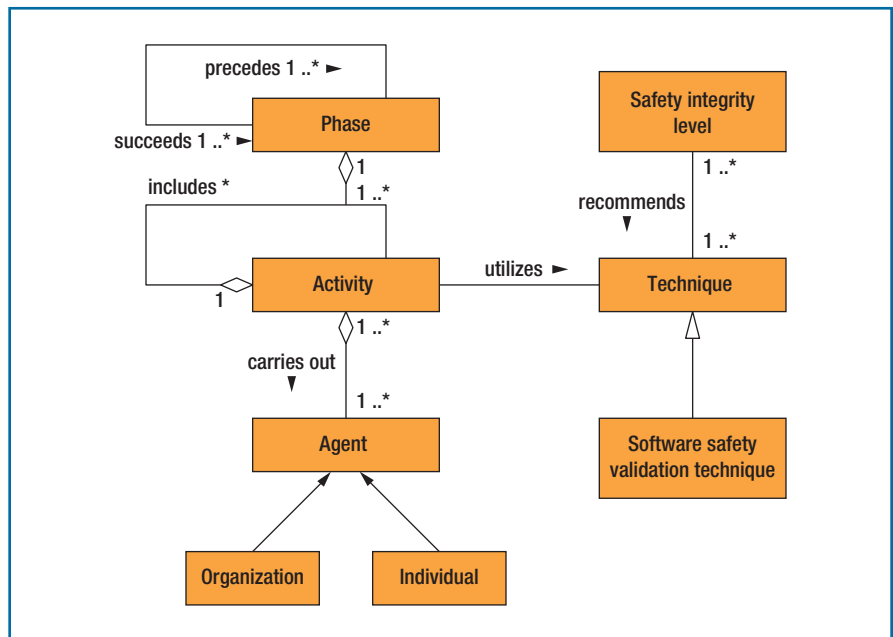


FIGURE 2. A fragment of the IEC 61508 information model. The main concepts of the safety standard and their relationships are formally described in a UML class diagram. The information model acts as the foundation upon which the questionnaire is created.

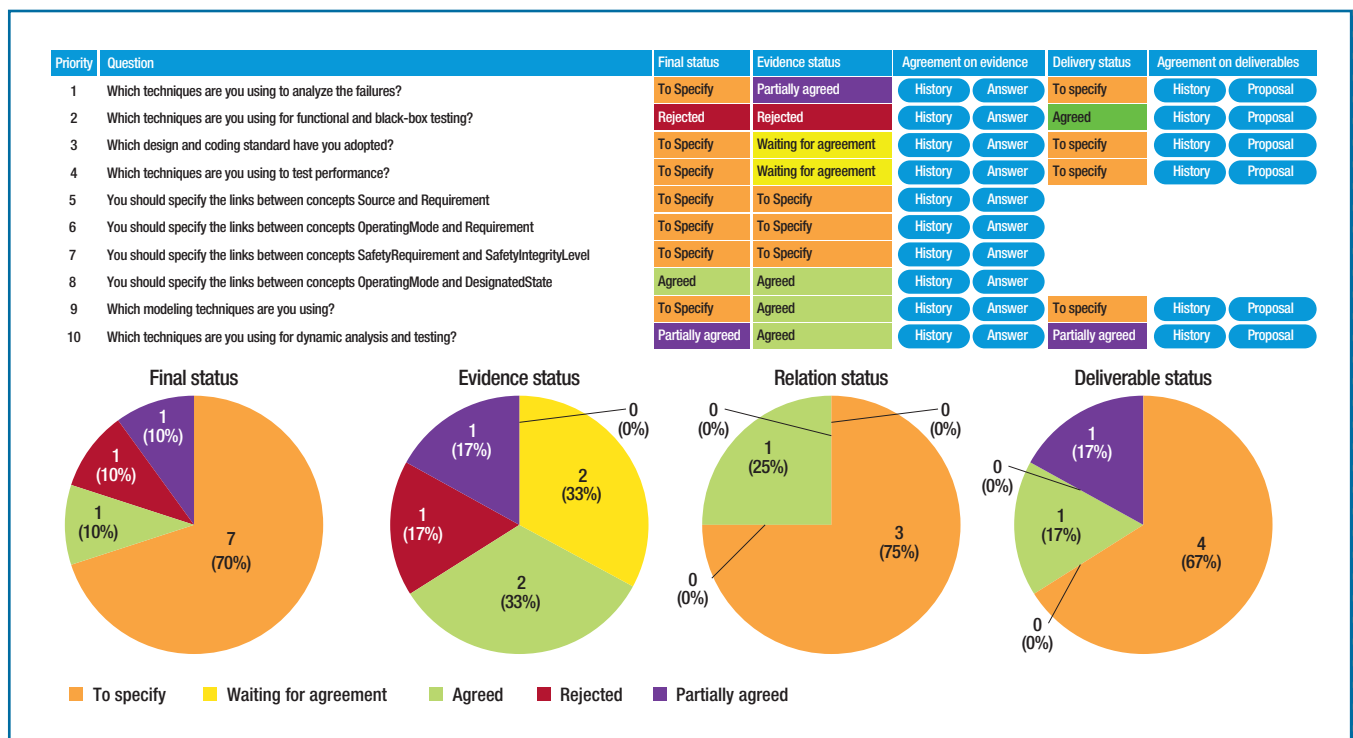


FIGURE 3. The EvidenceAgreement tool interface for monitoring agreement-completion status. On one side, the supplier can easily identify which questions need an answer—namely, those that haven't yet been answered or those to which the supplier rejects the answer. On the other side, the certifier can easily identify which of the supplier's answers still require a revision.

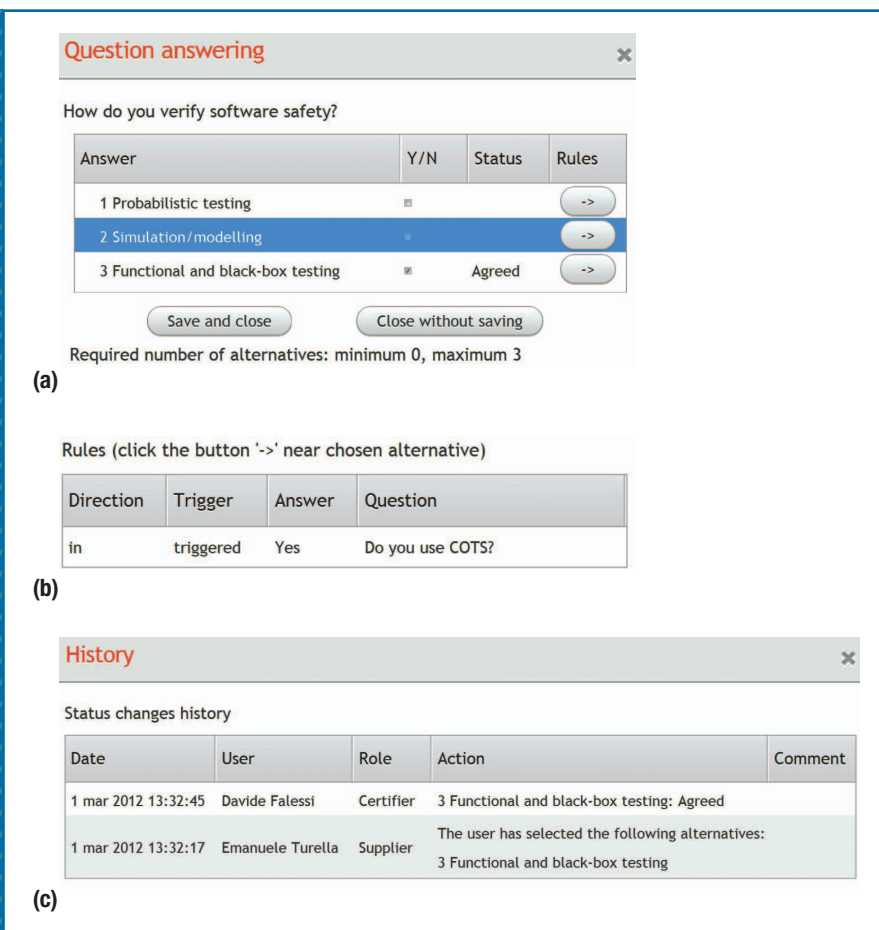


FIGURE 4. Example of exclusion rule type. For a given question (“How do you verify software safety?”), the tool shows (a) the possible answers, (b) the rules related to a selected answer, and (c) a log report of the agreement reached. The rules related to the answer selected in (a) are listed one per line, with the Direction column indicating whether the exclusion is incoming from, or outgoing to, the answer of another question. In this example, the highlighted answer “Simulation and modeling” in (a) has one incoming rule (from an answer to the question “Do you use COTS?”) that excludes it from being a possible answer of “How do you verify software safety?”

evidence required. The administrator creates at least one question of this type for all the information model classes. For example, “Which are the adopted techniques for software safety validation?” is a question that could be associated with the “software safety validation technique” class.

Answers to these questions describe the possible specializations of the evidence. For instance, the alternative answers to the safety validation technique question in Figure 2 include “probabilis-

tic testing,” “simulation and modeling,” and “functional and black-box testing.”

The supplier can answer these questions by selecting from predefined answers or proposing new ones. A given questionnaire’s information model stores the predefined answers. The certifier must agree on all answers to questions about evidence concepts and, if necessary, can suggest additional answers.

Relations between evidence concepts. After the supplier answers the questions

about the model classes, it must elaborate on the relations between the classes. The questions for a given relation are automatically derived from the answers provided for the related class pairs. Answers are of the open text type. For example, once a supplier answers the questions for “agent” and “software safety validation technique” types of evidence in Figure 2, it can specify which agent will be in charge of applying which safety validation technique.

Deliverables. The certifier and supplier must agree on how to deliver the evidence. For each proposed evidence concept, the supplier must therefore answer the question, “Which deliverable(s) will provide this type of evidence?” Deliverables include artifacts (such as a given type of documentation) and actions (such as a review meeting). In the EvidenceAgreement tool, we use the Det Norske Veritas (DNV) plan-approval documentation types to populate the list of possible deliverables.⁵

The supplier can choose from predefined answers or propose new ones. An agreement must have at least one agreed deliverable per evidence concept.

Rules and Inconsistent States

We use rules to enforce consistency, completeness, and traceability in the questionnaire answers. The administrator defines the rules, and the EvidenceAgreement tool checks them at runtime. There are two types of rule, which specify the constraints a questionnaire must meet:

- A *multiplicity rule* prescribes the minimum number of answers that the supplier must propose for a given question. For example, the standard might require that at least two different techniques (answers) are adopted for software safety validation.
- An *exclusion rule* excludes the co-existence of two specific questions to the same or different questions. For

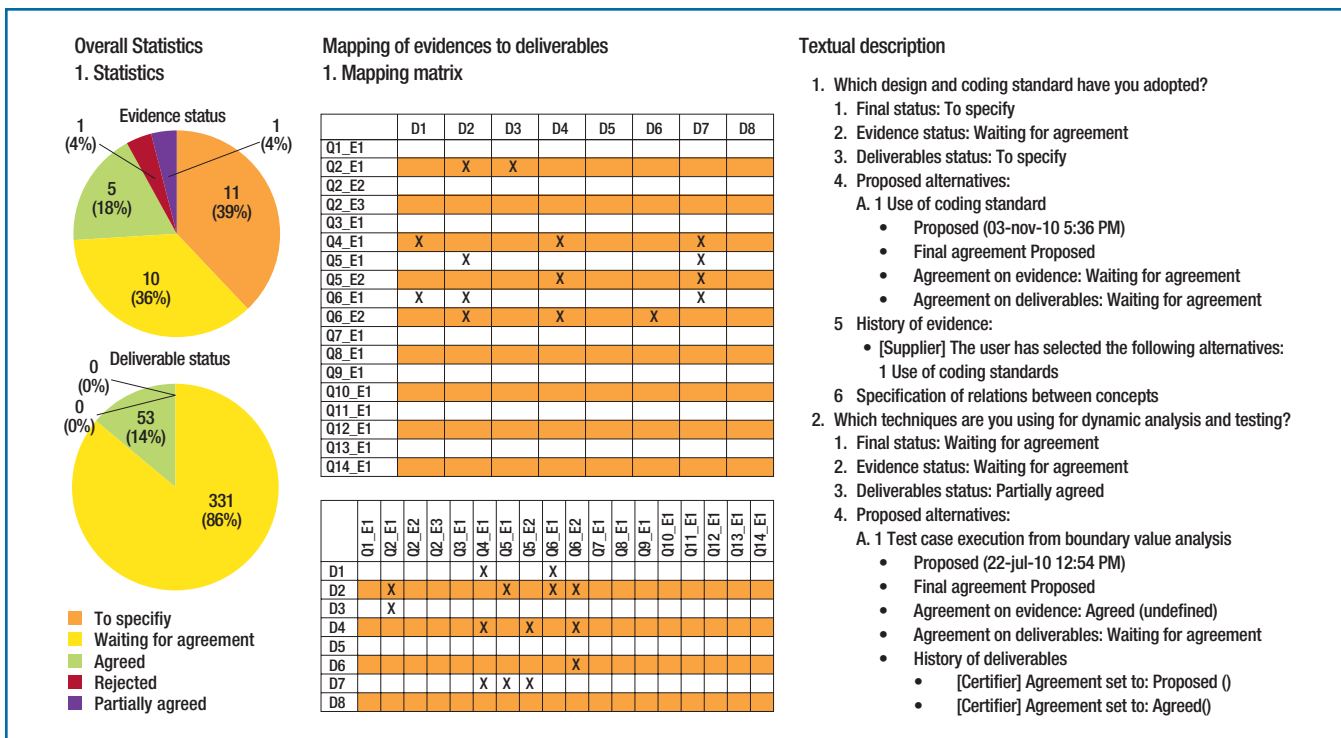


FIGURE 5. An excerpt of a PDF document produced by EvidenceAgreement. Both supplier and certifier can sign a printed document describing which evidence to provide; this is the result of the agreement procedure.

example, the rule excludes the answer “simulation and modeling” for software safety validation when the answer “COTS (commercial off-the-shelf) technology” has been selected.

The interactions between exclusion and multiplicity rules can lead to inconsistent states. For example, the supplier might not be able to meet the multiplicity constraint of one question because the answers available for it are excluded by answers to other questions. To illustrate, consider the example of IEC 61508, which specifies four safety integrity levels (SILs), with SIL 1 being the lowest level and SIL 4 the highest. At SIL 4, the supplier must often choose at least two testing techniques (answers) for software safety validation. If two of the three possible answers—for example, “probabilistic testing” and “simulation and modeling”—are excluded by answers to other questions, then an inconsistency occurs

because only one technique is available for software safety validation.

Inconsistent states require the user to backtrack and change the answers to one or more of the previously answered questions so that a feasible answer to the current question is no longer excluded. The EvidenceAgreement tool helps manage inconsistent states in two ways.

First, it prioritizes the questions at runtime according to the likelihood that each question will become infeasible to answer while respecting the multiplicity rules (see the first column in Figure 3). It estimates the likelihood as proportional to the number of required answers and the number of exclusion rules per answer. By following the priority order suggested by the tool, the user answers the most constrained questions first.

Second, for each alternative answer, the tool shows both all the alternatives that the answer excludes and all the alternatives that exclude the answer. The

example in Figure 4a lists a question and three possible answers (excluded answers are hidden). The “Y/N” checkmark indicates the selected answer—in this case, “Functional and black-box testing.” The blue highlight and the information in Figure 4b result from clicking the “Rules” button for that selection—in this case, the highlighted answer “Simulation and modeling” is excluded by an answer of the question “Do you use COTS?” In fact, the Direction column in Figure 4b is “in” (incoming) if the exclusion is incoming from an answer of another question and “out” (outgoing) if the selected answer excludes an answer of another question.

In Figure 4c, the tool reports the agreement according to the date, user, role, and the specific action performed regarding the question.

Output

The tool generates a PDF document as output. This customizable report is in-



the IEEE Computer Society. Contact him at d.falessi@ieee.org.

DAVIDE FALESSI is an adjunct research scientist in the Certus Center at the Simula Research Laboratory and an adjunct lecturer in the Department of Informatics, Systems and Production engineering at the University of Rome, Tor Vergata. His research interests focus on software and system engineering methodologies, from requirements to design, quality to reuse, and small academic controlled experiments to large-scale industrial case studies. Falessi has a PhD in computer engineering from the University of Rome, Tor Vergata. He's a member of



Contact him at mehrdad@simula.no.

MEHRDAD SABETZADEH is a research scientist in the Certus Center at Simula Research Laboratory. His research interests include requirements engineering, model-based development, and safety analysis and certification. Sabetzadeh received a PhD in computer science from the University of Toronto. He's a member of the IEEE Computer Society.



Society. Contact him at lionel.briand@uni.lu.

LIONEL BRIAND heads the software verification and validation laboratory at the University of Luxembourg's Interdisciplinary Centre for Security, Reliability, and Trust. Briand received a PhD in computer science, with high honors, from the University of Paris XI, France. He's co-editor in chief of *Empirical Software Engineering* and serves on the editorial boards of *Software and Systems Modeling* and *Software Testing, Verification, and Reliability*. He's an IEEE Fellow and recently received the prestigious Harlan Mills award from the IEEE Computer



Contact him at eturella@gmail.com.

EMANUELE TURELLA is a software architect at the National Institute of Nuclear Physics in Rome, Italy. His research interests include authentication and authorization infrastructure for Web-based applications. Turella has an MS summa cum laude from the University of Rome, Tor Vergata, with a thesis based on the EvidenceAgreement tool.



Contact him at thierry.coq@dnv.com.

THIERRY COQ is project director of Maritime & Energy at DNV IT Global Services. His research interests and industry experience center in the development and certification of software-dependant systems. Coq has an MS in engineering from Ecole centrale de Paris. Contact




Contact her at rpanesar@simula.no.

RAJWINDER KAUR PANESAR-WALAWEGE is a PhD scientist in the Certus Center at Simula Research Laboratory and in the University of Oslo's Department of Informatics. Her research interests include model-driven development, quality assurance of safety-critical systems, and empirical software engineering. Panesar-Walawege has an MSc in computer science from the University of Victoria, Canada.

tended mainly as an appendix to the certification contract. Figure 5 shows example output including statistical pie charts of the evidence and deliverable statuses, a matrix mapping the evidence to deliverables, and a textual description of the questionnaire results.

Our approach to constructing and specializing information models for safety standards lets certifiers and suppliers develop a negotiated and structured agreement about the evidence necessary for compliance. Although we haven't formally evaluated the tool, we note that it's not a major (and disruptive) break from current practice but, instead, a more effective way to do what's already being done manually with a plethora of different checklists and spreadsheets.

In the future, we plan to derive data schemas from the agreements generated by our approach and use them to construct and manage safety case databases that can be analyzed automatically. 

Acknowledgments

This work was conducted as part of the ModelME! Project, which is a joint research effort between Det Norske Veritas (DNV) and Simula Research Laboratory.

References

1. T.P. Kelly, *Arguing Safety - A Systematic Approach to Managing Safety Cases*, Univ. of York, 1998.
2. T. Spencer, "The A400M Military Transport Aircraft," *Defense Viewpoints*, 11 Dec. 2011; www.defenceviewpoints.co.uk/defence-industry/the-a400m-military-transport-aircraft.
3. Functional Safety of Electrical/Electronic/Programmable Electronic Safety-Related Systems (IEC 61508), Int'l Electrotechnical Commission, 2005.
4. R.K. Panesar-Walawege et al., "Characterizing the Chain of Evidence for Software Safety Cases: A Conceptual Model Based on the IEC 61508 Standard," *Proc. 2010 3rd Int'l Conf. Software Testing, Verification, and Validation*, IEEE CS, 2010, pp. 335-344.
5. *Recommended Practice DNV-RP-A201-Plan Approval Documentation Types*, DNV, 2010; <http://exchange.dnv.com/publishing/Codes/download.asp?url=2010-04/rp-a201.pdf>.