

Deriving Safety Cases for Hierarchical Systems in Model-based Development

Nurlida Basir¹, Ewen Denney², and Bernd Fischer¹

¹ ECS, University of Southampton, Southampton, SO17 1BJ, UK
(nb206r, b.fischer)@ecs.soton.ac.uk

² SGT/NASA NASA Ames Research Center
Mountain View, CA 94035, USA
Ewen.W.Denney@nasa.gov

Abstract. Model-based development and automated code generation are increasingly used for actual production code, in particular in mathematical and engineering domains. However, since code generators are typically not qualified, there is no guarantee that their output satisfies the system requirements, or is even safe. Here we present an approach to systematically derive safety cases that argue along the hierarchical structure of systems in model-based development. The safety cases are constructed mechanically using a formal analysis, based on automated theorem proving, of the automatically generated code. The analysis recovers the system structure and component hierarchy from the code, providing independent assurance of both code and model. It identifies how the system safety requirements are broken down into component requirements, and where they are ultimately established, thus establishing a hierarchy of requirements that is aligned with the hierarchical model structure. The derived safety cases reflect the results of the analysis, and provide a high-level argument that traces the requirements on the model via the inferred system structure to the code. We illustrate our approach on flight code generated from hierarchical Simulink models by Real-Time Workshop.

Keywords: Model-based software development, automated code generation, formal proofs, formal analysis, safety case, automated theorem proving.

1 Introduction

Model-based development and automated code generation are increasingly used for actual production code, in particular in mathematical and engineering domains. For example, NASA's Project Constellation uses Real-Time Workshop for its Guidance, Navigation, and Control (GN&C) systems and subsystems. However, since code generators are typically not qualified, there is no guarantee that their output is correct or even safe, and additional evidence of its safety is required. In previous work [5], we have thus constructed safety cases from information collected during the formal verification of the generated code. We also have constructed safety cases that correspond to the formal proofs found by automated theorem provers of the verification conditions, and reveal the underlying proof argumentation structure and top-level assumptions [6].

This paper is a continuation of our previous work, but here we systematically derive safety cases that argue along the hierarchical structure of systems in model-based

development. The safety cases are constructed mechanically using a formal analysis, based on automated theorem proving, of the automatically generated code. The analysis is driven by a set of formal safety requirements on the model which express as logical formulas the properties that the (software sub-) system's output signals must satisfy for the (overall) system to be safe. As an example, consider the GN&C subsystem as part of an overall spacecraft system where $vel_1 :: vel(\text{ECI})$ might express the requirement that the signal vel_1 must contain a velocity in the right (i.e., ECI) frame. These safety requirements result from a hazard analysis of the overall system, for example uncovering that the spacecraft can get into an unsafe state if the signal vel_1 is not representing a velocity measurement in the right frame; however, in this paper we assume the safety requirements as given. While Leveson *et al.* [14] rightly argue that most accidents involving software are due to errors in the software requirements specification, it is important to note that the safety requirements are not the same as the software requirements specification, precisely because of their origin in a hazard analysis.

We illustrate our work using the verification of two safety requirements for a spacecraft navigation system that was generated from a Simulink model by Real-Time Workshop [3]. Each requirement induces a slice of the system architecture. The analysis recovers the hierarchical structure of these slices from the code. It identifies how the system safety requirements are broken down into component safety requirements, and where they are ultimately established, thus deriving a hierarchy of requirements that is aligned with the hierarchy of the components in the system. It also identifies requirements that rely on additional externally given assumptions. The requirements determine the interface between the software system safety cases derived here (where they are the root elements) and the overall system safety case (where they are leaf nodes). The derived safety cases reflect the results of the program analysis, and provide a high-level argument that explains how the system slices establish the corresponding safety requirements. They help tracing the safety requirements from the model via the inferred system structure to the code, thus providing independent assurance of both model and code.

Here we describe the hierarchical structure of the safety cases and show how they are derived systematically from information uncovered by the analysis phase. They provide a traceable safety argument that shows in particular where the code, subsystem, and system depend on any internal and external assumptions. We believe they highlight the claims, key safety requirements, and evidence that are required to trust the generated code. We expect that this further alleviates distrust in code generators, which remains a problem for their use in safety-critical applications.

2 Background

2.1 Model-based Software Development

Model-based software development comprises a number of techniques that focus on creating and transforming domain-specific abstractions or *models* rather than algorithmic concepts or even code. In *model-based design* [3, 18], mathematical or, more often, visual methods are used to create an initial model of the system design. It is commonly used in the control systems domain, where block diagrams provide an accepted notation. Blocks can represent arbitrary computations and can be nested hierarchically,

which helps countering system complexity. They are connected by wires that represent the flow of signals through the system. A number of academic and commercial tools support model-based design in this domain. We focus on MathWorks Simulink [3], which is used by many NASA projects for at least some of their modeling and code development, particularly for GN&C problems. Simulink comes with a large library of standard modeling blocks that provide mathematical operations and signal routing suitable for control systems and complex operations.

Model-based code generation [18, 16] usually complements model-based design, and translates specifications in the form of a model into a program in a high-level programming language such as C or ADA. The translation process is often organized as a sequence of model transformations, where the last model is equivalent to the program. The final source code generation step can then be implemented using a simple template engine. Here we focus on a commercial generator, Mathworks Real-Time Workshop Embedded Coder 5.1 [3]. Real-Time Workshop generates ANSI/ISO compliant C and C++ code from MathWorks Simulink and Stateflow models. Embedded Coder adds various features which are useful for generating C code tuned for embedded devices.

2.2 Formal Program Analysis using AutoCert

The techniques described here are based on the AUTOCERT code analysis tool [9], which takes a set of requirements, and formally verifies that the code satisfies them. AUTOCERT can verify execution-safety requirements (e.g., array bounds), as well as individual mathematically specified requirements. AUTOCERT thus supports certification by formally verifying that auto-generated code is free of certain safety violations and complies with domain-specific safety requirements as those mentioned.

AUTOCERT follows the Hoare logic approach to verification, which needs *annotations*, i.e., logical assertions of program properties, at key locations in the code. These annotations are constructed automatically by a post-generation inference phase that exploits the idiomatic nature of auto-generated code and is driven by a generator- and domain-specific set of idioms. The inference algorithm builds an abstracted control-flow graph (CFG), collapsing the code idioms into single nodes. It then traverses the CFG from *use* nodes (where a requirement must hold) backwards to all corresponding *definitions* (where the relevant properties are established) and annotates the statements along the paths as required [9]. The definitions typically correspond to blocks (more precisely, to fragments of the code implementing a block), which can use assumptions on the properties of their input signal to establish the requirement. Hence, the inference algorithm must recurse over the variables corresponding to the input signals, derive the assumptions, and establish them as new requirements. As byproduct, the inference derives the subsystem and component interfaces (i.e., the requirements placed on them, and the assumptions made by them) as well as the system's overall assumptions, which need to be established by its context. This chain of requirements on variables and their definitions constitutes the backbone of our safety argument. A verification condition generator (VCG) processes the annotated code, feeding a set of verification conditions (VCs) into an automated theorem prover (ATP); their proofs guarantee that the code satisfies the requirements and also validate the definitions identified by the analysis, and thus the derived architecture. In the safety case, the proofs serve as evidence.

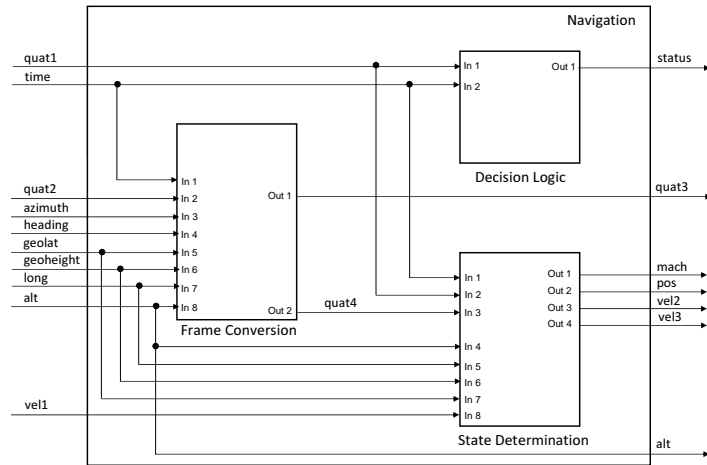


Fig. 1. High-level Architecture of Navigation System

During the course of analysis, AUTOCERT records various facts, such as the locations of uses and definitions, which are later used as input to the safety case generation process. Here, we extended the existing mechanism to record additional information from which we can reconstruct the system architecture slices.

2.3 Guidance, Navigation, and Control Systems

Spacecraft are typically decomposed into a number of different systems such as the power, thermal protection, or guidance, navigation, and control (GN&C) systems [22]. The GN&C system is a necessary element of every spacecraft. Here, we focus on the navigation subsystem within the GN&C system. It is used to determine a spacecraft's orientation and position, which is challenging from a safety perspective, due to its complex and mathematical nature. We give a simplified description of the subsystem where we also have changed the names of components and signals from the original.

The Navigation subsystem (cf. Fig. 1 for its architecture) takes several input signals, representing various physical quantities, and computes output signals representing other quantities, such as Mach number, angular velocity, position in a specified frame of reference, and so on. Signals are generally represented as floating point numbers or as quaternions and have an associated physical unit and/or frame which correctness are critical to the safety of the system. However, the units and frames are usually not made explicit in the model, and instead are expressed informally in comments and identifiers.

The Navigation subsystem is comprised of three sub-subsystems or components, a decision logic that computes a status value irrelevant to the requirements we consider here, a frame conversion, and a state determination. Frame Conversion first converts the frames of the incoming signals from a vehicle-based coordinate system to an earth-based coordinate system. The transformations of the coordinate systems are done by converting quaternions to direction cosine matrices (DCMs), applying some matrix algebra, and then converting them back to quaternions [20]. State Determination then performs the calculations to determine the vehicle state (e.g., position, attitude,

and attitude rate) from these signals. It is defined in terms of the relevant physical equations. Note that there are no individual blocks within the Navigation subsystem, but only within the components and thus all computation happens there.

In the following we simplify the description by denoting Navigation as a “system” rather than a “subsystem” and Frame Conversion and State Determination as “components” rather than “sub-subsystems”. However, both components are specified by independent models, so that we indeed work with a hierarchical model.

3 Deriving Safety Cases from the Formal Analysis of Hierarchical Systems

3.1 Formalization of the Requirements

We illustrate our work using the results of the safety certification for the code generated from the Simulink model of the navigation system. We concentrate on the following two safety requirements, (1) Signal $quat_3$ is a quaternion representing a transformation from the the Earth-Centered Inertial (ECI) frame to the body fixed frame, and (2) Signal vel_2 is in the velocity in the ECI frame. Since we are working with a formal, logic-based analysis framework, we need to formalize these requirements using a domain theory, as (1.) $quat_3 :: quat(ECI, Body)$, and (2.) $vel_2 :: vel(ECI)$ Here, ECI and Body are constants denoting the respective frames, $quat$ and vel are functions denoting transformations of or quantities in those frames, and $::$ is a predicate that asserts that the signal represents a transformation between (resp. quantity in) the required frame(s).

Obviously, the formalization of the safety requirements is safety-relevant: a wrong formalization can invalidate the assurance provided by the proofs [4, 15]. It thus needs to be called out and justified in the safety case.

3.2 Architecture Recovery

In order to certify the requirements on the navigation system, and to build a comprehensible safety case, we need to know where in the system they are established, and which parts of the system contribute to them. Intuitively, we can see in the system architecture (cf. Fig. 1) that the first requirement should be established by Frame Conversion, since the signal $quat_4$ comes straight out of that component (and similarly for vel_2 and State Determination in the case of the second requirement). However, this view is too simplistic. First, without looking into the component models it is not clear whether the requirement is indeed established within a component, or simply passed through (cf. for example alt in Navigation), and which of the component’s input signals (if any), or more precisely which assumptions on them, are used in establishing the requirement. However, simply expanding the component models destroys the hierarchical structure of the system. More importantly, the safety of the system ultimately depends on the safety of the code rather than the model, but because we cannot trust the code generator to translate the model correctly we cannot derive any trust from the model.

Instead, we analyze the code and recover the slice of the system architecture that is relevant to a given safety requirement. We record when the control flow based analysis




Fig. 2. Architecture Slices Recovered for Example Requirements

enters resp. leaves a component (implemented by RTW as a parameter-free procedure), and then remove the part of the requirements-definition chain that is contained within the component. The key to obtaining precise architecture slices is to identify situations in which the control flow path just passes through a component, without encountering a definition. In these cases, we can ignore the component altogether. We then assemble the slices from the signals involved in the recorded requirements-definitions chains and from the retained component. Note that the recovered architecture slices are not based on the call graph (even though the analysis follows the control flow) but on the implicit dataflow. They are thus very similar to the also dataflow-based Simulink models from which the code was originally generated. The main difference is that the slices use the generated variable names (here simplified to capitalized versions of the corresponding signal names) while the model of course uses the signal names.

Fig. 2 shows the architecture slices recovered for both requirements. In both cases, the irrelevant Decision Logic component has been removed by the analysis. For the first requirement, it has further identified that $Quat_3$ is unaffected by the call to the State Determination procedure, and consequently removed that component as well. For the second requirement, the analysis has identified $Quat_4$ as the (global) variable through which the two components communicate. In addition, although not shown in Fig. 2, it has derived the property placed as an assumption on this variable by State Determination, i.e., $Quat_4 :: \text{quat}(\text{NED}, \text{Body})$. This becomes a subordinate requirement to the original safety requirement, reflecting the hierarchical model structure. The requirements hierarchy is completed by the assumptions placed on the variables i.e., Vel_1 and i.e., $Quat_2$ corresponding to the components' input signals.

The property derived for $Quat_4$ also becomes part of the interfaces of both components that are connected through this link, as assumption on the State Determination and as safety requirement on Frame Conversion. By regrouping the analysis results by component rather than by original safety requirement, we thus obtain full component interfaces. They give a complete functional specification of the component, including all assumptions, as far as it is required to satisfy the given system-level safety requirements. The interfaces also serve as starting point for verifying the components independently, hence allowing a compositional (and therefore scalable) verification.

The recovered system architecture and requirements hierarchy already constitute a core safety argument: Navigation satisfies the safety requirement (2') if the components Frame Conversion and State Determination satisfy their respective interfaces, and the requirements for Vel_1 , $Quat_2$, and $Quat_4$ hold. This argument can serve as blueprint for a full-fledged safety case. In addition, the derived component interfaces serve as starting

points for the construction of independent safety cases for the components, yielding a hierarchy of safety cases that is aligned with the system’s hierarchy of models.

3.3 Arguing from System-Level Safety Requirements to Component-Level Safety Requirements

We now describe (in simplified form) the safety case derived from the information provided by the formal analysis. Its root goal is to show that the Navigation system satisfies the given safety requirements. We use the Goal Structuring Notation (GSN) [13] as technique to explicitly represent the logical flow of the safety argument. Its main elements are goals, strategies, and evidence. A strategy describes how a goal is achieved, reducing it to new, simpler subgoals, while evidence shows why a goal is met by the system. Contexts explain additional information to support the argument, assumptions outline what has to be relied upon in order to make an argument valid and justifications describe the rationale of the argument.

The upper part of the safety case argues the safety of the method of formal reasoning that we use but also points out the important provisos that we abstract away from real-time, and numerical issues. This is a straightforward modification of our previous work on programs without hierarchical system structure (cf. of Fig. 3 Tier I: Explaining the Safety Notion in [5]). Here, we thus focus on the lower part of the safety case that explains that, and how, the generated source code `Nav.cpp` satisfies the given safety requirements by providing formal proofs as evidence (see Fig. 3).

The key argument strategy here is to argue over each individual requirement that contributes to the program safety. The additional information that is required for the strategy to be understood and valid is identified and explained. This concerns the independent validity of the safety requirements and the logical consistency of the assumptions. We thus assume that no safety requirement is available for use as a (logical) assumption in the safety proofs, which prevents vacuous proofs based on mutually recursive dependencies between requirements and assumptions. We further assume that the given and derived assumptions together are consistent, again to prevent vacuous proofs. Each assumption is justified by a valid justification (e.g., the consistency can be checked by theorem prover).

As a result of this strategy we get as many subgoals as there are safety requirements given. Here we focus on the goal (R2) corresponding to the second requirement, i.e., that vel_2 is a velocity in the ECI frame. Context nodes with hyperlinks outline additional evidence in the form of documents, containing, for example, a detailed description of the system and requirement, and also the result of the hazard analysis.

The next step of the argument transitions from the informal level to a formalized safety requirement. This step helps in showing that the formal verification runs over the correct requirement, based on the right formula and variable, and thus provides a relevant proof of the program. We use an explicit strategy to describe this transition, which spawns three subgoals. As already discussed in Section 3.1, the first subgoal (F2) demonstrates that the formal proof is based on an appropriate formalization of the requirement, and the safety case points to the documentation of the logical domain theory as evidence of this. The second subgoal (L2) “glues together” model and code levels, which allows us to build a safety case for the model based on the analysis of the code.

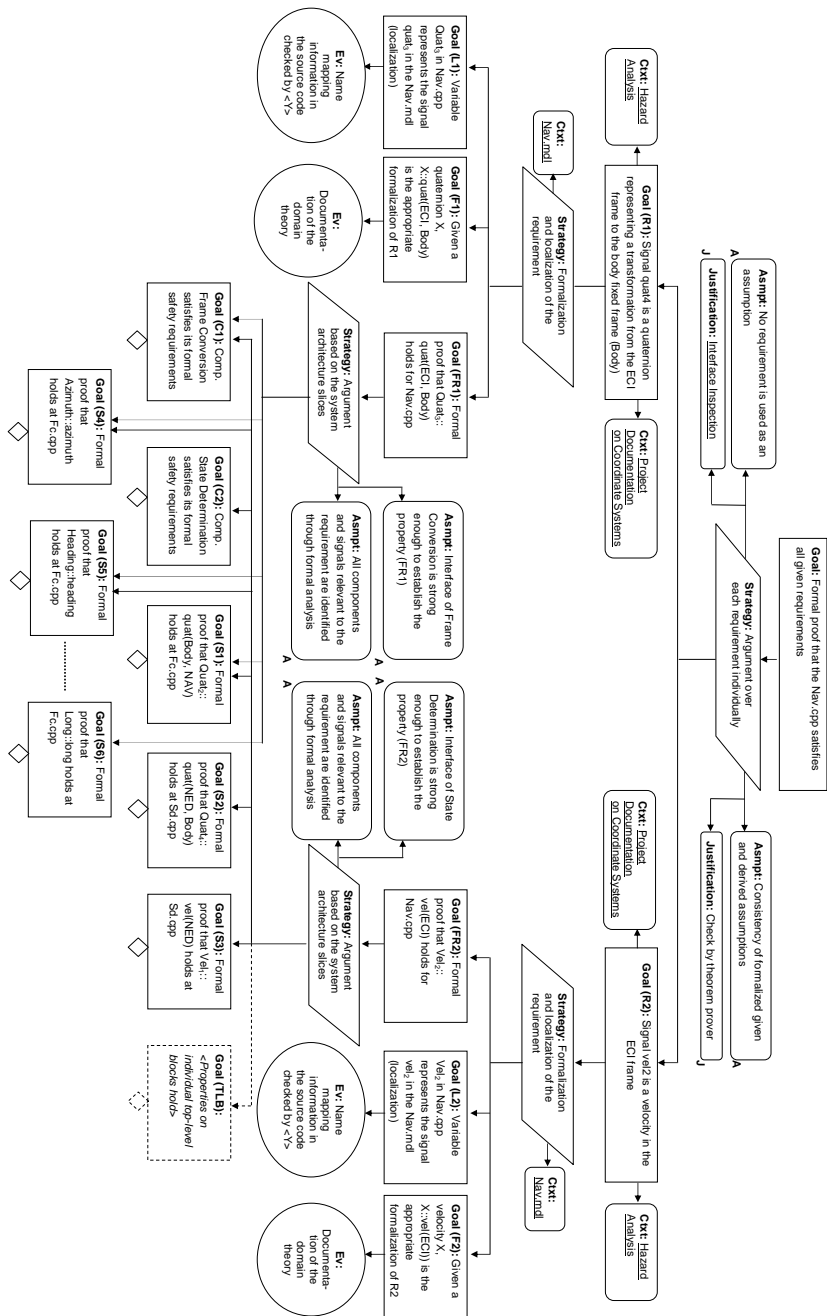


Fig. 3. Arguing from System-Level Requirements to Component-Level Requirements

In particular, as discussed in Section 3.2, we need to show the mapping between the signal names used in the model and the corresponding variable names used in the source code, which cannot be recovered by our analysis but must be given externally. Here, the safety case points to the mapping information given in the source code, and that it has been checked by a reviewer, as evidence. In addition, at this goal we also have to show the mapping between the model and code files, and in particular, in which code file the property formalized in (F2) has to be shown. In our example, this is straightforward, but for larger systems the localization needs more evidence.

With the results of (F2) and (L2) we can now construct the final subgoal (FR2) of our strategy, which shows that the fully formalized safety requirement $Vel_2:: vel$ (EC1) holds after execution of the code in `Nav.cpp`. This requirement eventually needs to be proven formally. However, at this level of abstraction, the safety case does not use an argument based on the full formal proofs. Instead, we use an argument based on the system architecture, or more precisely, on the recovered system architecture slices. It shows how the system level requirements are broken down into the component level requirements i.e., properties of the part of the system that is relevant to satisfy the requirement (FR2). The strategy is based on the assumption that the formal analysis has identified all relevant components and signals. We thus reduce (FR2) to a number of (delayed) subgoals for the components and signals in the architecture slice. For each component, we need to show that it satisfies the safety requirements specified in its interface (cf. subgoals (C1) and (C2)). This induces a further assumption on the strategy, namely that the interface is strong enough to show the requirement (FR2). Delaying the subgoals allows us to reuse the component-level safety cases. This way, we achieve a hierarchical structure for the system safety case that mirrors the hierarchy embedded in the system architecture. If the system contains top-level blocks in addition to the components (which is not the case in our example), we need to reason about their properties as well. This is indicated by the dashed subgoal (TLB). For each variable representing a signal, we need to show that it satisfies the safety requirements derived by the analysis (cf. subgoals (S1) to (S5)). This guarantees that the components' assumptions are met. These subgoals are delayed here as well, to keep the safety case compact. Their expanded structure again follows the lines of our previous work [5], and uses the argumentation shown in Fig. 5 (Tier III) of the safety case there with small modifications; in particular, the notion of safety condition needs to be replaced by that of safety requirement. Note that we make no distinction at this level between subgoals that are established by the components (S2) and those that are reduced to assumptions about the system's input signals and thus have trivial formal proofs, e.g., (S4).

3.4 Arguing from Component-Level Safety Requirements to Source Code

In the next step of our hierarchical development, we argue about the safety of the components wrt. their identified interfaces. The component-level safety cases also argue about a set of requirements, but there are two significant differences to the system-level safety cases. First, the component-level requirements are already formalized, due to the use of the formal analysis, so that we do not need to argue about the safety of the formalization and localization any more. Second, the argument will generally go down

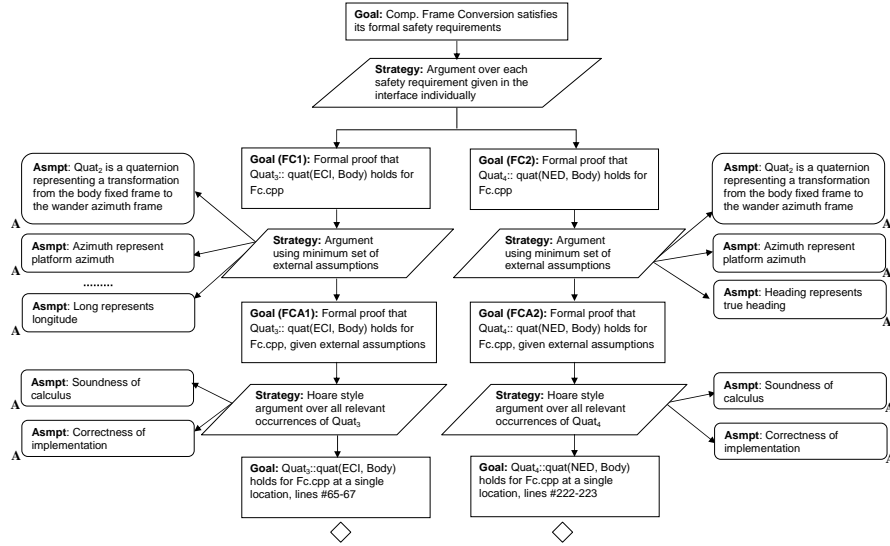


Fig. 4. Component-level safety case for Frame Conversion

to the the level of the generated code, with the proofs of the VCs as evidence; obviously, however, another layer of hierarchy is introduced if a component contains further (sub-system) components.

Fig 4 shows the safety case for Frame Conversion component. For each component, the strategy is to argue over each individual safety requirement stated in its interface. Here, we have two requirements, (FC1) which is used to discharge the (essentially identical) system-level goal (FR1) via (C1), and (FC2), which is used to discharge the signal subgoal (S2). Even though they serve different purposes in the system-level safety case we treat them the same at the component level. We focus on (FC2) here.

The component interfaces also list the assumptions that the component itself makes about the environment. However, not all assumptions are used for all requirements, so we use an explicit strategy to argue only using the minimal set of external (i.e., on the system's input signals) assumptions. Note that the use of internal assumptions (e.g., on $Quat_4$), which have been identified as subgoals in the system-level safety case (cf. (S2) in Fig. 3) will be made explicit further down in the component-level safety case.

The next strategy finally transitions from the safety argument to a program correctness proof, using a Hoare-style argument over all relevant occurrences of the variable. In this case, it leads to a single subgoal, proving that the safety requirement holds at the given source location. This is predicated on the assumptions that the applied Hoare-calculus is sound, and that the VCG is implemented correctly, which need to be justified elsewhere (cf. Tier I of the safety case presented in [5]). The structure of the Hoare-style argument is determined by the structure of the program. Since rest of the safety case is constructed as described in our previous work [5], we do not expand it here any further.

Showing the safety of the component is thus reduced to formally showing the validity of the VCs associated with each requirement in the interface. If (and only if) proofs for all corresponding VCs can be found, then the property holds for the entire program. The construction of safety cases from the proofs is described in our previous work [6].

3.5 Combining System-Level and Component-Level Safety Cases

Splitting the argument into system-level and component-level makes it easier to follow and allows to factor out common sub-arguments, but in order to obtain a complete argument we need to combine the system- and component-level safety cases. However, simply attaching the entire component-level safety cases to the corresponding component goals would introduce redundancies. Clearly, not every safety requirement on the system level relies on the full set of requirements established by the components, for example, (FR2) only uses the requirement derived for $Quat_4$ (i.e., goal (FC2) in Fig 4).

We thus replace each component goal only by the “branches” of the component-level safety case that are required; this information is provided by the program analysis. For component goals that are shared between different requirements this will lead to an “unsharing”. For example, (C1) will be replaced by the branch rooted in (FC1) below (FR1) and by the one rooted in (FC2) below (FR2). However, common subgoals at the level of the Hoare-style argument, which are based on computations contributing to different requirements, can remain shared.

Additional changes occur elsewhere in the system-level safety case. The assumptions to the architecture-based strategy solving (FR1) and (FR2) can be removed because the detailed argumentation in the component-level safety case provides the necessary evidence. Further the subgoals associated with the system’s input signals (i.e., (S1) and (S3)–(S6)) can be removed because corresponding subgoals still appear as leafs in the component-level safety case, where they are discharged by the assumptions. The subgoals on the connecting signals (here only (S2)) will be replaced by the root goals of the corresponding branches in the component-level safety case (i.e., (FC2)) at the appropriate position in the Hoare-style argument for the client component (i.e., State Determination).

4 Safety Case Construction

The safety cases described here quickly become too large for manual development. Fortunately, the bulk of the argument is based on information provided by AUTOCERT’s formal program analysis, and the argument structure follows the program and analysis structure, so that a largely automated safety case construction is possible. However, some information cannot be produced by the program analysis, such as environment constraints, external assumptions, list of related documents, or model names. This information must be specified externally by a safety engineer. This also applies to the formalization of the top-level safety requirements that drive AUTOCERT’s analysis and their integration with the system-wide hazard analysis and safety case.

In order to support the automated safety case construction, we integrate AUTOCERT’s formal program analysis with an existing commercial safety case tool, Adelaar’s ASCE v3.5 tool [1]. We extended AUTOCERT to extract the manually specified information from its own input and to structure this together with all information derived by the analysis into an XML format. The XML file records all the relevant information needed for the safety case construction. Subsequently, an XSLT program is used to transform this into a second XML format that logically represents the structure

of the safety case as defined by safety case templates underlying the examples shown above. Here, the templates were designed so that the same argument structure can easily be adapted to other programs and systems. Finally, we use a custom Java program to present the safety case using GSN. The Java program helps to set the position of the nodes in the safety case which involved some mathematical calculations and to represent the argument to follow the standard Adelard ASCE file format. This architecture avoids a tight integration of the analysis (i.e., AUTOCERT) and presentation (i.e., ASCE) tools, and provides enough flexibility to change the latter with little effort.

The integration is largely completed; in particular, we have already fully automated the construction of the component-level safety cases that argue down to the code structure, and make up the overwhelming fraction of the combined safety case. However, the print quality of these large safety cases is insufficient for presentation, so we choose to recreate them in Microsoft Word here. The implementation of the integration of system-level and component-level safety cases, as described in Section 3.5, requires further tool support not provided by ASCE and is left for future work.

5 Related Work

The development and acceptance of a safety argument or safety case is a key element of safety regulation in most of safety-critical sectors [19]. For example, Weaver [21] in his thesis presents arguments that reflect the contribution of software to the safety-critical system. Audsley *et al.* [4] present an argument based on correctness of the specification mapping i.e., translation from the system specifications into a model and subsequently into a code. Our work in contrast focuses on deriving a safety case that argues along the hierarchical structure of systems in model-based development and traces the safety requirements on the model via the inferred system structure to the code.

With the increased use of model-based development in safety-critical applications, the integration of safety cases into such approaches has become an important research topic. For example, Chen *et al.* [7] introduce an integration of model-based engineering with safety analysis and safety case to help in assessing decisions in system design of automotive embedded systems. Hause and Thom [12] describe how SysML and UML can be used to model system requirements and how the safety requirement and other system elements identified in system design were used to construct the safety case. However, the focus here is on extending the modelling framework to represent safety cases using the applied notation. Rushby [17] also uses automated theorem proving technology (based on the Yices SMT solver) to make a safety argument, but does not construct a detailed safety case. Moreover, his analysis starts with a manually constructed logic-based model of the system, whose connection to the underlying code remains unclear. In contrast, we focus on showing safety of the system on the code level and recover the slices of the system architecture to identify where in the system the safety requirements are established.

Most safety cases are constructed manually as no advanced tools are available to support the automatic safety case construction [11]. However, a manual safety case construction [8] is far from satisfactory as it is a time-consuming and error-prone process. Most existing safety case construction tools only provide basic drawing support a

la “boxes and arrows”. For example, GSN: ASCE v3.5 from Adelard [1], the University of York Freeware Visio Add-on and GSNCaseMaker [2] are graphical tools for creating a safety case by means of a drag and drop interface based on a commercial drawing tool. Obviously, tools supported by automated analyzers such as AUTOCERT are needed to produce the complex safety arguments for software. In our work, we integrate formal analysis with the commercial safety case tool (i.e., Adelard’s ASCE tool [1]) to automatically construct the safety case. Parallel to the work on safety cases described here, we have also used the same underlying information to create safety explanations in a textual form suitable for code reviews [10]. However, this work does not yet extend to the model-based reasoning level described here.

6 Conclusions and Future Work

We have described an approach where the hierarchical structure of systems in model-based development drives the construction of a hierarchical safety case for the generated code. Here, assurance is not implied by the trust in the generator but follows from a formal analysis of the code. The analysis is based on a set of formal safety requirements and provides formal proofs as evidence. We believe greater confidence in the assurance claim can be placed if the rationale behind the validity of the transition from the model to the program can be shown. We thus make explicit argument over the correct transition from the model level representation to the source level representation, including an argument over the formalization of the requirement. We show how the external assumptions on the systems input signals are used in establishing the safety of the program wrt. the given safety requirement. Like Rushby [17], we believe that a safe design will have ensured that the assumptions are valid. Moreover, Littlewood *et al.* in [15] describe that the probability of a claim, which has been shown by a formal proof, being false, is very low, when the assumptions and evidence are valid. We thus believe that formal methods can provide the highest level of assurance when they are combined with explicit safety arguments as the ones we derived here.

The work described here is still in progress. We are currently in the process of completing the automatic construction of the safety cases. Currently, we only consider nominal component behavior, but our approach could also be applied to the off-nominal case, provided that appropriate safety requirements for the off-nominal modes can be identified. So far, we have applied our technique only to flight code generated by Real-Time Workshop from hierarchical Simulink models. We are confident that the same approach can be applied to other modelling systems and generators as well. Future work will focus on complementary safety cases that argue the safety of the certification framework itself, in particular the safety of the underlying safety logic (the language semantics and the safety policy). We believe that the result of our research will clearly communicate the safety claims, key safety requirements, and evidence required to trust the generated code.

All necessary clearances for the publication of this paper have been obtained. If accepted, the authors will prepare the final manuscript in time for inclusion in the conference proceedings and will present the paper at the conference.

References

1. ASCE home page (2007), <http://www.adelard.com/web/hnav/ASCE>
2. CET GSNCaseMaker (2007), <http://www.esafetycase.com>
3. Real-TimeWorkshop Embedded Coder 5.1. (2007), <http://www.mathworks.com/products/rtwembedded>
4. Audsley, N. C., Bate, I. J., Crook-Dawkins, S. K.: Automatic Code Generation for Airborne Systems. In: Proc. of the IEEE Aerospace Conference, pp. 8-15, IEEE (2003).
5. Basir, N., Denney, E., Fischer, B.: Constructing a Safety Case for Automatically Generated Code from Formal Program Verification Information, In: Proc. SAFECOMP 08, LNCS 5219, pp. 249-262, Springer, (2008).
6. Basir, N., Denney, E. and Fischer, B.: Deriving safety cases from automatically constructed proofs. In: Proc. 4th IET Intl. Conf. on System Safety. London, (2009).
7. Chen, D.-J. , Johansson, R., Lönn, H., Papadopoulos, Y. , Sandberg, A., Törner, F. and Törngren, M.: Modelling support for design of safety-critical automotive embedded Systems. In: Proc. SAFECOMP 2008, LNCS 5219, pp. 72-85, Springer, 2008.
8. Cockram, T. and Lockwood, B.: Electronic Safety Cases: Challenges and Opportunities. In: Proc. of Safety Critical Systems Symposium'03, Springer, (2003).
9. Denney, E., Fischer, B.: A generic annotation inference algorithm for the safety certification of automatically generated code. In: Proc. 5th Conf. Generative Programming and Component Engineering, pp. 121-130, ACM, (2006).
10. Denney, E., Fischer, B.: A Verification-Driven Approach to Traceability and Documentation for Auto-Generated Mathematical Software. ASE 2009 pp. 560-564, (2009).
11. EUROCONTROL: The EUR RVSM Pre-Implementation Safety Case. (2001).
12. Hause, M.C. and Thom, F.: Integrated safety strategy to model driven development with SysML. In: Proc. 2th IET Conf., 2007(CP532):124-129, (2007).
13. Kelly, T. P.: Arguing Safety a Systematic Approach to Managing Safety Cases. PhD Thesis, University of York, (1998).
14. Leveson, N.G., Cha, S.S. and Shimeall, T.J.: Safety Verification of ADA Programs using Software Fault Trees. IEEE Software, 8(4):48-59, (1991).
15. Littlewood, B. and Wright, D.: The Use of Multilegged Arguments to Increase Confidence in Safety Claims for Software-Based Systems: A Study Based on a BBN Analysis of an Idealized Example, IEEE Trans. Software Eng. 33(5):347-365, (2007).
16. O'Halloran, C.: Model based code verification. ICFEM 2003, LNCS 2885, pp. 16-25, (2003).
17. Rushby, J.: A Safety-Case Approach For Certifying Adaptive Systems. In: AIAA Infotech@Aerospace Conference, (2009).
18. Schloegel, K., Oglesby, D., Engstrom, E., and Bhatt, D.: Composable code generation for model-based development. In: Proc. Intl. Workshop on Software Compilers for Embedded Systems, LNCS 2826, pp. 211-225, (2003).
19. UK Ministry of Defence: 00-56 Safety Management Requirements for Defence Systems, Issue 4, (2007).
20. Vallado, D. A.: Fundamentals of Astrodynamics and Applications, 2nd ed., Microcosm Press and Kluwer Academic Publishers, (2001).
21. Weaver, R. A.: The Safety of Software-Constructing and Assuring Arguments. PhD Thesis, University of York, (2003).
22. Weiss, K.A.: Component-Based Systems Engineering for Autonomous Spacecraft. MSc Thesis, Massachusetts Institute of Technology, (2003).