

# Combining MILS with Contract-Based Design for Safety and Security Requirements

Alessandro Cimatti<sup>1</sup>, Rance DeLong<sup>2</sup>, Davide Marcantonio<sup>1</sup>,  
and Stefano Tonetta<sup>1</sup> (✉)

<sup>1</sup> FBK-irst, Trento, Italy

{cimatti,marcantonio,tonettas}@fbk.eu

<sup>2</sup> The Open Group, Reading, UK

r.delong@opengroup.org

**Abstract.** The distributed MILS (D-MILS) approach to high-assurance systems is based on an architecture-driven end-to-end methodology that encompasses techniques and tools for modeling the system architecture, contract-based analysis of the architecture, automatic configuration of the platform, and assurance case generation from patterns. Following the MILS (“MILS” was originally an acronym for “Multiple Independent Levels of Security”. Today, we use “MILS” as a proper name for an architectural approach and an implementation framework, promulgated by a community of interested parties, and elaborated by ongoing MILS research and development efforts.) paradigm, the architecture is pivotal to define the security policy that is to be enforced by the platform, and to design safety mechanisms such as redundancies or failures monitoring. In D-MILS we enriched these security guarantees with formal reasoning to show that the global system requirements are met provided local policies are guaranteed by application components. We consider both safety-related and security-related requirements and we analyze the decomposition also taking into account the possibility of component failures. In this paper, we give an overview of our approach and we exemplify the architecture-driven paradigm for design and verification with an example of a fail-secure design pattern.

**Keywords:** MILS · Contract-based design · Safety and security · Formal verification

## 1 Introduction

The *MILS architectural approach* [6] to the design and implementation of critical systems involves two principal phases: the development of an abstract architecture intended to achieve the stated purpose, and the implementation of that architecture on a robust technology platform. During the first phase, essential properties are identified that the system is expected to exhibit, and the contributions to the achievement of those properties by the architectural structure and by the behavioural attributes of key components are analyzed and justified.

© Springer International Publishing Switzerland 2015

F. Koornneef and C. van Gulijk (Eds.): SAFECOMP 2015 Workshops, LNCS 9338, pp. 264–276, 2015.

DOI: 10.1007/978-3-319-24249-1\_23

Safety and security are more and more intertwined problems. The potential impact of security threats on safety-critical systems is increasing due to the interconnections of systems. Safety, security, and dependability are emergent behavioural properties of a system interacting with its environment. The MILS approach leverages system architecture to support vital system-level properties. The architecture reflects an intended pattern of information flow and causality referred to as the *policy architecture*, while key components of the architecture enforce *local policies* through specific behavioural properties. By reasoning compositionally over the components about the policy architecture and the local policies, many useful system-level properties may be established.

The *MILS platform* provides the technology for the concrete realisation of an abstract system architecture. A *separation kernel* [31,33], the underlying foundational component of the MILS platform, is used to establish and enforce the system architecture according to its configuration data.

The assurance of a system’s properties depends not only on the analysis of its design but on the correct implementation and deployment of that design. The configuration of the separation kernel must faithfully implement the specified architecture. This is guaranteed by the *MILS platform configuration compiler* that is driven by a model of the architecture and the constraints of the target platform to synthesize viable and semantically correct configuration data corresponding to the specified architecture.

In this paper, we give an overview of the integration of the MILS approach with contract-based reasoning developed in the D-MILS project [1]. The approach relies on the OCRA tool [13] to formally prove that the global system requirements are met, provided local policies are guaranteed by application components. We consider both safety-related and security-related requirements and we analyze the decomposition also taking into account the possibility of component failures. We exemplify the architecture-driven approach on the Starlight Interactive Link example [5], extended with a safety mechanism in order to take into account the possibility of component failures.

The rest of the paper is organized as follows: in Sect. 2, we give an overview of D-MILS project; in Sect. 3, we detail how the MILS approach has been extended with a contract-based design of the architecture and the related tool support; in Sect. 4, we describe how we extended the Starlight example and the related analysis of contract refinement; in Sect. 5, we give an overview of the related work, while we conclude in Sect. 6.

## 2 Overview of D-MILS

The D-MILS concept extends the capacity of MILS to implement a single unified policy architecture to a network of separation kernels [28,29]. To accomplish this, each separation kernel is combined with a new MILS foundational component, the *MILS networking system* (MNS), producing the effect of a distributed separation kernel. In the D-MILS Project [1] we have employed *Time-Triggered Ethernet* (TTE) [32] as the MILS “backplane”, permitting us to extend the

robustness and determinism benefits of a single *MILS node* to the network of *D-MILS nodes*, referred to as the *distributed MILS platform*<sup>1</sup> [26,27].

Since D-MILS systems are intended for critical applications, assurance of the system's critical properties is a necessary byproduct of its development. In such applications, evidence supporting the claimed properties must often be presented for consideration by objective third-party system certifiers. To achieve assurance requires diligence at all phases of design, development, and deployment; and, at all levels of abstraction: from the abstract architecture to the details of configuration and scheduling of physical resources within each separation kernel and within the TTE interfaces and switches. Correct operation of the deployed system depends upon the correctness of the configuration details, of the component composition, of key system components<sup>2</sup>, and of the D-MILS platform itself. Configuration is particularly challenging, because the scalability that D-MILS is intended to provide causes the magnitude of the configuration problem to scale as well. The concrete configuration data and scheduling details of the numerous separation kernels and of the TTE are at a very fine level of granularity, and must be complete, correct, and coherent.

The only reasonable prospect of achieving these various aspects of correctness, separately and jointly, is through pervasive and coordinated automation as embodied in the D-MILS tool chain. Inputs to the tool chain include, a declarative model of the system expressed in our own MILS dialect of the Architecture Analysis and Design Language (AADL) [18], facts about the target hardware platform, properties of separately developed system components, designer-imposed constraints and system property specifications, and human guidance to the construction of the assurance case. Components of the tool chain perform parsing of the languages [19], transformations among the various internal forms [20,21], analysis and verification [24], configuration data synthesis and rendering [25], and pattern-based assurance case construction [22,23]. Outputs of the tool chain include, proofs of specified system properties, configuration data for the D-MILS platform, and an assurance case expressed in Goal Structuring Notation (GSN) [2]. We say that D-MILS provides not only a robust and predictable platform for system implementation, but also an end-to-end and top-to-bottom method supported by extensive automation.

### 3 Architecture-Driven Integration of the MILS Approach and Contract-Based Design

In this paper we focus on the integration of the MILS architectural approach with contract-based design and analysis. Both MILS and contract-based approaches

---

<sup>1</sup> Our D-MILS Platform is composed of the LynxSecure Separation Kernel from Lynx Software Technologies, France, and TTE from TTTech, Austria.

<sup>2</sup> The D-MILS Project regards proof of component correctness to a specification as a “solved problem” and focusses on the correctness of the composition of components’ specifications, and of the configuration of the D-MILS platform.

focus on architecture, and do so in a complementary way. MILS regards information flow policy as an abstraction of architecture, and seeks to maximize the correspondence between architectural structure and the desired information flow policy of a system, which may rely on the behavior of key components to enforce local policies that further restrict the maximal information flow permitted by the architecture. The contract-based approach employs formalization and a method to prove that the architecture decomposition represented in the set of contracts of the components is a proper refinement of the system requirements. Contracts specify the properties that the components individually, and the system as a whole, are expected to guarantee, and the assumptions that their respective environments must meet. Formal verification techniques are used to check that the derivation of the local policies from the system requirements is correct.

An architecture is only as valuable as the integrity of its components and connections. Recognizing the importance of integrity, MILS provides an implementation platform that can be configured to the “shape” of the architecture by initializing it with specific configuration data compiled to embody the global information flow policy.

The two methods are complementary and their combination yields a strong result. The contract-based method proves that the composition of components that satisfy their contracts will meet the system requirements, provided that their integrity is protected. The MILS platform guarantees the integrity of components and their configured connections, preventing interference that could cause a verified component to fail to satisfy its contract<sup>3</sup>.

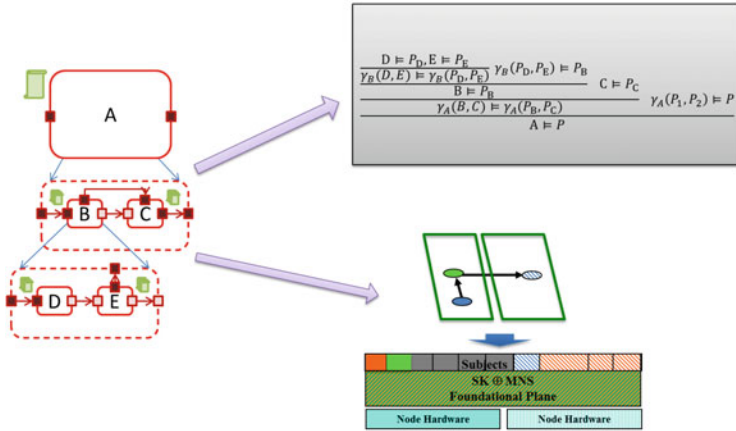
In Fig. 1, we show the approach applied to an abstract example. The system  $A$  is decomposed into subsystems  $B$  and  $C$ , and  $B$  in turn is decomposed into  $D$  and  $E$ . Each component is enriched with a contract (represented here by green scrolls). If the contract refinement is correct, we have associated with the architecture a formal proof that the system is correct provided that the leaf components ( $D$ ,  $E$ , and  $C$ ) satisfy their contracts. Namely, if  $D$  and  $E$  satisfy their contracts ( $D \models P_D, E \models P_E$ ) and the contract refinement of  $B$  is correct ( $\gamma_B(P_D, P_E) \preceq P_B$ ), then the composition of  $D$  and  $E$  satisfies the contract of  $B$  ( $\gamma_B(D, E) \models P_B$ ). Moreover, if  $C$  satisfies its contract ( $C \models P_C$ ) and the contract refinement of  $A$  is correct ( $\gamma_A(P_B, P_C) \preceq P_A$ ), then the composition of  $B$  and  $C$  satisfies the contract of  $A$  ( $\gamma_A(B, C) \models P_A$ ).

In MILS terms, the architecture defines three subjects ( $D$ ,  $E$  and  $C$ ) and prescribes that the only allowed communications must be the ones between  $D$  and  $E$  and between  $E$  and  $C$ . This is translated into a configuration for the D-MILS platform (taking into account other deployment constraints in terms of available resources), which in this example encompasses two MILS nodes.

### 3.1 Tool Support for Contract-Based Reasoning

In D-MILS, the architecture is specified in a variant of AADL, called MILS-AADL, similar to the SLIM language developed in the COMPASS project [7].

<sup>3</sup> For the purpose of our work we assume that components can be constructed and verified to satisfy their contracts.



**Fig. 1.** The architecture is used for (1) formal reasoning to prove that the system requirements are assured by the local policies, (2) configuration of the platform to ensure the global information flow policy and the integrity of the architecture.


The COMPASS tool set has been extended in order to support the new language and to enrich the components with annotations that specify different verification properties such as contracts. The language used to specify the component contracts is the one provided by the OCRA tool [13]. It consists of a textual human-readable version of a First-Order Linear-time Temporal Logic. The logic has been extended in D-MILS to support uninterpreted functions, i.e. functional symbols that do not have a specific interpretation but are used to abstract procedures and the related results (such as CRC checksum or encryption), or to label data with user-defined tags (such as “is\_high” or “low-level”, etc.).

Such a very expressive language required the development of effective techniques to reason about the contracts. To this purpose the engine undertakes to prove the contract refinement. The refinement is first translated by OCRA into a set of entailment problems in temporal logic. nuXmv [11] translates this into a liveness model-checking problem with a classic automata-theoretic approach [37]. The resulting problem requires proving that a certain liveness condition can be visited only finitely many times along an (infinite) execution. This problem is in turn reduced to proving an invariant on the reachable states with the K-liveness techniques described in [17]. This has been extended to infinite-state systems and to take into account real-time aspects in [15]. Finally, the invariant is proved with an efficient combination of induction-based reasoning, explicit-state search, and predicate abstraction, extending the IC3 algorithm [9] to the infinite-state case, as described in [14].

## 4 Starlight Example

### 4.1 Architecture

In this section, we exemplify the approach on an example taken from the literature [5,12]. The Starlight Interactive Link is a dispatching device developed by the Australian Defense Science and Technology Organization to allow users to establish simultaneous connections to high-level (classified) and low-level networks. The idea is that the device acts as a switch that the user can control to dispatch the keyboard output to either a high-level server or to a low-level server. The user can use the low-level server to browse the external world, send messages, or have data sent to the high-level server for later use.



**Fig. 2.** Architecture of the D-MILS Starlight example.

Figure 2 shows the architecture of the Starlight Interactive Link as formalized in D-MILS. The components  $H$  and  $L$  represent respectively the high-level and low-level networks. The low-level network can exchange information with the external world. The component  $D$  represents the Starlight device, which receives commands from the user and dispatches the commands to  $H$  or to  $L$  based on an internal state. The state is changed with two *switch* commands, namely *switch\_to\_low* and *switch\_to\_high*. The original architecture has only the blue components, with  $D$  in place of  $E$ . We extended this architecture with a safety mechanism to make the system “fail-secure” with respect to failures of the dispatcher: the dispatcher is extended with a monitor  $M$ ; the communication of the dispatcher to  $L$  is filtered by  $M$  that in case of failure of  $D$  blocks the communication. To avoid confusion we refer to the actual device that is filtered by  $M$  as the dispatcher ( $D$ ), while to the component consisting of  $D$  and  $M$  as the extended dispatcher  $E$ .

## 4.2 System Contract

The architecture has been enriched with contracts that formalize the functional requirements to ensure that the system responds correctly to the user commands, and the security requirement that there is no leakage of high-level data. Here, we focus on the latter, which says:

**Req-Sys-secure:** No high-level data shall be sent by  $L$  to the external world.

The architecture ensures Req-Sys-secure assuming the following requirement on the user:

**Req-User-secure:** The user shall switch the dispatcher to high before entering high-level data.

Moreover, we consider the following safety requirement:

**Req-Sys-safe:** No single failure shall cause a loss of Req-Sys-secure.

We formalized the requirements of the system and of the components using OCRA contracts. In the following, we use the concrete syntax accepted by the tool. We briefly clarify the used notation: “and”, “or”, “not”, “implies” are standard Boolean operators; “always”, “never”, “since”, “in the past” are standard temporal operators of LTL with past also referred to with the mathematical notation  $G$ ,  $G\neg$ ,  $S$ ,  $O$ ; “last\_data” is a built function to refer to the last data passed by the event of a event data port; italics names refer to ports or uninterpreted functions declared in the model.

The requirements Req-Sys-secure and Req-User-secure have been formalized into the FO-LTL formulas:

**Formal-Sys-secure:** never  $is\_high(last\_data(outL))$

**Formal-User-secure:** always  $((is\_high(last\_data(cmd)))$  implies  $((not$   $switch\_to\_low)$  since  $switch\_to\_high))$

Note that the formalization of Req-User-secure improves the informal requirement, which is not precise. A literal formalization would be:

**Formal-User-secure-wrong:** always  $((is\_secure(last\_data(cmd)))$  implies (in the past  $switch\_to\_high))$

but this is wrong, because we have to ensure that the last switch was a  $switch\_to\_high$ , without a more recent  $switch\_to\_low$ <sup>4</sup>. We can actually improve the informal requirement as:

**Req-User-secure-new:** Whenever the user sends commands with high data, she shall previously issue a  $switch\_to\_high$  and no  $switch\_to\_low$  since the last  $switch\_to\_high$ .

which is formalized by Formal-User-secure.

---

<sup>4</sup> As suggested by one of the reviewers, in an alternative model, we could use only one event data instead of two switch events and ensure that the last switch was low.

Note that while Req-Sys-secure is a requirement on the implementation of the Starlight system, Req-User-secure is actually a requirement on its environment (the user). This is reflected by the system contract specification, which sets Formal-Sys-secure as the guarantee and Formal-User-secure as the assumption of the system contract.

### 4.3 Component Contracts

The dispatcher ensures the system security requirement with the following local requirement:

**Req-D-low-mode:** The dispatcher shall send commands to  $L$  only if the last switch was a *switch\_to\_low* and the input command has been received after.

formalized into:

**Formal-D-low-mode:** always ( $cmdL$  implies (((not *switch\_to\_high*) since *switch\_to\_low*) and ((not *switch\_to\_low*) since  $cmd$ )))

In order to fulfill requirement Req-Sys-safe, we also filter the commands to  $L$  by a monitor  $M$ , which has a requirement **Req-M-low-mode** identical to Req-D-low-mode, and formalized in the same way. Thus,  $D$  passes also the switches to the monitor and must ensure the following requirement:

**Req-D-fw-switch:** Whenever the dispatcher receives a *switch\_to\_high*, it shall pass it to  $M$  before doing any other actions and it sends a *switch\_to\_low* to  $M$  only if the last received switch was a *switch\_to\_low*.

formalized into:

**Formal-D-fw-switch:** always ((*switch\_to\_high* implies ((not ( $cmdH$  or  $cmdL$  or *return* or *monitor\_switch\_to\_low*)) until *monitor\_switch\_to\_high*)) and (*monitor\_switch\_to\_low* implies ((not *switch\_to\_high*) since *switch\_to\_low*)));

Finally, in order to make the refinement correct, we must require all components to not invent high data. We express this by requiring that  $D$ ,  $M$ , and  $L$  only pass the data that they have received. Thus, for  $D$ , we require that:

**Req-D-data:**  $D$  shall pass to  $cmdL$  only the data that has been received with last  $cmd$ .

formalized into:

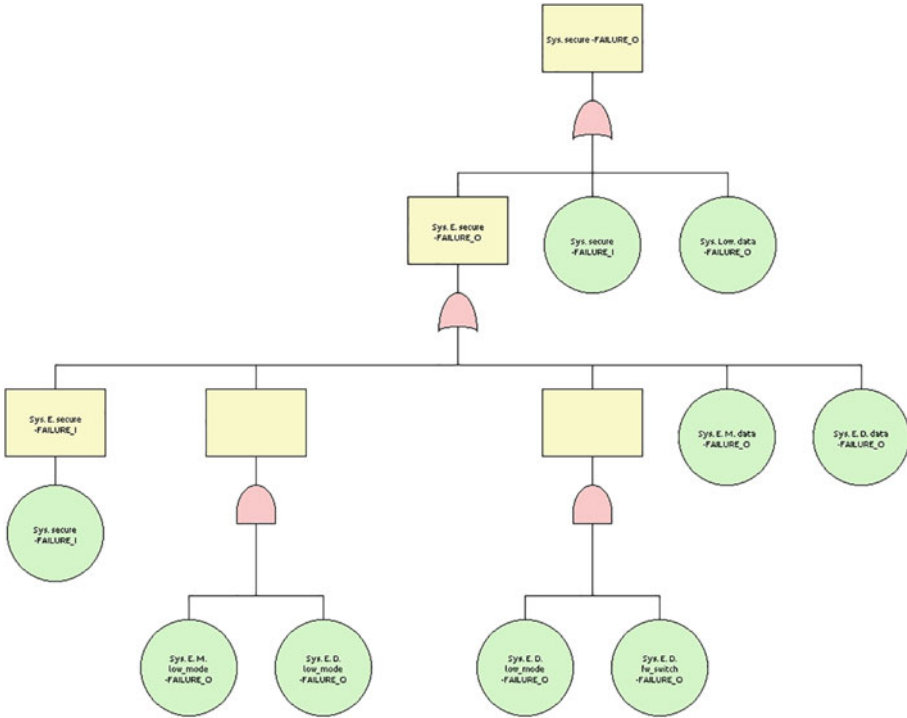
**Formal-D-data:** always (( $cmdL$  implies ((in the past  $cmd$ ) and (last\_data( $cmdL$ ) = last\_data( $cmd$ ))))))

The requirements **Req-M-data** and **Req-L-data**, of  $M$  and  $L$  respectively, are analogous. Note that these formulas are actually guarantees of corresponding contracts, without assumptions (i.e. assumptions equal to *true*).



### 4.4 Analysis Results

Given the above contract specifications, OCRA can prove the system Req-Sys-secure assuming Req-User-secure is correctly refined by the contracts of  $D$ ,  $M$ , and  $L$  (see [16] for more details on the technique). One can also show that by using Formal-User-secure-wrong instead of Formal-User-secure the refinement is not correct and yields a counterexample trace execution.



**Fig. 3.** Fault-tree generated from the contract refinement. Events are labeled with the name of the component instance followed by the name of the contract, followed either by FAILURE.O, which represents the failure of the component to satisfy the contract, or by FAILURE.I, which represents the failure of the component environment to satisfy the contract assumption.

In order to prove Req-Sys-safe, we use OCRA to produce a fault tree showing the dependency of the system failure on the failure of the components (see [8] for more details on the technique). The generated fault tree is exhibited in Fig. 3. It shows that neither Req-D-low-mode nor Req-M-low-mode are single points of failure. Instead, Req-D-data, Req-M-data, Req-L-data are single points of failure. While the failure of Req-L-data does not represent real threats since  $L$  never receives high data, the failure of Req-D-data and Req-M-data could result

in  $D$  or  $M$  sending information that had been temporary stored in a buffer used for handling multiple requests or in a cache for improving performance. This can be solved for example by ensuring that such memories are deleted before every *switch\_to\_low* is completed.

## 5 Related Work

Security-by-contract is an approach proposed in [30] to increase the trust in code downloaded on mobile applications. The work proposes a framework where downloaded code can be checked according to a security contract. With respect to this work, there is no focus on the system architecture, the refinement of contracts, or safety analysis taking into account component failures.

Information flow contracts are supported in SPARK, a safety-critical subset of Ada [3, 4]. The SPARK contracts are specified at software level on procedures. So, in principle, they are complementary to our approach, which focuses more on the system-level architecture. As for future work, we will consider to extend the approach with information flow contracts. Currently, the information flow can be specified only at coarse level with the connections in the architecture. To our knowledge, there are no works combining SPARK information flow contracts with safety analysis.

In [10], an avionic case-study architecture is formalized in Alloy and analyzed with respect to safety and security requirements. Similarly to our approach, first-order logic is used to formalize the requirements, although Alloy does not support temporal operators. The case study formalizes also security attacks that are not present in our example. Different to our approach, the failures and security attacks are explicitly modeled, while in our case we exploit a feature of OCRA to automatically inject the failures starting from the nominal contract specification. Our conjecture is that the same case study of [10] can be formalized in MILS-AADL or directly in OCRA with the possibility of checking contract refinement and performing contract-based fault-tree analysis.

Another case study on validation of safety and security requirements has been presented in [35], but it focuses on testing.

Fault trees and FMEA have been extended in [34, 36] to consider also security aspects. Different to our approach and other model-based safety analysis techniques, these works are not based on the automatic generation of fault trees and FMEA tables from the system design.

## 6 Conclusions

In this paper, we briefly overview the approach to safety and security undertaken in D-MILS and we describe a small example of the D-MILS approach to the verification of the system architecture with respect to safety and security requirements. The example is based on the Starlight device that switches commands between high-level and low-level servers. The requirements of the system and its components have been formalized using OCRA contracts, their refinement has

been verified and analyzed taking into account the failure of components. This is a preliminary application of the methodology, which will be further evaluated in the D-MILS project demonstrators. In the future, we would like to integrate contracts and their analysis with finer-grained information flow properties as do the SPARK contracts discussed in [3,4].

**Acknowledgments.** This work was performed on the D-MILS project (“Distributed MILS for Dependable Information and Communication Infrastructures”, European Commission FP7 ICT grant no. 318772), with our partners fortiss, Verimag, RWTH Aachen, U of York, Frequentis, Lynx, TTTech, and INRIA, funded partially under the EC’s Seventh Framework Programme.

## References

1. D-MILS Project. <http://www.d-mils.org/>
2. GSN community standard. Technical report, Origin Consulting (York) Limited (2011)
3. Amtoft, T., Hatcliff, J., Rodríguez, E.: Precise and automated contract-based reasoning for verification and certification of information flow properties of programs with arrays. In: Gordon, A.D. (ed.) ESOP 2010. LNCS, vol. 6012, pp. 43–63. Springer, Heidelberg (2010)
4. Amtoft, T., Hatcliff, J., Rodríguez, E., Robby, Hoag, J., Greve, D.: Specification and checking of software contracts for conditional information flow. In: Hardin, D.S. (ed.) Design and Verification of Microprocessor Systems for High-Assurance Applications, pp. 229–245. Springer, New York (2010)
5. Anderson, M., North, C., Griffin, J., Milner, R., Yesberg, J., Yiu, K.: Starlight: interactive link. In: 12th Annual Computer Security Applications Conference, pp. 55–63 (1996)
6. Boettcher, C., DeLong, R., Rushby, J., Sifre, W.: The MILS component integration approach to secure information sharing. In: 27th AIAA/IEEE Digital Avionics Systems Conference, St. Paul, MN, October 2008
7. Bozzano, M., Cimatti, A., Katoen, J., Nguyen, V.Y., Noll, T., Roveri, M.: Safety, dependability and performance analysis of extended AADL models. *Comput. J.* **54**, 754–775 (2011)
8. Bozzano, M., Cimatti, A., Mattarei, C., Tonetta, S.: Formal safety assessment via contract-based design. In: Cassez, F., Raskin, J.-F. (eds.) ATVA 2014. LNCS, vol. 8837, pp. 81–97. Springer, Heidelberg (2014)
9. Bradley, A.R.: SAT-based model checking without unrolling. In: Jhala, R., Schmidt, D. (eds.) VMCAI 2011. LNCS, vol. 6538, pp. 70–87. Springer, Heidelberg (2011)
10. Brunel, J., Rioux, L., Paul, S., Faucogney, A., Vallée, F.: Formal safety and security assessment of an avionic architecture with alloy. In: ESSS, pp. 8–19 (2014)
11. Cavada, R., Cimatti, A., Dorigatti, M., Griggio, A., Mariotti, A., Micheli, A., Mover, S., Roveri, M., Tonetta, S.: The NUXMV symbolic model checker. In: Biere, A., Bloem, R. (eds.) CAV 2014. LNCS, vol. 8559, pp. 334–342. Springer, Heidelberg (2014)
12. S. Chong and R. Van Der Meyden, Using architecture to reason about information security (2014). arXiv preprint [arXiv:1409.0309](https://arxiv.org/abs/1409.0309)

13. Cimatti, A., Dorigatti, M., Tonetta, S.: OCRA: a tool for checking the refinement of temporal contracts. In: ASE, pp. 702–705 (2013)
14. Cimatti, A., Griggio, A., Mover, S., Tonetta, S.: IC3 modulo theories via implicit predicate abstraction. In: Ábrahám, E., Havelund, K. (eds.) TACAS 2014 (ETAPS). LNCS, vol. 8413, pp. 46–61. Springer, Heidelberg (2014)
15. Cimatti, A., Griggio, A., Mover, S., Tonetta, S.: Verifying LTL properties of hybrid systems with K-LIVENESS. In: Biere, A., Bloem, R. (eds.) CAV 2014. LNCS, vol. 8559, pp. 424–440. Springer, Heidelberg (2014)
16. Cimatti, A., Tonetta, S.: Contracts-refinement proof system for component-based embedded systems. *Sci. Comput. Program.* **97**, 333–348 (2015)
17. Claessen, K., Sörensson, N.: A liveness checking algorithm that counts. In: FMCAD, pp. 52–59 (2012)
18. Specification of MILS-AADL. Technical report D2.1, Version 2.0, D-MILS Project, July 2014. <http://www.d-mils.org/page/results>
19. D2.2 translation of mils-aadl into formal architectural modeling framework. Technical report D2.2, Version 1.2, D-MILS Project, February 2014. <http://www.d-mils.org/page/results>
20. Intermediate languages and semantics transformations for distributed mils - part 1. Technical report D3.2, Version 1.2, D-MILS Project, February 2014. <http://www.d-mils.org/page/results>
21. Intermediate languages and semantics transformations for distributed mils - part 2. Technical report D3.3, Version 1.0, D-MILS Project, July 2014. <http://www.d-mils.org/page/results>
22. Compositional assurance cases and arguments for distributed mils. Technical report D4.2, Version 1.0, D-MILS Project, April 2014. <http://www.d-mils.org/page/results>
23. Integration of formal evidence and expression in mils assurance case. Technical report D4.3, Version 0.7, D-MILS Project, March 2015. <http://www.d-mils.org/page/results>
24. Compositional verification techniques and tools for distributed mils—part 1. Technical report D4.4, Version 1.0, D-MILS Project, July 2014. <http://www.d-mils.org/page/results>
25. Distributed mils platform configuration compiler. Technical report D5.2, Version 0.2, D-MILS Project, March 2014. <http://www.d-mils.org/page/results>
26. Extended separation kernel capable of global exported resource addressing. Technical report D6.1, Version 2.0, D-MILS Project, March 2014. <http://www.d-mils.org/page/results>
27. Mils network system supporting TTEthernet. Technical report D6.3, Version 1.0, D-MILS Project, March 2014. <http://www.d-mils.org/page/results>
28. R. DeLong, Commentary on the MILS Network Subsystem Protection Profile. Technical report, Version 0.31, September 2011
29. DeLong, R., Rushby, J.: Protection Profile for MILS Network Subsystems in Environments Requiring High Robustness, Version 0.31, September 2011
30. Dragoni, N., Massacci, F., Walter, T., Schaefer, C.: What the heck is this application doing? - a security-by-contract architecture for pervasive services. *Comput. Secur.* **28**, 566–577 (2009)
31. Information Assurance Directorate, National Security Agency, U.S. Government Protection Profile for Separation Kernels in Environments Requiring High Robustness, Fort George G. Meade, MD 20755–6000, Version 1.03, June 2007

32. Kopetz, H., Ademaj, A., Grillinger, P., Steinhammer, K.: The time-triggered ethernet (TTE) design. In: 8th IEEE International Symposium on Object-oriented Real-time distributed Computing (ISORC), Seattle, Washington (2005)
33. Rushby, J.: The design and verification of secure systems. In: Eighth ACM Symposium on Operating System Principles, Asilomar, CA, December 1981, pp. 12–21 (1981). (ACM Operating Systems Review, Vol. 15, No. 5)
34. Schmittner, C., Gruber, T., Puschner, P., Schoitsch, E.: Security application of failure mode and effect analysis (FMEA). In: Bondavalli, A., Di Giandomenico, F. (eds.) SAFECOMP 2014. LNCS, vol. 8666, pp. 310–325. Springer, Heidelberg (2014)
35. Sojka, M., Krec, M., Hanzálek, Z.: Case study on combined validation of safety & security requirements. In: SIES, pp. 244–251 (2014)
36. Steiner, M., Liggesmeyer, P.: Combination of safety and security analysis - finding security problems that threaten the safety of a system. In: SAFECOMP Workshop DECS (2013)
37. Vardi, M.Y.: An automata-theoretic approach to linear temporal logic. In: Moller, F., Birtwistle, G. (eds.) LC 1995. LNCS, vol. 1043, pp. 238–266. Springer, Heidelberg (1995)