

Formal Safety and Security Assessment of an Avionic Architecture with Alloy

Julien Brunel
ONERA/DTIM
F-31055 Toulouse, France
julien.brunel@onera.fr

Laurent Rioux Stéphane Paul
Thales Research & Technology
F-91767 Palaiseau, France
firstname.lastname@thalesgroup.com

Anthony Faucogney Frédérique Vallée
All4Tec
F-53001 Laval, France
firstname.lastname@all4tec.net

We propose an approach based on Alloy to formally model and assess a system architecture with respect to safety and security requirements. We illustrate this approach by considering as a case study an avionic system developed by Thales, which provides guidance to aircraft. We show how to define in Alloy a metamodel of avionic architectures with a focus on failure propagations. We then express the specific architecture of the case study in Alloy. Finally, we express and check properties that refer to the robustness of the architecture to failures and attacks.

1 Introduction

Nowadays, safety and security are commonly identified disciplines in system and software engineering. However it often covers the same meaning in the head of non-experts. Even if the end target may be considered similar, *i.e.*, “protect the system against unspecified behaviors during operation”, the approaches proposed by safety and security and their evaluation criteria are sometimes equivalent but sometimes quite different or even opposite.

For example, at risk management level, on the one hand, a security policy can require keeping a door closed to secure the data saved in a room, and on the other hand, the safety policy can require keeping this door open in case of fire. How is the relative importance of both concerns managed? Does safety serve security or does security serve safety or are they at the same level? Another example at development level: the safety and security disciplines are also different, handled by different teams, which most of the time do not talk to each other and where their work is based on different vocabulary and standards. What would the common language be for making better communication between safety and security teams?

Besides, the mindset of safety and security teams may be different, even opposite when we focus on knowledge sharing. The first one tries to share the most valuable knowledge in order to reduce impacts due to systems failures; the second one tries to protect its security mechanisms against publishing in order to reduce the potential exploitation probability. How can safety and security team mindsets may be harmonized and better coordinated?

A collaborative ITEA2 project named MERgE¹ has been setup to identify the discipline divergences and convergences to then demonstrate the benefit of a co-engineering activity sup-

¹<http://merge-project.eu>

ported by a collaborative modeling tool platform. In this article we focus on a solution promoted by the project to perform early validation of systems with respect to safety and security properties, taking benefit from lightweight formal methods, and more specifically Alloy [3].

This article is organized as follows. In Sect. 2, we present the case study, a system that provides guidance to aircraft. In Sect. 3, we show how to model in Alloy the architecture of the case study. Then, we express and check the expected safety properties in Sect. 4 and the security properties in Sect. 5.

2 Case study description

An approach is the last leg of an aircraft's flight, before landing. Approaches can be performed under Visual Flight Rules (VFR) or under Instrument Flight Rules (IFR). VFR are a set of regulations under which a pilot operates an aircraft in weather conditions clear enough to allow the pilot to operate the aircraft with visual reference to the ground, and by visually avoiding obstructions and other aircraft. For the aircraft operating under IFR, an Instrument Approach Procedure (IAP) defines a series of predetermined instrument-supported manoeuvres from the beginning of the initial approach to the landing (or to a point from which a landing may be made visually). There are two main classifications for IAPs: precision approaches (PA) and non-precision approaches (NPA). Precision approaches utilize both lateral and vertical information. Non-precision approaches provide lateral course information only. Instrument approaches have traditionally been supported by land-based navigation aids (i.e. the so-called "navaids"), typically Localizers, VHF Omnidirectional Range (VOR) beacons, Non-Directional Beacons (NDB), Distance Measuring Equipment (DME), or Instrument Landing System (ILS).

Since the 1960s, Global Navigation Satellite Systems (GNSS) have been experimented and deployed to bypass land-based equipment. GNSS is a system of satellites that provide autonomous geo-spatial positioning with global coverage. It allows small electronic receivers to determine their location (longitude, latitude, and altitude) to high precision (within a few metres) using time signals transmitted along a line of sight by radio from satellites. Currently, the United States NAVSTAR Global Positioning System (GPS) and the Russian GLONASS are the only global operational GNSSs. The European Union's Galileo positioning system is a GNSS in initial deployment phase, scheduled to be fully operational by 2020. Localizer Performance with Vertical guidance (LPV) is the highest precision GNSS aviation instrument approach procedure currently available without specialized aircrew training requirements. LPV is designed to provide 16 meter horizontal accuracy and 20 meter vertical accuracy 95 percent of the time. LPV landing minima are usually similar to those in an ILS.

This case-study is concerned by the architecting of a new Thales Avionics aircraft embedded system designed to support an LPV approach. Its architecture is represented by Fig. 1.

We can summarize the behavior of this sub-system as follows. Both satellites constellations (American GPS and GALILEO constellations) send the signal to SBAS processing functions. After correlations of both positions information, the SBAS sends the plane position (lateral and vertical) to two occurrences of the LPV processing function. The data produced by LPV processing functions are sent to three displays (three occurrences of a function Acquire. In each display, a comparison of the data received from LPV1 and LPV2 is performed. In case of inconsistency, an alarm is triggered by a function Monitor. The crew chooses which of the two LPV processing is used by each display (function SelectSource, not represented in Fig 1).

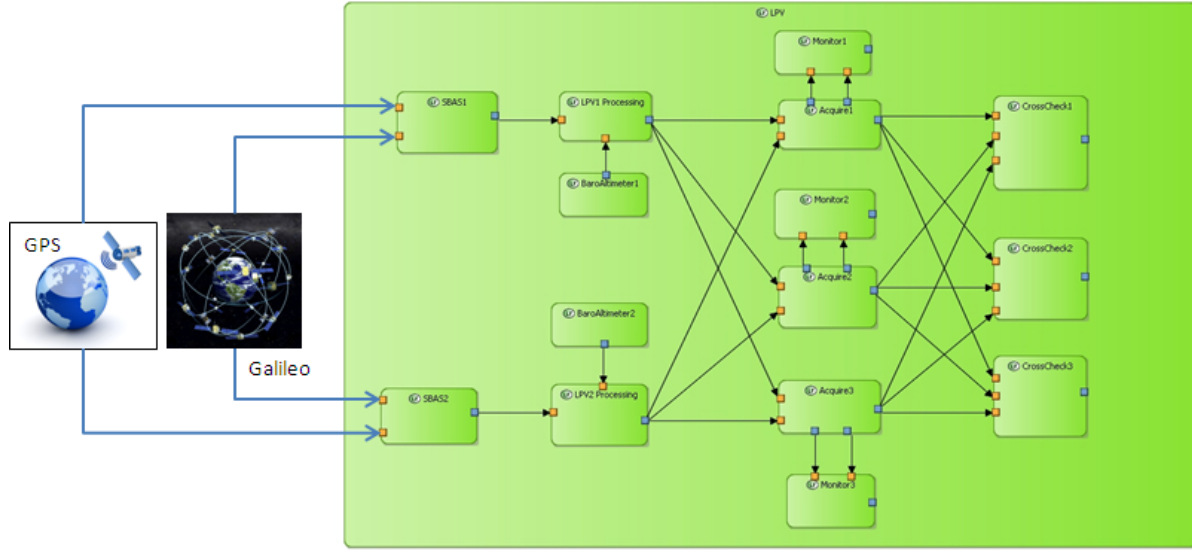


Figure 1: LPV architecture

Besides, each display receives the data computed by the two other displays. Then, the function Crosscheck compares the data of the current display with the data of the two others and resets the current display in case it differs from both other displays.

This initial architecture was designed taking into account a number of safety requirements. Two of these requirements are recalled below.

Safety requirements:

- Safety 1* *Loss of LPV capability.* No single failure must lead to the loss of LPV capability.
- Safety 2* *Misleading information integrity.* The architecture must control the value of the LPV data provided by each calculator and between each screen and find mitigation in case of erroneous data.

In this paper we will first show how to model this architecture using Alloy (Sect. 3) and how to formally prove that this architecture indeed complies to the aforementioned safety requirements (Sect. 4).

In a second step, we want to ensure that the above architecture is also resilient to a number of malevolent attacks (Sect. 5). Seven attacks have been considered as listed below.

Security attacks:

- Attack 1* One malicious GPS signal (a fake signal that SBAS considers to come from GPS).
- Attack 2* One constellation satellite signal is scramble.
- Attack 3* The RNAV ground station is neutralized, meaning that no more RNAV signal can be send to the plane.
- Attack 4* An attack combining Attack 1 and Attack 2 scenarios.
- Attack 5* An attack combining Attack 2 and Attack 3 scenarios.
- Attack 6* An attack combining Attack 1 and Attack 3 scenarios.
- Attack 7* An attack combining Attack 1, Attack 2 and Attack 3 scenarios.

The paper shows how the proposed architecture fails to comply with all the security requirements, and how the architecture can be enhanced to provide full compliance to both safety and security requirements.

3 Modelling the case study with Alloy

Alloy [3] is a formal system-modelling language amenable to automatic analyses. Alloy is implemented as a cost-less free-software tool, the Alloy Analyzer, which is programmed in Java and hence runnable on the majority of platforms. Alloy has recently been used in the context of security assessment, for instance to model JVM security constraints [5], access control policies [6], or attacks in cryptographic protocols [4]. Besides, we proposed in earlier work a preliminary study of the safety assessment of the LPV system [2]. In this article, we propose to enrich this study by considering a more complete architecture and by expressing and verifying both safety and security objectives.

The AltaRica [1] language, which is widely use for safety assessment, would have been another possible choice. However, we decided to take benefit from the model-based aspect of Alloy and its expressiveness for the specification of the properties to check. Indeed, Alloy allows to define easily the metamodel of the avionic architectures we will analyze instead of encoding them in terms of AltaRica concepts. Moreover, the specification of the properties we want to check are expressed in relational first-order logic with many features adapted to model-based reasoning.

3.1 General presentation of modelling in Alloy

Alloy is a general-purpose modelling language that is nonetheless very well adapted to the following (non-exhaustive) list of activities: abstract modelling of a problem or of a system; production of a metamodel (model corresponding to a viewpoint); analysis of a model using well-formedness or formal semantic rules; automatic generation of an instance conforming to a model, possibly according to supplementary constraints; finding interesting instances of a model. Models designed in Alloy can deal with static aspects only, or integrate also dynamic aspects, so as to check behavioral properties.

We now give a brief glance at the main concepts of the language using a simple example. Each Alloy file contains one module made of several declarations, the order of which is not important. The most important declaration is the signature which introduces a structural concept. This may be seen as a class or entity in modelling parlance. A signature should be interpreted as a set of possible instances (as it is sometimes said that a class “is” the set of objects defined by itself). Finally, every signature comes with fields that may be seen, as a first approximation, as class attributes.

```
sig Data { consumedBy : some System }
sig System {}
sig Criticality {
  concernedData : one Data,
  concernedSystem : one System
}
```

Here, we defined 3 concepts : Data, System and Criticality. As explained earlier, Alloy advocates not to delve into unnecessary details and only give information on things we want to understand or analyze. Thus, here, a system is just defined to be a set of “things”, but we do not say anything about the exact nature of its elements.

The keywords *some* or *one* give details on the multiplicity of the relation, as *1..** and *1* in UML. Here the field declarations mean that: every datum is consumed by at least one (some) system; every criticality concerns exactly one (one) data and one system. Other possible multiplicities are: *lone* which means at most one; and *set* which means any number (including 0).

Then, we can add constraints on possible instances of our model. For instance, we would like to state that every system consumes at least one datum. This can be done by writing additional facts (facts state properties that always hold, so as few facts as possible should be stated in order to avoid over-specification):

```
fact {
  // every system consumes at least one datum
  all s : System | some consumedBy.s
  // for any system which consumes a given datum, the said datum and system
  // should belong to a same unique criticality
  all d : Data | all s : System | one c : Criticality |
    c.concernedData = d and c.concernedSystem = s
}
```

Given all these definitions, the Alloy Analyzer can

- provide graphical representation of the metamodel or of an instance
- carry out some explorations (the command `run` builds instances that satisfy a given statement)
- check whether a given constraint is satisfied by all instances of the model (command `check`)

3.2 LPV model in Alloy

To represent the functional architecture of LPV (see Fig 1) in Alloy, we identified the major concepts at stake and defined an Alloy signature for each of them: Function, Port, IPort (input port), OPort (output port). Each Function has a set input of IPorts and a set output of OPort as attributes. An OPort of a function is related to a set of IPort of other functions through a flow. For a given function, the relation between its IPort and OPort, for instance expressing how failures propagate, can be given directly by an Alloy formula (a *fact* in Alloy terminology) that constrains the model. Each function and each port hold a dysfunctional status (OK, Lost, or Err). If the status of a port is Lost, it means that this port does not produce any data. If the status is Err, it means that it yields an undetected erroneous data. Moreover, each port has a value. Notice that, in our model, the value will only be useful for certain ports, representing the pilot selection, the discrepancy between both LPV processing, and the reset of displays. Concerning the other ports, we only deal with their status and the value is just ignored (a finer modelling would have been possible here but it would have led to a worthless, more complex model). The following Alloy code corresponds to these concept definitions (see also Fig. 2).

```

enum Status { OK, Err, Lost }
abstract sig Port {
    status : Status,
    value : Value
}
abstract sig IPort extends Port {}
abstract sig OPort extends Port {
    flow : set IPort,
}
abstract sig Function {
    input : set IPort,
    output : set OPort,
    status : Status
}
    
```

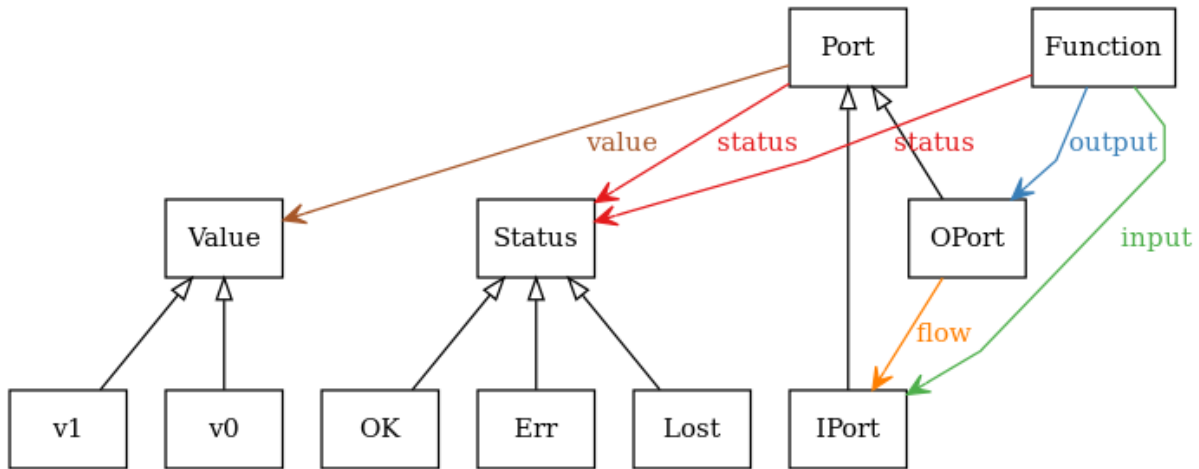


Figure 2: Case study metamodel (simplified)

We then define instances of these concepts corresponding to the LPV functional architecture. The function instances take into account the selection of the source by the crew (SelectSource), the satellite data position (GPS and Galileo), two occurrences of SBAS positioning (ComputeSBAS1, ComputeSBAS2), two occurrences of LPV processing (ComputeLPV1, ComputeLPV2), three occurrences of displays (Acquire i , $i \in \{1..3\}$), three occurrences of display resetters (Crosscheck i , $i \in \{1..3\}$) and of monitors in order to trigger an alarm, (Monitor i , $i \in \{1..3\}$). We also define the different ports of each function, and the way ports are related to each other via flows. For instance, the following Alloy code is an excerpt of the flow definition, expressing that the output port oSBAS1 is related to the input port iSBAS1 via a flow (idem for oSBAS2 and iSBAS2).

```

flow = oSBAS1 → iSBAS1 + oSBAS2 → iSBAS2 + ...
    
```

We also define some global constraints the architecture must satisfy, such as the fact that two ports related by a flow share the same status and the same value:

```

all p1, p2 : Port | p1 → p2 in flow implies p1.status = p2.status and p1.value = p2.value
    
```

We now define the relation between input and output ports inside each function in term of failure propagation. For instance, the following code expresses that the status of the output port `oDeviation1` of function `ComputeLPV1` is equal to the status of the input if the function `ComputeLPV1` is OK, is equal to `Lost` if the function `ComputeLPV1` is lost, and is erroneous (`Err`) otherwise.

```
let st = ComputeLPV1.status |
oDeviation1.status = {
    st = OK implies iSBAS1.status
    else st = Lost implies Lost
    else Err
}
```

The following code defines the status of the output port `oSelected1` of function `Acquire1`. If this function is OK, `oSelected1` the status is either equal to the status of its first input or to the status of its second input, depending on the selection made by the pilot. If the function `Acquire1` is lost, then the status of output `oSelected1` is lost. Otherwise, it is erroneous.

```
let st = Acquire1.status |
let v = iSelection1.value |
oSelected1.status = {
    st = OK and v = v0 implies iDeviation11.status
    else st = OK and v = v1 implies iDeviation21.status
    else st = Lost implies Lost
    else Err
}
```

Similarly, we define, for each function, output ports status (and value in the case of pilot selection, discrepancy, and display reset) from input ports status (and value).

Then, Alloy Analyzer proposes, with the command `run` to produce automatically an instance that satisfies all the definitions we expressed, and the global constraints. Of course, this is only possible if there is no inconsistency in the definitions. We can customize the way the instance is graphically displayed. Fig. 3 shows an excerpt of an instance produced by the Analyzer, where functions are represented by blue parallelograms, output ports by green circles and input ports by red circles.

4 Safety assessment

We are now able to check formally various properties expressed in Alloy. A first kind of properties we can check consists of structural properties about the model. For instance, the fact that a port can only belong to one function and the fact that flows correspond to one-to-many communication is expressed by the following properties:

```
assert model_structure {
    input in Function one → IPort
    output in Function one → OPort
    flow in OPort one → IPort }
```

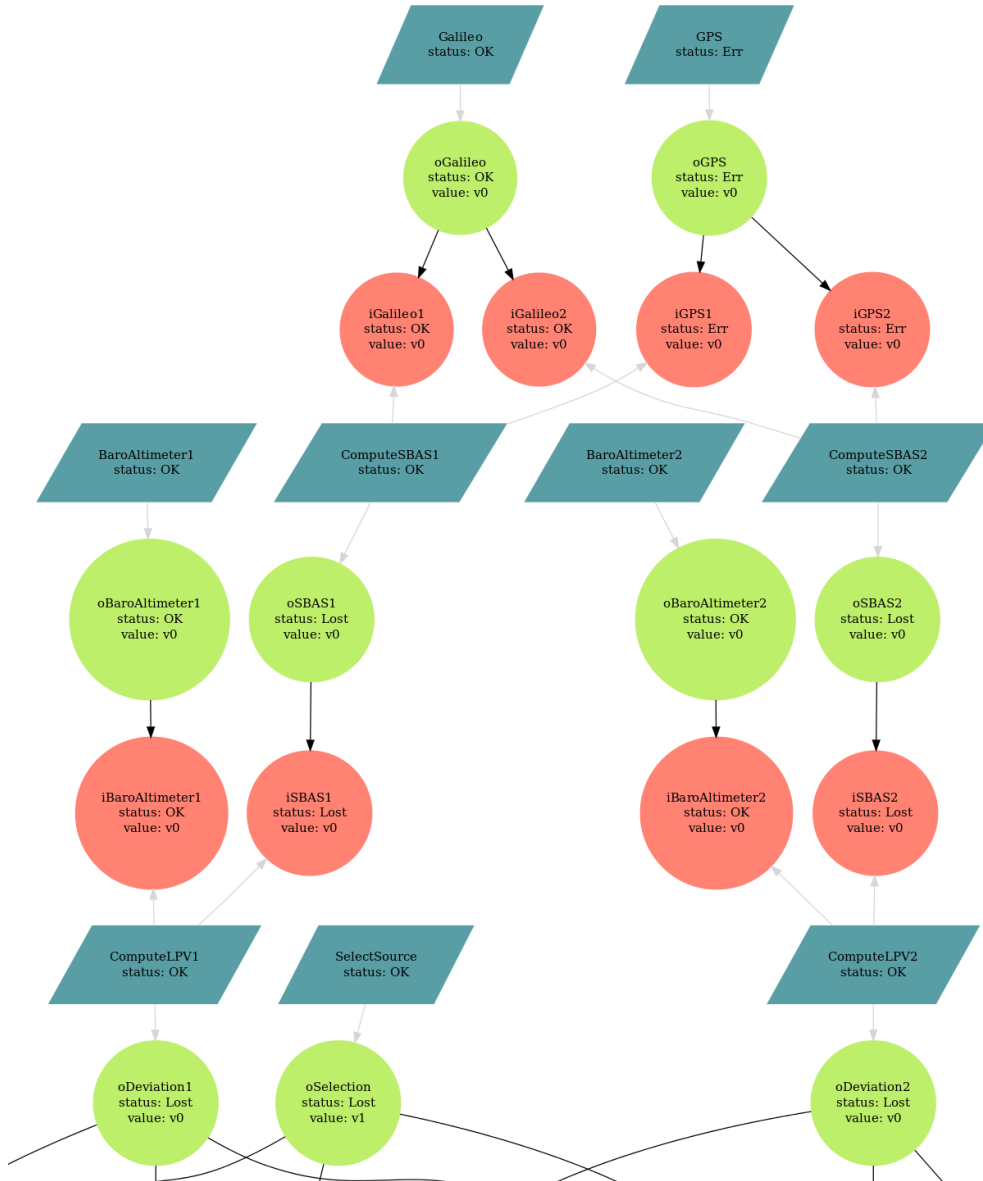


Figure 3: Visualization of LPV functional architecture in the Alloy Analyzer (excerpt)

The command `check model_structure` verifies that, *up to a certain bound*, all possible instances of the model satisfy these properties. If it is not the case, it yields a counter-example instance.

Now we want to validate the safety objectives expressed in Sect. 2. Concerning the *Loss of LPV capability* constraint, we express the following property:

- *If one (and only one) LPV processing is lost, and if the pilot makes a correct selection, then the data sent by the three displays are still correct.* This corresponds to the following Alloy code (for the loss of LPV1):

```
assert one_computer_lost {
  (all f: Function | (f ≠ ComputeLPV1 implies f.status=OK)
    and ComputeLPV1.status=Lost and oSelection.value=v1)
  implies oSelected1.status = OK
    and oSelected2.status = OK
    and oSelected3.status = OK
}
```

The command `check one_computer_lost` verifies that the model satisfies this property.

Concerning the *Misleading information integrity* constraint, we expressed two properties to be satisfied by the functional architecture:

- *If one LPV processing produces erroneous data, then an alarm (modeled by the variables `oDiscrepancy`) is launched on the three displays.* This corresponds to the following Alloy code (for the LPV1 in erroneous failure mode):

```
assert one_computer_erroneous {
  (all f: Function | (f ≠ ComputeLPV1 implies f.status=OK)
    and ComputeLPV1.status=Err)
  implies oDiscrepancy1.value=v1
    and oDiscrepancy2.value=v1
    and oDiscrepancy3.value=v1
}
```

- *If one display returns an erroneous data, it resets itself.* This corresponds to the following Alloy code for display 1 (`Acquire1`) in erroneous failure mode:

```
assert one_display_erroneous {
  (all f: Function | (f ≠ Acquire1 implies f.status=OK)
    and Acquire1.status=Err)
  implies oReset1.value=v1
}
```

These safety objectives are validated by the Alloy Analyzer.

5 Security assessment

Let us now consider the security objectives. For each attack identified in Sect. 2, we will express that the attack has no bad influence on the LPV capability, *i.e.*, the data produced by the three displays (represented by ports `oSelectedi`) are correct. So, we have the following expression for the first two identified attacks.

```

//Attack1
assert one_satellite_corrupted {
  (all f: Function | (f ≠GPS implies f.status=OK) and GPS.status=Err)
  implies oSelected1.status = OK and oSelected2.status = OK and oSelected3.status = OK
}
//Attack2
assert one_satellite_lost{
  (all f: Function | (f ≠GPS implies f.status=OK) and GPS.status=Lost)
  implies oSelected1.status = OK and oSelected2.status = OK and oSelected3.status = OK
}

```

It turns out that one of these properties is false with the current version of the model, which does not include RNAV function for lateral guidance. Alloy Analyzer produces counter examples that help understanding why the LPV capability is lost. Actually, without RNAV, the loss of SBAS satellite data is sufficient to lose the whole LPV capability. After the first attack, *i.e.*, a fake GPS signal, the SBAS does not manage to consolidate both satellite data, so the LPV processing does not have any input data, which induces a loss of the signal that is sent to the displays. However, the second attack, which consists in scrambling one satellite data (the corresponding signal is lost) is not sufficient to lose the LPV capability since SBAS sends to LPV processing the other satellite data.

In order to resist to the first attack, we propose to consider, in addition to satellite data, RNAV signal together with a Baro-altimeter, and to provide an alarm signal to the pilot when the difference between the information of SBAS and of the couple RNAV-Baro-altimeter is too important.

The definition of the output of LPV processings becomes more complicated, since it takes into account SBAS data, RNAV and baro-altimeter. Besides, it provides an alarm to the pilot in case of inconsistency between SBAS data on the one hand and RNAV/baro-altimeter data on the other hand. These definitions for the first LPV processing (the second is similar) are given as follows.

```

let st = ComputeLPV1.status | oDeviation1.status = {
  st=Lost implies Lost else st=Err implies Err else iSBAS1.status = OK implies OK
  else iSBAS1.status = Lost and iRNAV1.status=OK and iBaroAltimeter1.status=OK
  implies OK
  else iSBAS1.status=Lost and (iRNAV1.status=Lost or iBaroAltimeter1.status=Lost)
  implies Lost
  else iSBAS1.status=Lost and (iRNAV1.status=Err or iBaroAltimeter1.status=Err) implies Err
  else iSBAS1.status=Err and (iRNAV1.status=Lost or iBaroAltimeter1.status=Lost) implies Err
  else iSBAS1.status=Err and iRNAV1.status=OK and iBaroAltimeter1.status=OK implies Lost
  else iSBAS1.status=Err and (iRNAV1.status=Err or iBaroAltimeter1.status=Err) implies Lost
  else Err }

let st = ComputeLPV1.status | LPV1_alarm.value = {
  st=OK and iSBAS1.status ≠Lost and iRNAV1.status ≠Lost and iBaroAltimeter1.status ≠Lost
  and (iSBAS1.status=Err or iRNAV1.status=Err or iBaroAltimeter1.status=Err)
  implies v1
  else st=OK and iSBAS1.status= Err and (iRNAV1.status=OK or iBaroAltimeter1.status=OK)
  implies v1
  else st=OK and (iSBAS1.status = Lost)

```

```

    and (iRNAV1.status=Lost or iBaroAltimeter1.status=Lost) implies v1
  else v0 }

```

Now we can define the security objectives corresponding to all the attacks identified in Sect. 2.

```

//Attack 3
assert RNAV_lost {
  (all f: Function | (f ≠RNAV1 and f ≠RNAV2 implies f.status=OK)
    and RNAV1.status=Lost and RNAV2.status=Lost)
  implies oSelected1.status = OK and oSelected2.status = OK and oSelected3.status = OK
}

//Attack 4
assert one_satellite_lost_one_satellite_corrupted {
  (all f: Function | (f ≠GPS and f ≠Galileo) implies f.status=OK
    and GPS.status=Err and Galileo.status=Lost)
  implies LPV1_alarm.value=v1
}

//Attack 5
assert one_satellite_lost_RNAV_lost {
  (all f: Function | (f ≠GPS and f ≠RNAV1) implies f.status=OK)
    and GPS.status=Lost and RNAV1.status=Lost
  implies oSelected1.status = OK and oSelected2.status = OK and oSelected3.status = OK
}

//Attack 6
assert one_satellite_corrupted_RNAV_lost {
  (all f: Function | (f ≠GPS and f ≠RNAV1 and f ≠RNAV2) implies f.status=OK)
    and GPS.status=Err and RNAV1.status=Lost and RNAV2.status=Lost
  implies LPV1_alarm.value=v1
}

//Attack 7
assert one_satellite_corrupted_one_satellite_lost_RNAV_lost {
  (all f: Function | f ≠GPS and f ≠Galileo and f ≠RNAV1 and f ≠RNAV2
    implies f.status=OK)
    and GPS.status=Err and Galileo.status=Lost
    and RNAV1.status=Lost and RNAV2.status=Lost
  implies LPV1_alarm.value=v1
}

```

Notice that for properties relative to attacks 4, 6 and 7, we only require the alarm to be launched (the information that come to the displays are incorrect or absent in these cases). For the other properties, we require that the information displayed are correct.

All these properties are validated by the Alloy Analyzer.

6 Conclusion

In this article, we modeled the architecture of an avionic system in the Alloy language. We showed how Alloy allows to easily define a simple metamodel of avionic architectures with

failure propagation, which particularly fits the representation of the architecture of our case study.

We then expressed safety and security properties and enriched the initial model in order to fulfill all the objectives. We checked with the Alloy Analyzer that these properties are fulfilled by the model.

Since the safety and the security properties that we expressed are based on the same failure conditions, it would be possible to mix the assessment of both concerns, in order for instance to study the impact of attacks on safety objectives, or the other way around, the impact of unintentional failures on security objective. Such kinds of analyses of safety and security together, for which lightweight formal methods and Alloy in particular are promising, will be studied in particular in the remainder of the MERgE project.

References

- [1] André Arnold, Gérald Point, Alain Griffault & Antoine Rauzy (1999): *The AltaRica Formalism for Describing Concurrent Systems*. *Fundamenta Informaticae* 40(2,3), pp. 109–124, doi:10.3233/FI-1999-402302. Available at <http://iospress.metapress.com/content/W2506433232111V5>.
- [2] Julien Brunel, David Chemouil, Nicolas Mélédo & Vincent Ibanez (2014): *Formal Modelling and Safety Analysis of an Avionic Functional Architecture with Alloy*. In: *Embedded Real Time Software and Systems (ERTSS 2014)*, Toulouse, France.
- [3] Daniel Jackson (2006): *Software Abstractions: Logic, Language, and Analysis*. The MIT Press.
- [4] Amerson Lin, Mike Bond & Jolyon Clulow (2010): *Modeling Partial Attacks with Alloy*. In Bruce Christianson, Bruno Crispo, James A. Malcolm & Michael Roe, editors: *Security Protocols, Lecture Notes in Computer Science 5964*, Springer Berlin Heidelberg, pp. 20–33. Available at http://dx.doi.org/10.1007/978-3-642-17773-6_4.
- [5] Mark C. Reynolds (2010): *Lightweight Modeling of Java Virtual Machine Security Constraints*. In Marc Frappier, Uwe Glässer, Sarfraz Khurshid, Régine Laleau & Steve Reeves, editors: *Abstract State Machines, Alloy, B and Z, Lecture Notes in Computer Science 5977*, Springer Berlin Heidelberg, pp. 146–159. Available at http://dx.doi.org/10.1007/978-3-642-11811-1_12.
- [6] M. Toahchoodee & I. Ray (2009): *Using alloy to analyse a spatio-temporal access control model supporting delegation*. *Information Security, IET* 3(3), pp. 75–113, doi:10.1049/iet-ifs.2008.0074.