
AltaRica 3.0 in ten modelling patterns

Michel Batteux

IRT SystemX,
8 Avenue de la Vauve, 91120 Palaiseau,
Paris-Saclay, France
Email: michel.batteux@irt-systemx.fr

Tatiana Prosvirnova*

Laboratoire Genie Industriel,
CentraleSupélec,
Université Paris-Saclay,
8-10, rue Joliot-Curie,
91190 Gif-sur-Yvette, France
Email: tatiana.prosvirnova@centralesupelec.fr
*Corresponding author

Antoine B. Rauzy

MTP,
Norwegian University of Science and Technology,
S.P. Andersens veg 5,
7491 Trondheim, Norway
Email: antoine.rauzy@ntnu.no

Abstract: AltaRica 3.0 is an object-oriented modelling language dedicated to probabilistic risk and safety analyses. It is a prominent representative of modelling formalisms supporting the so-called model-based approach in reliability engineering. In this article, we illustrate the key features of the AltaRica 3.0 technology by presenting the implementation of ten very common modelling patterns. We demonstrate in this way the expressive power of the language as well as its elegance and simplicity of use.

Keywords: probabilistic risk and safety assessment; modelling languages; modelling patterns; AltaRica 3.0.

Reference to this paper should be made as follows: Batteux, M., Prosvirnova, T. and Rauzy, A.B. (2019) 'AltaRica 3.0 in ten modelling patterns', *Int. J. Critical Computer-Based Systems*, Vol. 9, Nos. 1/2, pp.133–165.

Biographical notes: Michel Batteux is currently a Research Engineer at the Technological Research Institute SystemX (France) and a Project Manager of the OpenAltaRica Project. After a Master degree in Mathematics and Computer Science at the Paris Diderot University, and a PhD in Computer Science at the Paris-Sud University, he had several research

engineer positions in different academic or industrial research laboratories: the Laboratory of Model driven engineering for embedded systems of CEA, the Computer Science Laboratory of the Ecole Polytechnique, Thales Research and Technology. His topics of interest are around the model-based design and assessment of complex technical systems, with a special focus on the performance assessment.

Tatiana Prosvirnova currently works as a Research Engineer at CentraleSupélec (Paris, France). She has a PhD in Computer Science. She is graduated from Ecole Polytechnique. She has MS of Science of Ecole Polytechnique in Systems Engineering. She has been working as a Software Developer at Dassault Systemes (the largest French software editor) for three years. Her research interests are model based safety assessment, systems engineering, formal methods, software development, system architecture modelling and model synchronisation.

Antoine B. Rauzy has currently a Full Professor position at Norwegian University of Science and Technology (NTNU, Trondheim, Norway). He is also the head of the chair Blériot-Fabre, sponsored by the group SAFRAN, at CentraleSupélec (Paris, France). During his career, he was a Researcher at French National Centre for Scientific Research, a CEO of the start-up company ARBoost Technologies and The director of the R&D Department on Systems Engineering at Dassault Systemes. He works in the reliability engineering and system safety field for more than 20 years. He extended his research topics to systems engineering since about ten years. He published over 150 articles in international conferences and journals. He developed state-of-the-art algorithms and software for probabilistic safety analyses.

1 Introduction

The aim of this article is to present the key features of the object-oriented, formal, modelling language AltaRica 3.0. The AltaRica 3.0 technology is dedicated to probabilistic risk and safety analyses of complex technical systems. We shall use here the expression ‘probabilistic risk analysis’ or equivalently ‘probabilistic risk assessment’ in a broad sense, which encompasses processes as diverse as safety and reliability assessments, optimisations of maintenance policies, assessments of the expected production level over a given period and so on. In a word, we shall speak here about assessments of operational performance of technical systems subject to random events such as mechanical failures, operator errors.

A long journey has been made since the publication of the WASH 1,400 report, (Rasmussen, 1975). Probabilistic risk and safety analyses are nowadays used on a daily basis in virtually all industries presenting significant risks for their operators, the public or the environment. Safety standards such as IEC 61508 (2010), ARP 4761 (2004) or ISO 26262 (2012) recommend this approach.

As of today, these analyses rely mainly on fault trees, event trees, reliability block diagrams or a combination of those, see, e.g., Kumamoto and Henley (1996), Andrews and Moss (2002) and Rausand and Høyland (2004) for reference books. These modelling formalisms present a good compromise between the ease of use, the

accuracy of descriptions and the computational complexity of assessments. However, they have important drawbacks as well. First, they lack expressive power. Cold and warm redundancies, exclusive failures, reconfigurations, control mechanisms and many other increasingly common features of technical systems can only be approximated, leading to over-pessimism and inaccurate ranking of causes of risks, (see, e.g., Epstein and Rauzy, 2005). Second, models designed with these formalisms are very distant from system specifications. Retrieving system specifications from the safety models is nearly impossible. For this reason, these models are hard to share with stakeholders and to maintain through the life-cycle of systems.

These drawbacks were compensated by advantages when systems at stake were purely mechanical, or at least when their mechanical parts were dominant. It is not the case anymore for software intensive, high integrity, critical systems. In these systems, the control (and reconfigurations) play a central role. We can characterise them as deformable: their architecture changes throughout their mission. New methods have to be developed to describe systems' behaviour more accurately.

To get more expressive power, one needs to leave combinatorial (Boolean) formalisms for states/events formalisms such as Markov chains or stochastic Petri nets (Ajmone-Marsan et al., 1994). Explicit representations of the state space, such as Markov chains, suffer from the exponential blow-up of the number of states and transitions. Implicit representations, such as stochastic Petri nets, are thus highly preferable. They are however not sufficient in themselves to reduce the distance between system specifications and safety models. With that respect, the lack of structure of stochastic Petri nets is an important drawback (see, e.g., Signoret et al., 2013).

The promise of the so-called model-based risk and safety assessment approach is to provide analysts with modelling formalisms that have both a high expressive power and suitable structuring mechanisms. It is possible in this way to design models that reflect the functional and physical architectures of the system under study. Safety models are thus closer to systems specifications, which makes them both easier to share with non-specialists and to maintain. Moreover, a single model can be used to assess several safety goals.

Modelling formalisms that support this approach can be classified into three categories. The first category consists of specialised profiles of model-based systems engineering formalisms such as SysML, (see, e.g., David et al., 2010; Yakymets et al., 2014; Mauborgne et al., 2016; Mhenni et al., 2016). The objective here is however more to introduce a safety facet into models of system architecture than to design actual safety models. The second category consists of extensions of fault trees or reliability block diagrams so to enrich their expressive power. This category includes dynamic fault trees (Dugan et al., 1992; Bouissou and Bon, 2003), multistate systems (Lisnianski and Levitin, 2003; Natvig, 2010; Papadopoulos et al., 2011), and some other proposals (Signoret et al., 2013). The third category, which aims at taking fully advantage of the model-based approach, consists of modelling languages such as SAML (Güdemann and Ortmeier, 2010), Figaro (Bouissou et al., 1991) and AltaRica (Point and Rauzy, 1999). SAML is oriented towards probabilistic model checking (and compiled into PRISM descriptions (Kwiatkowska et al., 2011)). Figaro has been historically the first modelling language dedicated to probabilistic risk and safety analyses. It is a rule-based systems (Bouissou et al., 2002; Bouissou and Houdebine, 2002). Since the very first version of AltaRica, the choice has been made to rely on the more natural and mathematically clearer notion of state automata.

AltaRica 3.0 is, as its name suggests, the third version of the language. The design of this new version started in 2012 to take advantage of more than 10 years of academic and industrial experience accumulated with the previous versions (Prosvirnova et al., 2013). AltaRica 3.0 can be described by the following equation.

$$\text{GTS} + \text{S2ML} = \text{AltaRica 3.0} \quad (1)$$

The above equation¹ summarises an idea that goes much beyond risk and safety analyses and that can be stated as follows. Any behavioural description language is made of two parts: a mathematical framework to describe behaviours and a set of constructs to structure models. In the case of AltaRica 3.0, the mathematical framework is the notion of guarded transition systems (GTS), see Rauzy (2008); Batteux et al. (2017) for in depth presentations. GTS are state automata. They have been designed to increase as much as possible the expressive power of the language without increasing the computational cost of assessment algorithms. S2ML stands for system structure modelling language (Batteux et al., 2015). S2ML gathers in a coherent way structuring constructs stemmed from object-oriented programming, (see, e.g., Abadi and Cardelli, 1998), and prototype-oriented programming, (see, e.g., Noble et al., 1999). The combination of GTS and S2ML results in a powerful, versatile language which exploits in an optimum way assessment algorithms.

It is of primary importance, in order to make the modelling process efficient, to reuse as much as possible modelling components within models and between models. In languages such as Modelica (Fritzson, 2015), this goal is achieved via the design of libraries of on-the-shelf ready-to-use modelling components. Reusing components is also possible in probabilistic risk and safety analyses, but to a much lesser extent. The reason is that these analyses represent systems at a high level of abstraction. Modelling components, except for very basic ones, tend thus to be specific to each system. In AltaRica 3.0, reuse is mostly achieved by the design of modelling patterns, i.e., examples of models representing remarkable features of the system under study. Once identified, patterns can be duplicated and adjusted for specific needs, see, e.g., Kehren et al. (2004) for a preliminary study and Kloul and Rauzy (2017) for a recent application. Patterns are pervasive in engineering. They have been developed for instance in the field of technical system architecture (Maier, 2009), as well as in software engineering (Gamma et al., 1994). Patterns are not only a mean to organise and to document models, but also and more fundamentally a way to reason about systems under study.

It is probably too early to design a taxonomy of modelling patterns encountered in operational performance analyses. For the time being, we classify patterns into two categories: behavioural patterns that aim at describing the behaviour of a single component or a small group of components, and architectural patterns that aim at describing the whole model organisation. We present and discuss here five patterns of each category.

The contribution of this article is thus twofold: first, it introduces the reader with the AltaRica 3.0 technology; second, it studies some of the most common modelling patterns encountered in probabilistic risk and safety analyses.

The remainder of this article is organised as follows. Section 2 introduces basic concepts of the language via the behavioural pattern ‘repairable unit’ and two fundamental architectural patterns: the ‘structural decomposition’ and ‘hierarchical

block diagram' patterns. Section 3 describes four common behavioural patterns: the 'periodically tested unit', 'warm redundancy', 'shared resource' and 'common cause failure' patterns. Section 4 describes three advanced architectural patterns: the 'reliability network', 'production tree', and 'monitored system' patterns. Finally, Section 5 gives a snapshot on assessment tools and concludes the article.

2 Getting started

2.1 Guarded transition systems

GTS, introduced in Rauzy (2008), are at the core of AltaRica 3.0. Formally, a guarded transition system is a quintuple $\langle V, E, T, A, \iota \rangle$, where:

- V is a set of variables. Each variable v of V has a type, i.e., can take its value in a certain set of constants (Booleans, integers, reals or set of symbolic constants), called its domain and denoted by $\text{dom}(v)$. V is actually the disjoint union of two subsets S and F , where S is the subset of state variables and F is the subset of flow variables.
- E is a set of events. Each event e of E can be associated with a non-decreasing invertible function delay_e from $[0, 1]$ to $\mathbb{R}^+ \cup \{+\infty\}$. The inverse delay_e^{-1} of this function is a cumulative probability distribution.
- T is a set of transitions. A transition t is a triple $\langle e, g, a \rangle$, denoted by $g \xrightarrow{e} a$, where e is an event of E , g is a Boolean condition on variables of V called the guard of the transition, and a is an instruction, called the action of the transition, that changes the value of (some of) the state variables.
- A is an instruction, called the assertion, that calculates the values of flow variables from the values of state variables.
- Finally, ι is a function that gives the initial value of state variables and the default value of flow variables.


As an illustration, we shall consider our first behavioural pattern.

Modeling Pattern 1 (repairable unit): We call repairable unit a component that may fail and be repaired.

The guarded transition system for such a unit is graphically represented Figure 1. The corresponding AltaRica 3.0 code is given Figure 2.

This modelling component is reused in many different models and often many times within a model. For this reason, it is declared as a class in the code of Figure 2 (lines 3–12). A class is an on-the-shelf modelling component. It can be instantiated as many times as necessary into models.

The state of a unit is represented by a variable named `_state` (declared line 4) that takes its value in the domain `RepairableUnitState` (declared line 1). This domain consists of the two symbolic constants `WORKING` and `FAILED`. Initially, the unit is working, so the attribute `init` of the variable `_state` is set to `WORKING`. This attribute indicates also that `_state` is a state variable.

Figure 1 Graphical representation of the guarded transition system for a repairable unit**Figure 2** AltaRica 3.0 code for the guarded transition system pictured Figure 1 (see online version for colours)

```

1  domain RepairableUnitState {WORKING, FAILED}
2
3  class RepairableUnit
4    RepairableUnitState _state (init = WORKING);
5    event failure (delay = exponential(lambda));
6    event repair (delay = Dirac(tau));
7    parameter Real lambda = 1.0e-4;
8    parameter Real tau = 12;
9    transition
10     failure: _state==WORKING -> _state := FAILED;
11     repair: _state==FAILED -> _state := WORKING;
12 end

```

As the unit is repairable, it has two transitions: a failure transition (declared line 10) that goes from the state working to the state failed and a repair transition (declared line 11) that goes the reverse way. In the code of Figure 2, events *failure* and *repair* (declared respectively lines 5 and 6) are associated with respectively a delay obeying the inverse of a negative exponential distribution of parameter *lambda* and constant delay *tau*. These parameters are declared respectively lines 7 and 8. AltaRica 3.0 provides several built-in distributions, like Dirac, exponential, Weibull as well as empirical distributions (given as a list of points).

AltaRica 3.0 comes with a standard library that declares a number of classes such as `RepairableUnit` to represent various types of components. These classes encode behavioural patterns. We shall see now examples of structural patterns.

2.2 Composition

The class `RepairableUnit` involves no flow variable and therefore no assertion. Flow variables are mainly a mean to connect components. The composition of two (or more) GTS is actually a guarded transition system.

Formally, let $M_1 : \langle V_1, E_1, T_1, A_1, \iota_1 \rangle$ and $M_2 : \langle V_2, E_2, T_2, A_2, \iota_2 \rangle$ be two GTS. Then $M_1 \otimes M_2$ is simply the guarded transition system $\langle V, E, T, A, \iota \rangle$ such that $V = V_1 \cup V_2$, $E = E_1 \cup E_2$, $T = T_1 \cup T_2$, $A = A_2 \circ A_1$ and $\iota = \iota_2 \circ \iota_1$.

This means that models can be obtained by composing smaller models.

The structural decomposition pattern is a typical example of composition of components. It is widely used to represent functional and physical breakdowns of systems (see, e.g., Krob, 2017). It works as follows.

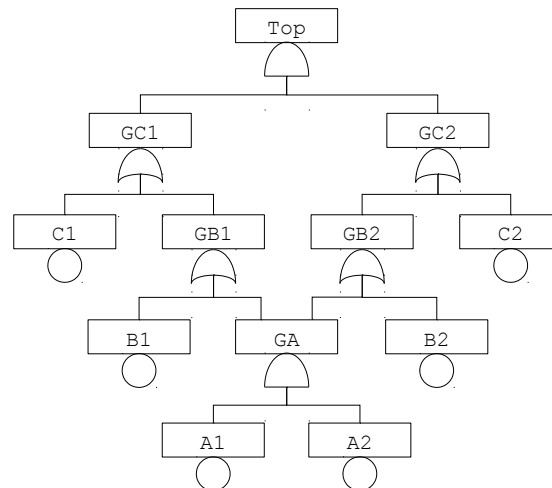
Modelling Pattern 2 (structural decomposition): A structural decomposition consists of:

- A hierarchy of components, each with its own behaviour represented by means of states and transitions. Each component in the hierarchy has a number of parents (possibly none) and a number of children (possibly none).
- A description of interactions between parents and children components, which is represented by flow variables and assertions.

Components are thus organised into a directed acyclic graph. Components are nodes of the graph. Links parent-child are edges of the graph. The graph is acyclic, therefore there is no loop in the hierarchy. If each component has at most one parent, then the hierarchy is a tree.

Fault trees are examples of models that obey the structural decomposition pattern. Consider for instance the fault tree pictured Figure 3.

Figure 3 A fault tree



This fault tree is made of 6 basic events: A1, A2, B1, B2, C1 and C2 and 6 internal events: top (the top event), GC1, GC2, GB1, GB2 and GA. Events are actually organised into a directed acyclic graph (and not just a tree because the event GA has two parents: GB1 and GB2).

Fault trees are a degenerated case of structural decomposition: first, only leaves of the hierarchy carry out behaviours; second, the communication goes only from children to parents. For this reason, it is not necessary to represent internal events explicitly as components, assertions suffice.

Assuming basic events of our tree represent failed states of repairable units (as described above), the AltaRica 3.0 code for this fault tree could be as given in Figure 4.

This code declares first the class `BasicEventForRepairableUnit` to encode basic events (lines 1–6). This class inherits from the class `RepairableUnit` (via the `extends` clause, line 2). It means that a `BasicEventForRepairableUnit` is a `RepairableUnit` with additional properties. In this case, a Boolean flow variable `failed` is declared (line 3). Its default value is set to `false` via the attribute `reset`. This attribute indicates

also that `failed` is a flow variable. Finally, the assertion line 5 tells how the value of the flow variable is calculated from the value of the state variable. The flow variable `failed` exports the state of the component and will be used to communicate this state to the parent components.

Figure 4 AltaRica 3.0 code for the fault tree pictured Figure 3 (see online version for colours)

```

1  class BasicEventForRepairableUnit
2      extends RepairableUnit;
3      Boolean failed(reset = false);
4      assertion
5          failed := _state==FAILED;
6  end
7
8  block FaultTree
9      BasicEventForRepairableUnit A1, A2(lambda=1.0e-5);
10     BasicEventForRepairableUnit B1, B2(lambda=2.0e-5);
11     BasicEventForRepairableUnit C1, C2(lambda=1.0e-7,
12         tau=24);
13     Boolean Top, GA, GB1, GB2, GC1, GC2 (reset = false);
14     assertion
15         Top := GC1 and GC2;
16         GC1 := C1.failed or GB1;
17         GC2 := C2.failed or GB2;
18         GB1 := B1.failed or GA;
19         GB2 := B2.failed or GA;
20         GA := A1.failed and A2.failed;
21     observer Boolean failed = Top;
22 end

```

The fault tree itself is encoded as prototype, i.e., a modelling component with a unique occurrence (lines 8–22). Prototypes are introduced by the keyword `block`.

The block `FaultTree` declares as many instances of `BasicEventForRepairableUnit` as there are basic events in the fault tree (lines 9–12) and as many Boolean flow variables as there are internal events (line 13). The assertion consists then simply in one equation per internal event, telling how the value of the internal event is calculated from the values of its children (line 15–20). Variables declared inside a component (prototype or instance of class) are accessed via the dot notation: `C1.failed` denotes the variable `failed` of the component `C1`.

In the code of Figure 4, the values of parameters `lambda` and `tau` of basic events are redefined at instantiation, when necessary. It is even possible to change the distributions themselves at instantiation. This makes it possible to design generic classes for components.

Note that it would be possible to use other types for basic events. We may for instance imagine that some of the components are not repairable. More complex schemes will be presented in the next section.

Note also that the fault tree has been declared as a prototype. There are actually little chance, if any, for this modelling component to be reused somewhere. It is definitely specific to the system under study, conversely to the components `RepairableUnit` and `BasicEventForRepairableUnit` that are reused several times in this and other

models. We shall see that some constructs are only available on prototypes, making the distinction between prototypes and classes of further interest.

The code declares the Boolean observer `failed` (line 21). Observers do not play any role in the behavioural description of the model. They are updated after each transition firing. They are used to define and to calculate performance indicators. Observers can be seen as the interface of the model for the assessment tools. They make it possible to optimise the code, for instance by removing variables, without disturbing the assessment of indicators.

2.3 Semantics

The semantics of a guarded transition system $M : \langle V = S \uplus F, E, T, A, \iota \rangle$ is defined as the set of its possible executions. To define formally the executions, we need to introduce the notions of state and schedule.

A state σ of M is valuation of variables of V verifying $\sigma(F) = A(\sigma(S))$.

A schedule of M is a function from T to $\mathbb{R}^+ \cup \{+\infty\}$. A schedule Γ is compatible with a state σ and a date $d \in \mathbb{R}^+$ if for all transitions $t : G \xrightarrow{e} P$ of T , $d \leq \Gamma(t)$ if $G(\sigma) = true$ and $\Gamma(t) = +\infty$ otherwise.

Intuitively, an execution of M is a sequence:

$$\langle \sigma_0, d_0, \Gamma_0 \rangle \xrightarrow{t_1} \langle \sigma_1, d_1, \Gamma_1 \rangle \xrightarrow{t_2} \dots \xrightarrow{t_n} \langle \sigma_n, d_n, \Gamma_n \rangle$$

where $n \geq 0$, the σ_i 's are states of M , the d_i 's are dates, i.e., non-negative real numbers verifying $0 = d_0 \leq d_1 \leq \dots \leq d_n$, each Γ_i is a schedule compatible with σ_i and d_i and finally the t_i 's are transitions of M .

Formally, the set of valid executions is defined recursively as follows.

The empty execution $\langle \sigma_0, 0, \Gamma_0 \rangle$ is a valid execution if $\sigma_0(S) = \iota$, $\sigma_0(F) = A(\iota)$ and the schedule Γ_0 is such that for all transitions $t : G \xrightarrow{e} P$ of T :

- $\Gamma_0(t) = delay_e(t)$ for some $z \in [0, 1]$ if $G(\iota) = true$.
- $\Gamma_0(t) = +\infty$ if $G(\iota) = false$.

Now, if $\Lambda = \langle \sigma_0, d_0, \Gamma_0 \rangle \xrightarrow{t_1} \dots \xrightarrow{t_n} \langle \sigma_n, d_n, \Gamma_n \rangle$, $n \geq 0$, is a valid execution, then so is the execution $\Lambda \xrightarrow{t_{n+1}} \langle \sigma_{n+1}, d_{n+1}, \Gamma_{n+1} \rangle$ if the following conditions hold, assuming $t_{n+1} = G_{n+1} \xrightarrow{e_{n+1}} P_{n+1}$.

- $G_{n+1}(\sigma_n) = true$.
- $\sigma_{n+1} = A(P_{n+1}(\sigma_n))$, i.e., the firing of the transition t_{n+1} is performed in two steps: first, state variables are updated by means of the action P_{n+1} of the transition, then flow variables are updated by means of the assertion A .
- $d_{n+1} = \Gamma_n(t_{n+1})$ and there is no transition t of T such that $\Gamma_n(t) < \Gamma_n(t_{n+1})$.
- Γ_{n+1} is obtained from Γ_n by applying the following rules to all transitions $t : G \xrightarrow{e} P$ of T .

1 If $G(\sigma_{n+1}) = true$, then:

- a If $G(\sigma_n) = true$ and $t \neq t_{n+1}$, i.e., if the transition was already scheduled, then $\Gamma_{n+1}(t) = \Gamma_n(t)$, i.e., the previous firing date is kept.


- b Otherwise, $\Gamma_{n+1}(t) = d_{n+1} + \text{delay}_e(z)$ for some $z \in [0, 1]$, i.e., a new firing date is chosen.
- 2 If $G(\sigma_{n+1}) = \text{false}$, then $\Gamma_{n+1}(t) = +\infty$.

Note that executions are fully determined by the choices of the z 's.

Example

In our example, there are two types of delays: exponential and Dirac. Figure 5 shows how their values are calculated from a number $z \in [0, 1]$.

Figure 5 Delays as inverse of cumulative probability distributions, (a) exponential delay (b) Dirac delay



A possible execution could be as follows.

At time 0, all state variables take the value WORKING, all the flow variables take the value false and the six transitions failure are fireable. The initial schedule could be for instance as follows.

A1.failure:	5849.45	B1.failure:	7068.84	C1.failure:	3629.63
A2.failure:	7406.81	B2.failure:	227.47	C2.failure:	1037.08

As B2.failure has the earliest firing date, it is fired (at 227.47). After its firing, B2.state takes the value FAILED, B2.failed, GB2 and GC2 take the value true, and the transition B2.repair gets fireable and is scheduled at $227.47 + 12 = 239.47$. The other variables and transitions stay unchanged.

A1.failure:	5849.45	B1.failure:	7068.84	C1.failure:	3629.63
A2.failure:	7406.81	B2.repair:	239.47	C2.failure:	1037.08

As B2.repair has the earliest firing date, it is fired (at 239.47). After its firing, all state variables take back the value FAILED, all flow variables take back the value false, and the transition B2.failure gets fireable again and is scheduled for instance at $239.47 + 2788.84 = 3128.21$. The other transitions stay unchanged.

A1.failure:	5849.45	B1.failure:	7068.84	C1.failure:	3629.63
A2.failure:	7406.81	B2.failure:	3128.21	C2.failure:	1037.08

And so on.

2.4 Another fundamental architectural pattern

Another very common and fundamental architectural pattern consists in representing a system as a hierarchy of communicating blocks:


Modelling Pattern 3 (hierarchical block diagram): A hierarchical block diagram consists of two types of blocks:

- Basic blocks that represent components of the system and that carry out the behaviour of these components.
- Internal blocks that may contain other internal and basic blocks.

Blocks may have ports that represent their interface. Ports are connected via links. Such hierarchical block diagrams are pervasive in model-based systems engineering. They are used in modelling languages as diverse as Matlab/Simulink (Klee and Allen, 2011), Modelica (Fritzson, 2015), Lustre (Halbwachs et al., 1991), SysML (Friedenthal et al., 2011), and many others.

As an illustration, consider the system pictured Figure 6.

Figure 6 A hierarchical block diagram



This system is made of the two subsystems AB and C in series. The subsystem AB is made of four basic units A1, A2, B1 and B2. The subsystem C is made of two basic units C1 and C2. Components are connected from left to right: the inputs of A1 and A2 are connected to two inputs. The inputs of B1 and B2 are connected to both outputs of A1 and A2. The inputs of C1 and C2 are connected respectively to the outputs of B1 and B2. Finally, the output of the system aggregates the outputs of C1 and C2.

If the basic units of this hierarchical block diagram are repairable units described by our modelling pattern 1, then the above description is just a (hierarchical) reliability block diagram.

The AltaRica 3.0 code for this diagram is given Figure 7.

Figure 7 AltaRica 3.0 code for hierarchical reliability block diagram pictured Figure 6 (see online version for colours)

```

1  class BasicBlockForRepairableUnit
2  extends RepairableUnit;
3  Boolean input, output (reset = false);
4  assertion
5  output := _state==WORKING and input;
6  end
7
8  block Plant
9  Boolean input1, input2 (reset = true);
10 block AB
11   BasicBlockForRepairableUnit A1, A2, B1, B2;
12   assertion
13   B1.input := A1.output or A2.output;
14   B2.input := A1.output or A2.output;
15 end
16 block C
17   BasicBlockForRepairableUnit C1, C2
18   (lambda = lambdaC);
19   parameter Real lambdaC = 2.0e-6;
20 end
21 Boolean output (reset = false);
22 assertion
23 input1 := true;
24 input2 := true;
25 AB.A1.input := input1;
26 AB.A2.input := input2;
27 C.C1.input := AB.B1.output;
28 C.C2.input := AB.B2.output;
29 output := C.C1.output or C.C2.output;
30 observer Boolean failed = not C.output;
31 end

```

This code starts by declaring a class `BasicBlockForRepairableUnit` for basic blocks. This class is similar to the class `BasicEventForRepairableUnit`, except it describes the transfer function between the input of the block and its output.

The system is declared as a block (lines 8–31). The block `Plant` declares two sub-blocks `AB` and `C` (respectively lines 10–15 and 16–20). These sub-blocks declare in turn instances of the class `BasicBlockForRepairableUnit`: `A1`, `A2`, `B1` and `B2` for the sub-block `AB` (line 11), `C1` and `C2` for the sub-block `C` (line 18). The declaration of `C1` and `C2` modifies the value of the parameter `lambda`. Assertions at system and sub-block levels realise the connections.

It would have been indeed possible to declare classes for sub-blocks `AB` and `C` and then to instantiate them into the block `Plant`. However, as these sub-blocks have a unique occurrence, it is more natural to declare them as prototypes. Moreover, prototypes make it possible to edit different levels of the hierarchy in the same view. For instance, it is possible to modify connections in the inner block `AB` while editing the outer block `Plant`. This is not possible with the class/instance mechanism because it is not allowed to modify a class from one of its instances.

2.5 Discussion

AltaRica 3.0 is agnostic: names are used to make clear what blocks, variables and events represent, but a model would be assessed in the same way if objects would be named X, Y, Z. Moreover, the naming conventions we used here are just a personal choice of the authors (e.g., we write symbolic constants with capital letters, prefix state variables with an underscore, capitalise names of blocks and classes).

AltaRica 3.0 is primarily a textual language. Graphical representations, such as those of Figures 1, 3 and 6 are an excellent communication means. They made the success of languages such as MATLAB/Simulink (Klee and Allen, 2011). We use them as much as possible. However, as soon as the model gets complex, they cannot embed all of its details. Moreover, the same model can be looked at from different angles, therefore with different graphical representations. In other words, the text is the reference, even though modelling environments make it possible to create models by dragging and dropping graphical representations (icons). The interested reader can look at Fuhrmann (2011) for an interesting discussion on the pragmatics of graphical modelling.

3 Behavioural patterns

In this section, we shall review some very common patterns used to represent the behaviours of components, beside the `RepairableUnit` pattern we have already seen.

3.1 Advanced failure models

The `RepairableUnit` pattern corresponds well to continuously monitored components. In process industry, there are however many components that are only periodically inspected or tested, e.g., shutdown valves that are tested by means of partial stroke. Safety standards and best practice guides ((IEC 61508, 2010; ISO/TR 12489, 2013) pay a lot of attention to this kind of components. Hence our next modelling pattern.

Modelling Pattern 4 (periodically tested component): A periodically tested component is a component:

- That alternates operation and test phases.
- Whose failures are only revealed thanks to the tests.

It is moreover assumed that the phases are organised on a calendar base, i.e., they have a fixed duration and their chaining is decided once for all. This (possible) approximation makes it possible to consider each component independently.

The behaviour of periodically tested components is conveniently described by multiphase Markov chains like the one pictured Figure 8 (indeed when occurrences of failures and repairs obey the Markovian hypothesis).

The component switches periodically from operation to test phases. The operation and test phases last respectively π and τ hours. A first operation phase that lasts θ is introduced so to be able to represent staggered tests when considering a multi-components system. The component fails with a failure rate λ . It is assumed that the component cannot fail during the test. The state of the component is thus represented

by means of three values (represented from top to bottom on the figure): its actual state, its observed state and its phase. Stochastic transitions are represented with plain arrows while deterministic transitions are represented with dashed arrows.

Figure 8 Multiphase Markov chain for a periodically tested component

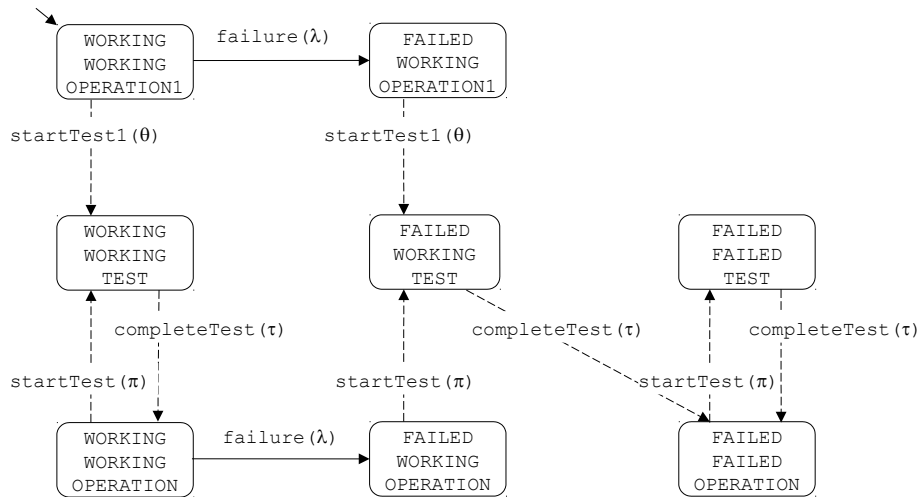


Figure 9 AltaRica 3.0 code for multiphase Markov chain pictured Figure 8 (see online version for colours)

```

1  domain State {WORKING, FAILED}
2  domain Phase {OPERATION1, OPERATION, TEST}
3
4  class PeriodicallyTestedComponent
5    State _actualState (init = WORKING);
6    State _observedState (init = WORKING);
7    Phase _phase (init = OPERATION1);
8    event failure (delay = exponential(lambda));
9    event startTest1 (delay = Dirac(theta));
10   event startTest (delay = Dirac(pi));
11   event completeTest (delay = Dirac(tau));
12   parameter Real lambda = 1.0e-6;
13   parameter Real theta = 2188;
14   parameter Real pi = 4378;
15   parameter Real tau = 2;
16   transition
17     failure: _actualState==WORKING and _phase!=TEST
18       -> _actualState := FAILED;
19     startTest1: _phase==OPERATION1 -> _phase := TEST;
20     startTest: _phase==OPERATION -> _phase := TEST;
21     completeTest: _phase==TEST -> {
22       _phase := OPERATION;
23       _observedState := _actualState;
24     }
25 end

```

The AltaRica 3.0 code for such a component is given Figure 9. This code is a direct translation of the multiphase Markov chain. It obeys the same principles as the one for repairable units, except that the state of the component is now described by means of three state variables: `_actualState`, `_observedState` and `_phase`.

The class `PeriodicallyTestedComponent` can be instantiated everywhere convenient (possibly after some adjustments), in particular in models built according to the structural decomposition pattern (pattern 2) and the hierarchical block diagram pattern (pattern 3).

More advanced models for periodically tested components are discussed in Cacheux et al. (2013). These models can be easily implemented in AltaRica 3.0 as well.

In the above pattern, the two phases alternate periodically after an initial phase. Many articles have been published that study phased mission systems, i.e., systems whose mission is decomposed into a finite number of successive phases, e.g. taxiing, take-off, cruise, landing and taxiing again for a commercial airliner. The first report work on this topic the authors could find was one by Esary and Ziehms (see Ziehms, 1974; Esary and Ziehms, 1975). In phased-mission systems, a component may be active only in some phases. Moreover, the causes of failure of the system may be different from one phase to the other. It is in general assumed that if the component is failed in one phase, then it remains failed in the subsequent phases. In any cases, it is fairly easy to adjust the above pattern to represent components of phased-mission systems.

3.2 Dependent behaviours

So far, the patterns we proposed make ‘only’ possible to represent combinatorial models such as fault trees or reliability block diagrams in a unified way and provide a mean to extend failure models for basic components. In a word, we stayed in the realm of combinatorial models.

Combinatorial models are not sufficient when dependencies amongst events have to be represented [although some approximations can be proposed, (see, e.g., Vaurio, 2001)]. Dependencies typically arise when a component is in cold or warm redundancy of another.

This leads us to our next pattern.

Modelling Pattern 5 (warm redundancy): A component is said in warm redundancy (of another one) if the following holds.

- The component is initially in standby mode. It is put in operation when there is a demand, i.e., typically when another component fails.
- The component may fail both in standby mode and in operation, although with different failure distributions, as a component in operation is more stressed than a component in standby.
- The attempt to start the component may be unsuccessful (failure on demand).
- The component may be repaired.
- The component is put back in standby mode after a repair or when it is not demanded anymore.

The above definition is quite general and can be indeed tuned to more specific cases, such as cold redundancies where the component is assumed to be safe when it is in standby mode.

Figure 10 shows the states/transitions diagram representing the behaviour of such component. Figure 11 gives the corresponding AltaRica 3.0 code.

Figure 10 States/transitions diagram for a component in warm redundancy

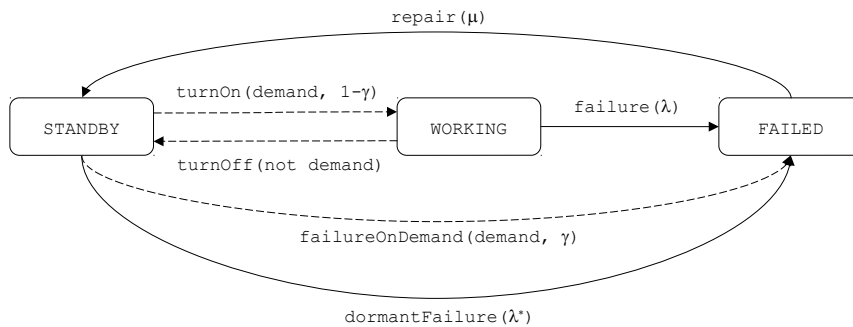


Figure 11 AltaRica 3.0 code for a component in warm redundancy (see online version for colours)

```

1  domain State {STANDBY, WORKING, FAILED}
2
3  class ComponentInWarmRedundancy
4    State _state (init = STANDBY);
5    Boolean demand (reset = false);
6    event turnOn (delay = Dirac(0),
7    expectation = 1-gamma);
8    event failureOnDemand (delay = Dirac(0),
9    expectation = gamma);
10   event turnOff (delay = Dirac(0));
11   event failure (delay = exponential(lambda));
12   event dormantFailure(delay = exponential(lambdaStar));
13   event repair (delay = exponential(mu));
14   parameter Real gamma = 0.02;
15   parameter Real lambda = 1.0e-4;
16   parameter Real lambdaStar = 1.0e-6;
17   parameter Real mu = 0.1;
18   transition
19     turnOn: _state==STANDBY and demand ->
20       _state:=WORKING;
21     failureOnDemand: _state==STANDBY and demand ->
22       _state:=FAILED;
23     turnOff: _state==WORKING and not demand ->
24       _state:=STANDBY;
25     failure: _state==WORKING -> _state := FAILED;
26     dormantFailure: _state==STANDBY ->
27       _state := FAILED;
28     repair: _state==FAILED -> _state := STANDBY;
29   end
  
```


This model involves a new type of stochastic transitions. When the component is in standby mode and gets demanded, technically when the input flow variable `demand` becomes true, the component reacts immediately: the two transitions `turnOn` and `failureOnDemand` (defined lines 20 and 20) become fireable. As they are labelled with events associated null (Dirac) delays (defined lines 7 and 9), one of them is instantaneously fired. The choice is non-deterministic: the attribute `expectation` can be seen as a weight. Assume transitions labelled with events e_1, \dots, e_n are in competition to be fired at the same exact date and that the `expectation`'s of these events are respectively w_1, \dots, w_n . Then, the transition labelled with the event e_i is fired with a probability $w_i / \sum_{j=1}^n w_j$. In our model, the weights of `turnOn` and `failureOnDemand` are determined by the parameter γ which can be seen as the probability of failure on demand.

Figure 12 AltaRica 3.0 code for trains in warm redundancy (see online version for colours)

```

1  class Train
2    ComponentInWarmRedundancy A, B;
3    Boolean demand, failed(reset = false);
4    assertion
5      A.demand := demand;
6      B.demand := demand;
7      failed := A._state==FAILED or B._state==FAILED;
8  end
9
10 block Plant
11   Train mainTrain, spareTrain;
12   Boolean failed(reset=false);
13   assertion
14     failed := mainTrain.failed and spareTrain.failed;
15     mainTrain.demand := not mainTrain.failed;
16     spareTrain.demand := mainTrain.failed;
17 end

```

Components in cold/warm redundancy are used in combination with other components, typically repairable components as defined by the pattern 1. They can be also used in combination to implement redundancies between subsystems, as illustrated by the code given Figure 12.

This code represents a system made of two trains: the main train and a spare train in warm redundancy. Each train is made of two components in series. The idea is then simple: the main train is demanded if it is not failed, while the spare train is demanded if the main train is failed. The same principle applies with any two subsystems.

Note that initially the two components of the main train are in standby. But, as the train is not failed, it is demanded. Therefore the two components are attempted to start.

3.3 Synchronised behaviours

In the above ‘warm redundancy’ pattern, when the main component fails, the spare component is immediately attempted to start. The failure of the former and start (or failure on demand) of the latter occur at the same date, although in order. There are cases however where two (or more) transitions occur exactly at the same time, without any ordering among them, typically because they encode local effects of a global event.

Figure 13 AltaRica 3.0 code for units sharing a maintenance team (see online version for colours)

```

1  domain UnitState {WORKING, FAILED, REPAIR}
2
3  class Unit
4      UnitState _state(init=WORKING);
5      event failure(delay = exponential(lambda));
6      event startRepair, completeRepair;
7      parameter Real lambda = 1.0e-4;
8      transition
9          failure: _state==WORKING -> _state:=FAILED;
10         startRepair: _state==FAILED -> _state:=REPAIR;
11         completeRepair: _state==REPAIR -> _state:=WORKING;
12 end
13
14 domain MaintenanceTeamState {STANDBY, WORKING}
15
16 class MaintenanceTeam
17     MaintenanceTeamState _state(init=STANDBY);
18     event startJob, completeJob;
19     transition
20         startJob: _state==STANDBY -> _state:=WORKING;
21         completeJob: _state==WORKING -> _state:=STANDBY;
22 end
23
24 block System
25     Unit U1, U2(startRepair.hidden = true,
26         completeRepair.hidden = true);
27     MaintenanceTeam M(startJob.hidden = true,
28         completeJob.hidden = true);
29     event startRepair1, startRepair2(delay = Dirac(0));
30     event completeRepair1,
31         completeRepair2(delay = exponential(mu));
32     parameter Real mu = 0.025;
33     transition
34         startRepair1:!U1.startRepair & !M.startJob;
35         startRepair2:!U2.startRepair & !M.startJob;
36         completeRepair1:!U1.completeRepair & !M.completeJob;
37         completeRepair2:!U2.completeRepair & !M.completeJob;
38 end

```

As an illustration, consider two units A and B sharing the same resource R. R can be for instance the repair team or a spare part. When the unit A gets the resource R, R becomes simultaneously unavailable for B. A and R change of state at exactly the same time. This idea is formalised by the following pattern.

Modelling Pattern 6 (shared resource): This pattern consists in two or more units sharing the same set of resources.

- When one of the units gets one of the resources, this resource becomes simultaneously unavailable for the other units.

- Symmetrically, when the unit releases the resource, the resource becomes simultaneously available for the other units.

AltaRica 3.0 provides the powerful concept of synchronisation to encode the ‘shared resource’ pattern. This concept has been originally introduced by Arnold and Nivat to describe interactions between concurrent processes (Arnold, 1994). It is generalised in AltaRica 3.0.

Consider, for example, a system consisting in two units and a maintenance team. Assume that the maintenance team can work on only one unit at a time. If a unit fails, it must therefore wait until that the maintenance team is free to start being repaired. Once repaired, it releases the resource, i.e., the maintenance team. A possible code for this system is given in Figure 13.

In this code, the definitions of classes `Unit` and `MaintenanceTeam` are no surprise. The block `System` declares events (lines 29–31). The corresponding transitions (lines 31–37) synchronise transitions of components.

For instance, the transition `startRepair1` results of the simultaneous firing of transitions `startRepair` of unit `U1` and `startJob` of the repair team. The modality `!` prefixing the event name makes the corresponding transition mandatory: the global transition `startRepair1` can be fired only if local transitions `U1.startRepair` and `M.startJob` are fireable. In other words, this transition is just equivalent to the following one.

```
startRepair1: U1._state==FAILED and M._state==STANDBY
-> { U1._state := REPAIR;
      M._state := WORKING;
    }
```

It is however much more convenient to declare local effect locally.

Instances of units and maintenance team are declared with the attribute `hidden` of their events `startRepair` and `completeRepair`, respectively `startJob` and `completeJob`, set to `true` (lines 26–28). This indicates that these events cannot be fired individually, but only through synchronisations.

We shall see now a more complex synchronisation scheme, involving the modality ?.

Common cause failures are an important contributor to the risk in many technical systems, e.g., in nuclear power plants, (see, e.g., Mosleh et al., 1998). There are many types of common cause failures. We shall restrict our attention to the following category (other types of common cause failures can be represented in AltaRica 3.0 as well, but with other modelling patterns).

Modelling Pattern 7 (common cause failures): A common cause failure is an event:

- Impacting simultaneously several basic components;
- Failing all impacted components that are not already failed.

Fire or flooding are typical such events. Strictly speaking, common cause failures as defined above cannot be described at basic component level since they involve several of them. Their description in AltaRica 3.0 relies on a very important construct of the language, namely the synchronisation of events.

The AltaRica 3.0 code that describes a common cause failure acting on three repairable units is given Figure 14.

Figure 14 AltaRica 3.0 code describing a common cause failure to three repairable units (see online version for colours)

```

1  block Plant
2    RepairableUnit A, B, C;
3    event ccf (delay = exponential(ccfRate));
4    parameter Real ccfRate = 1.0e-6;
5    transition
6      ccf: ?A.failure & ?B.failure & ?C.failure;
7  end

```

The event and the transition representing the common cause failure are declared in the block that aggregates the three components (lines 3 and 6). The transition line 6 synchronises the three events `A.failure`, `B.failure` and `C.failure`, i.e., attempts to fire these three events simultaneously. The modality `?` indicates that event it prefixes is fired only if possible (a modality `!` would indicate that event it prefixes is mandatory). The transition labelled with `ccf` is fireable if at least one the three events is fireable, i.e., if at least one of the guards of the individual transitions is satisfied in the current state. Its firing consists in performing actions of individual transitions that can be fired. Eventually, it is thus equivalent to the following code.

```

ccf: A._state==WORKING or B._state==WORKING or
C._state==WORKING -> {
  if A._state==WORKING then A._state := FAILED;
  if B._state==WORKING then B._state := FAILED;
  if C._state==WORKING then C._state := FAILED;
}

```

3.4 Discussion

The patterns presented in this section involve components that can be in more than two states, but that still fundamentally either working or failed. It is sometimes interesting to take into account degradation levels (and of course the transitions between these levels), as discussed in the abundant literature on so-called multistate systems, (see, e.g., Lisnianski and Levitin, 2003; Natvig, 2010). Tools like HiP-HOPS provide means to describe such systems (Papadopoulos et al., 2011). Rauzy and Yang proposed recently a unifying algebraic framework, (see Rauzy and Yang, 2018).

Multistate systems are straightforward to represent in AltaRica 3.0 (for this reason we do not introduce here a specific pattern). Figure 15 shows for instance how levels of degradation can be composed thanks to dedicated operators.

Once declared, operators (and functions) can be used in actions of transitions and assertions. Operators (and functions) in AltaRica 3.0 are similar to macro-instructions of programming languages. The switch expression is equivalent to a cascade of if-then-else expression: the conditions are looked at in turn and the first one which is satisfied determines the value of the expression.

Figure 15 AltaRica 3.0 code for aggregation operators working on the WDFState domain (see online version for colours)

```

1  domain WDFState {WORKING, DEGRADED, FAILED}
2
3  operator WDFState WDFMin(WDFState in1, WDFState in2)
4    switch {
5      case in1==FAILED or in2==FAILED: FAILED
6      case in1==WORKING and in2==WORKING: WORKING
7      default: DEGRADED
8    }
9  end
10
11 operator WDFState WDFMax(WDFState in1, WDFState in2,
12                          WDFState out)
13   switch {
14     case in1==FAILED and in2==FAILED: FAILED
15     case in1==WORKING or in2==WORKING: WORKING
16     default: DEGRADED
17   }
18 end

```

We presented here binary operators for a ternary logic for the sake of the simplicity. It is of course possible to extend them to any multivalued logic (and any number of arguments), (see, e.g., Malinowski, 2009) for a review.


4 Architectural patterns

In this section, we present three common architectural patterns to illustrate the importance of this concept and the ability of AltaRica 3.0 to implement advanced models.

4.1 Reliability networks

So far, all the models we considered are data-flow, i.e., that the information propagates only one way between components. There are however risk/safety assessment models in which the information can circulate in either directions, depending on the states of the components. Reliability networks belong to this category of models. Although much less popular than fault trees and reliability block diagrams, they have focused important research efforts, (see, e.g., Colbourn, 1987; Shier, 1991) for two monographs on this topics. The reason is that power and water distribution networks as well as several other types of infrastructures are typically analysed as reliability networks.

A reliability network is a graph with two kinds of nodes: source nodes that produce something (power, information, ...) and target nodes that consume and redistribute this something. Nodes of both types may fail according to some probability distribution. Edges are assumed to work perfectly.

Figure 16 A reliability network

As an illustration, consider the network pictured Figure 16. This network is made of the two source nodes $S1$ and $S2$ and six target nodes $T1, \dots, T6$.

In this network, edges are bidirectional. The information can thus circulate in different directions, depending on the state of components. For instance, the node $T6$ can be powered by the node $S1$ via the nodes $T2$ and $T4$. In case $T2$ and $S2$ are failed, $T4$ can be powered by the node $S1$ via the nodes $T1, T3, T5$ and finally $T6$.

Note that we do not lose any generality by assuming edges are perfect: an imperfect edge between two nodes A and B can be represented as an imperfect node N and two perfect edges form A to N and from N to B .

A typical question one may ask on such a networks is ‘what is the probability that a specific node is a powered at time t ’, i.e., what is the probability that there exists a working path from one of the operating source nodes to that particular node.

Answering this question raises difficult algorithmic issues. Specialised algorithms have been developed (Madre et al., 1994; Rauzy, 2003). As of today, AltaRica 3.0 is the only modelling language embedding natively mechanisms to represent and assess reliability networks, (see, Batteux et al., 2017) for in-depth explanations.

The solution is surprisingly simple and elegant. It relies on the following pattern.

Modelling Pattern 8 (reliability network): The reliability network pattern consists in a network of components such that:

- Each component has a unique Boolean input and a unique Boolean output.
- The input of a component is defined as a coherent Boolean formula of the outputs of the other nodes.

Whether each component is reachable, i.e., has at least one of its inputs true, depends on the global state of the network.

Recall that a Boolean formula is coherent if it is built only with ‘and’, ‘or’ and ‘k-out-of-n’ connectives.

Assuming that each node of the reliability network pictured Figure 16 is repairable, this network can be represented in AltaRica 3.0 using the class `BasicBlockForRepairableUnit` defined Section 2.4 and suitable assertions, as shown Figure 17.

The assertion of this model is ‘looped’, i.e., that variables depends eventually on each other. AltaRica 3.0 fixpoint mechanism to update the value of flow variables after each transition firing is able to deal efficiently with such dependencies (Batteux et al., 2017).

Figure 17 AltaRica 3.0 code for the reliability network of Figure 16 (see online version for colours)

```

1  block ReliabilityNetwork
2    BasicBlockForRepairableUnit S1, S2, T1, T2, T3, T4,
3    T5, T6;
4    assertion
5      S1.input := true;
6      S2.input := true;
7      T1.input := S1.output or T2.output or T3.output;
8      T2.input := S1.output or T1.output or T4.output;
9      T3.input := T1.output or T5.output;
10     T4.input := T2.output or T6.output;
11     T5.input := S2.output or T3.output or T6.output;
12     T6.input := S2.output or T4.output or T5.output;
13  end

```

4.2 Production systems

Many industrial processes can be represented as a hierarchical and oriented network of treatment units. Each unit treats the production of upstream units and transfers it to downstream units. The amount of products an unit is treating depends on its intrinsic capacity, on the production of upstream units and the demand of downstream units. There is thus a two-ways dependency between units.

When the organisation of production units is complex, optimising the production is also a complex problem, involving typically the determination of the maximum flow that can travel from source nodes to target nodes. Relatively efficient algorithms exist to do so, e.g., the Ford-Fulkerson algorithm (Cormen et al., 2001), but calling such an algorithm after each transition firing would be probably of a prohibitive cost. When the production plant can be decomposed hierarchically, the following pattern, that has been introduced and studied in Kloul and Rauzy (2017), makes it possible to solve the problem elegantly.

Modelling Pattern 9 (production tree): A production tree consists in a hierarchy (a tree) of components such that:

- Each component has a certain production capacity that depends on its intrinsic capacity, its state, and the capacities of its subcomponents.
- Each component has a certain production objective and is in charge of determining the production objectives of its subcomponents so to reach its own objective.
- The production objective of a component never exceeds its production capacity.

As an illustration, consider a sub-system S of a certain system consisting of two subsystems C and D. Assume that production policy consists in allocating the production objective in priority to C. Assume finally that S does not fail at the same rate when it is producing something and when it does not produce anything.

An AltaRica 3.0 code describing the system S is given in Figure 18.

Figure 18 AltaRica 3.0 code implementing a priority node in a production tree (see online version for colours)

```

1  block S
2    extends ComponentInWarmRedundancy;
3    block C
4      // description of C
5    end
6    block D
7      // description of D
8    end
9    parameter Real intrinsicCapacity = 110.0;
10   Real capacity(reset=0);
11   Real objective(reset=0);
12   assertion
13     capacity := switch {
14       case _state==FAILED: 0.0
15       case C.capacity+D.capacity>intrinsicCapacity:
16         intrinsicCapacity
17       default: C.capacity+D.capacity
18     };
19   if objective>C.capacity then {
20     C.objective := C.capacity;
21     D.objective := objective-C.capacity;
22   }
23   else {
24     C.objective := objective;
25     D.objective := 0;
26   }
27   demand := objective>0;
28 end

```

The block S inherits from the class `ComponentInWarmRedundancy` (pattern 5). It composes the two blocks C and D. Finally, it declares the parameter `intrinsicCapacity` and the two flow variables `capacity` and `objective` to represent the corresponding quantities.

The capacity of S should be 0 if S is failed, and the minimum of its intrinsic capacity and the sum of the capacities of C and D otherwise. This is directly translated into the assertion line 13.

S must split the production objective between C and D, giving the priority to C. This is again directly translated into the assertion line 19.

Finally, S must be started if its objective is not null and stopped otherwise, which is done by setting the inherited variable `demand` accordingly (line 27).

It is very easy to adjust the code of Figure 18 to represent hierarchical components that compose more than two components and/or that implement other production policies such as production with units in series, or units in parallel (and a pro-rata allocation of the production objective).

It is also worth to notice that by changing real flows to Boolean flows it is possible to use the production tree pattern to encode dynamic fault trees, (see, e.g., Dugan et al., 1992), as well as Boolean driven Markov process (Bouissou and Bon, 2003) applying the translation proposed in Rauzy (2015).

4.3 Monitored systems

Maintenance policies have a strong impact on operational performance of systems, for at least two reasons: first, a high reliability can only be achieved by means of maintenance interventions; second these interventions, that often require to shutdown at least partly the production, are in practice a major contributor to the system down times. One of the today's challenges of process industry is to move from calendar- or time-based maintenance policies to condition-based maintenance policies, in order to reduce down-times and beyond costs. If assessing the formers is well mastered, assessing the latter raises difficult modelling issues.

Our next and last pattern is dedicated to the analysis of monitored systems. It has been recently introduced and studied in Zhang et al. (2018).

Modelling Pattern 10 (monitored system): The monitored system pattern consists in organising the model into four modules:

- A module dedicated to the description of the behaviour of units of the system.
- A module that calculates the actual state of the system from the states of the units.
- A module that represents the diagnostic made on the state of the system.
- Finally, a controller that makes operational decisions, including the scheduling of maintenance interventions, based on the diagnostic and that broadcasts orders to units so that they implement these decisions.

This pattern not only helps to organise models of monitored systems, but more fundamentally provides a methodology to reason about these systems.

As an illustration, consider a two-out-of-three system made of periodically tested units. The system is thus failed if at least two out of the three units are failed. In this case, the production must be stopped immediately and a maintenance intervention scheduled. There may be a significant delay between the time a maintenance intervention is scheduled and the time it starts effectively. The production is stopped during maintenance intervention. If one unit is failed, the system is degraded. In this case, the production can go on, but a maintenance intervention must be scheduled. As a unit is out of order during a test, it is dangerous to test a unit while another one is failed because only the third unit remains in service. It is however necessary to take that risk, because it is the only mean to detect that a second unit may be failed, and therefore that the production must be stopped immediately. For the same reason, it is better to stagger tests of the units.

Modelling this kind of systems is by no means easy. One has to handle intricate stochastic and deterministic events, to represent both the actual state of the system and the diagnostic made on this state, to represent decision rules for maintenance interventions and to implement them effectively in units, and so on. Additional difficulties come when there are different types of maintenance interventions, applying on different groups of units.

The implementation in AltaRica 3.0 of the condition-diagnostic-decision pattern on our example is performed into two steps.

The first step consists in augmenting the periodically-inspected-component pattern so to represent maintenance interventions. To do so, it suffices to add

a phase MAINTENANCE to the component as well as two transitions: a transition `startMaintenance` from a phase other than MAINTENANCE to the phase MAINTENANCE and a transition `completeMaintenance` from the phase MAINTENANCE to the phase OPERATION1. The transition `completeMaintenance` may be slightly different depending on whether the component is as good as new after the maintenance or not.

The second step consists in describing the system as a whole, as shown Figure 19 and Figure 20.

Figure 19 AltaRica 3.0 code implementing the condition-diagnostic-decision pattern (see online version for colours)

```

1  domain Diagnostic {WORKING, DEGRADED, FAILED}
2
3  block Plant
4    block Units
5      PeriodicallyTestedUnit U1(tau=6, theta=2190-3*tau,
6        pi=4380-tau);
7      PeriodicallyTestedUnit U2(tau=6, theta=2190-2*tau,
8        pi=4380-tau);
9      PeriodicallyTestedUnit U3(tau=6, theta=2190-1*tau,
10       pi=4380-tau);
11   end
12   block ActualStateCalculator
13     embeds main.Units.U1._actualState as s1;
14     embeds main.Units.U2._actualState as s2;
15     embeds main.Units.U3._actualState as s3;
16     Boolean failed(reset=false);
17     assertion
18       failed := #(s1==FAILED, s2==FAILED, s3==FAILED)>=2;
19   end
20   block DiagnosticCalculator
21     embeds main.Units.U1._observedState as o1;
22     embeds main.Units.U2._observedState as o2;
23     embeds main.Units.U3._observedState as o3;
24     Diagnostic diagnostic(reset=WORKING);
25     assertion
26       diagnostic := switch {
27         case #(o1==FAILED, o2==FAILED, o3==FAILED)>=2:
28           FAILED
29         case #(o1==FAILED, o2==FAILED, o3==FAILED)==1:
30           DEGRADED
31         default: WORKING
32       };
33   end
34   block Controller
35     ...
36   end
37   observer Boolean dangerousState =
38     ActualStateCalculator.failed and
39     not Controller.productionStopped;
40   observer Boolean safeState =
41     Controller.productionStopped;
42   observer Boolean maintenanceOnGoing =
43     Controller.maintenanceOnGoing;
44 end

```

Figure 20 AltaRica 3.0 code implementing the controller (see online version for colours)

```

1  block Controller
2    embeds main.Units.U1 as U1;
3    embeds main.Units.U2 as U2;
4    embeds main.Units.U3 as U1;
5    embeds main.DiagnosticCalculator.diagnostic as
6      diagnostic;
7    Boolean productionStopped (init=false);
8    Boolean maintenanceOnGoing (init=false);
9    event stopProduction (delay=Dirac(0));
10   event startMaintenance (delay=Dirac(delta));
11   event completeMaintenance (delay=Dirac(mu));
12   parameter Real delta = 720;
13   parameter Real mu = 48;
14   transition
15     stopProduction:
16     diagnostic==FAILED and not productionStopped ->
17       productionStopped := true;
18     startMaintenance:
19     !U1.startMaintenance
20     & !U2.startMaintenance
21     & !U3.startMaintenance
22     & diagnostic!=WORKING and not maintenanceOnGoing ->
23       maintenanceOnGoing := true;
24     completeMaintenance:
25     !U1.completeMaintenance
26     & !U2.completeMaintenance
27     & !U3.completeMaintenance
28     & maintenanceOnGoing -> {
29       maintenanceOnGoing := false;
30       productionStopped := false;
31     }
32 end

```

The block Plant declares four sub-blocks corresponding to the four modules of the pattern. We have here the illustration of a very general principle: the ‘good’ architecture for a model does not necessarily mimics the architecture of the system, at least its physical architecture.

The block Units (line 4) gathers the declaration of individual units. Parameters τ (duration of tests) and π (interval between two tests) are redefined. Parameters θ ’s (date of the first test) are adjusted so to implement staggered tests.

The block ActualStateCalculator (line 12) calculates the actual state of the plant from the `_actualState`’s of the units. To do so, it aggregates these variables by means of the clause `embeds`. Aggregation should be seen as a reference to an element declared outside the aggregating block. This element is accessed by means of an absolute or a relative path. In the code of the example, it is an absolute path: the keyword `main` denotes the outer most block (or class) of the current hierarchy. Therefore, `main.Units.U1._actualState` denotes the variable `_actualState` of the unit `U1` of the block `Units` of the block `Plant`. It would have been possible to use instead the relative path `owner.Units.U1._actualState`. The keyword `owner` denotes the parent block.

The block `DiagnosticCalculator` (line 20) makes a diagnostic on the state of the plant from the `_observedState`'s of the units. It aggregates these variables.

The block `Controller` stops the production when the plant is diagnosed failed (line 15). It schedules and executes maintenance interventions. To do so, it synchronises its own events `startMaintenance` and `completeMaintenance` with those of units (lines 18 and 24).

Observers `dangerousState`, `safeState` and `maintenanceOnGoing` make it possible to count the number of times one of the states they describe has been encountered, to get the sojourn times in these states and other indicators of interest.

This pattern concludes our presentation of AltaRica 3.0.

5 Conclusions

In this article, we gave a snapshot of the expressive power of the AltaRica 3.0 modelling language by showing how to encode ten very common modelling patterns in probabilistic risk and safety analyses of complex technical systems. Modelling patterns are not only a very efficient means to architect and to document models, they are more fundamentally a way to reason about systems under study. We distinguished here between behavioural patterns, that stand at component level, and architectural patterns, that stand at model level. This taxonomy may be imperfect and will probably evolve in the future. Behavioural patterns are rather well mastered as they are used in a variety of contexts and with other modelling formalisms than AltaRica. On the contrary, the study of architectural patterns is still in its infancy (or its adolescence). The authors are deeply convinced that, together with the introduction of machine learning techniques to obtain degradation profile of components, they are a game changer in probabilistic risk and safety analyses.

An integrated modelling environment for AltaRica Wizard (AltaRica 3.0) is currently under development as joint effort of the OpenAltaRica team at IRT-SystemX (Paris, France) and the Norwegian University of Science and Technology. Industrial partners (Airbus, Safran and Thales) support this project. A versatile set of assessment tools has been developed, which includes:

- A step by step simulator making possible to play ‘what-if’ scenarios and to validate models. This simulator implements abstract interpretation techniques so to simulate faithfully stochastic and timed executions (Batteux et al., 2018).
- A compiler of AltaRica models into fault trees. This compiler relies on advanced algorithmic techniques (Prosvirnova and Rauzy, 2015). Fault trees are then assessed with XFTA (Rauzy, 2012), which is one of the most efficient available calculation engines.
- A compiler of AltaRica models into Markov chains. This compiler produces Markov chains that approximate the original model while staying of reasonable sizes (Brameret et al., 2015). Markov chains are then assessed with Mark-XPR, as very efficient calculation engine (Rauzy, 2004).
- A generator of critical sequences (still under development at the time we write these lines).

- A stochastic simulator. Stochastic simulation is itself a versatile tool to assess complex models, (see, e.g., Zio, 2013).

These tools make the AltaRica 3.0 technology versatile and efficient. They make it possible cross-verification. They prefigure what will be the next generation of modelling environments for the assessment of operational performance of complex technical systems.

Tutorial material, courses for both primary and continuing education are also under development.

In a word, the AltaRica 3.0 technology is now mature or close to be. Deploying it in industry will indeed take time, but everything is ready for. With that respect, training analysts is of course a key issue. New technologies are often deployed by new generation of engineers. This will be probably the case for the model-based approach in risk and safety analyses. Using high-level modelling languages such as AltaRica 3.0 requires some familiarity with information technology in general and programming in particular. New generations of engineers have undoubtedly these skills.

Safety analyses are often thought as a cost. They are performed because regulation bodies require them. As we are moving from a product-centered industry to a service-oriented industry, their role is however changing: they are used not only to assess whether systems are safe and available enough to be operated, but also to establish contractual relations between providers and clients of services. In other words, they are becoming an economical driver. With that respect, we can expect that the industrial impact of new technologies such as AltaRica 3.0 goes much beyond traditional safety analyses.

Acknowledgements

The authors would like to thank the reviewers of this article for their nice and useful comments.

References

- Abadi, M. and Cardelli, L. (1998) *A Theory of Objects*, Springer-Verlag, New York, USA.
- Ajmone-Marsan, M., Balbo, G., Conte, G., Donatelli, S. and Franceschinis, G. (1994) *Modelling with Generalized Stochastic Petri Nets*, Wiley Series in Parallel Computing, John Wiley and Sons, New York, USA.
- Andrews, J.D. and Moss, R.T. (2002) *Reliability and Risk Assessment*, 2nd ed., ASM International, Materials Park, Ohio 44073-0002, USA.
- Arnold, A. (1994) *Finite Transition Systems*, C.A.R Hoare. Prentice Hall, Upper Saddle River, New Jersey, USA.
- ARP 4761 (2004) *Guidelines and Methods for Conducting the Safety Assessment Process on Civil Airborne Systems and Equipment*, SAE International.
- Batteux, M., Prosvirnova, T. and Rauzy, A. (2015) *System Structure Modeling Language (S2ML)*, AltaRica Association, archive hal-01234903, version 1.

- Batteux, M., Prosvirnova, T. and Rauzy, A. (2017) 'AltaRica 3.0 assertions: the whys and the wherefore', *Journal of Risk and Reliability*, September, Vol. 231, No. 6, pp.691–700.
- Batteux, M., Prosvirnova, T. and Rauzy, A. (2018) 'Enhancement of the AltaRica 3.0 stepwise simulator by introducing an abstract notion of time', in *Proceedings of European Safety and Reliability Conference Safe Societies in a Changing World (ESREL 2018)*, Trondheim, Norway, June.
- Bouissou, M. and Bon, J-L. (2003) 'A new formalism that combines advantages of fault-trees and Markov models: Boolean logic-driven Markov processes', *Reliability Engineering and System Safety*, Vol. 82, No. 2, pp.149–163.
- Bouissou, M. and Houdebine, J-C. (2002) 'Inconsistency detection in KB3 models', in *Proceedings of European Safety and Reliability Conference, ESREL '2002*, ISdF, Lyon, France, March, pp.754–759.
- Bouissou, M., Bouhadana, H., Bannelier, M. and Villatte, N. (1991) 'Knowledge modelling and reliability processing: presentation of the FIGARO language and of associated tools', in Lindeberg, J.F. (Ed.): *Proceedings of SAFECOMP '91 – IFAC International Conference on Safety of Computer Control Systems*, Pergamon Press, Trondheim, Norway, pp.69–75.
- Bouissou, M., Humbert, S., Muffat, S. and Villatte, N. (2002) 'KB3 tool: feedback on knowledge bases', in *Proceedings of European Safety and Reliability Conference, ESREL '2002*, Lyon, France, ISdF, March, pp.114–119.
- Brameret, P-A., Rauzy, A. and Roussel, J-M. (2015) 'Automated generation of partial Markov chain from high level descriptions', *Reliability Engineering and System Safety*, July, Vol. 139, pp.179–187.
- Cacheux, P-J., Collas, S., Dutuit, Y., Folleau, C., Signoret, J-P. and Thomas, P. (2013) 'Assessment of the expected number and frequency of failures of periodically tested systems', *Reliability Engineering and System Safety*, October, Vol. 118, No. C, pp.61–70.
- Colbourn, C.J. (1987) *The Combinatorics of Network Reliability*, Oxford University Press, New York.
- Cormen, T.H., Leiserson, C.E., Rivest, R.L. and Stein, C. (2001) *Introduction to Algorithms*, 2nd ed., The MIT Press, Cambridge, MA, USA.
- David, P., Idasiak, V. and Kratz, F. (2010) 'Reliability study of complex physical systems using sysml.', *Reliability Engineering and System Safety*, Vol. 95, No. 4, pp.431–450.
- Dugan, J.B., Bavuso, S.J. and Boyd, M.A. (1992) 'Dynamic fault-tree models for fault-tolerant computer systems', *IEEE Transactions on Reliability*, September, Vol. 41, No. 3, pp.363–377.
- Epstein, S. and Rauzy, A. (2005) 'Can we trust PRA?', *Reliability Engineering and System Safety*, Vol. 88, No. 3, pp.195–205.
- Esary, J.D. and Ziehms, H. (1975) 'Reliability analysis of phased missions', in Barlow, R.E., Fussel, J.B. and Singpurwalla, N.D. (Eds.): *Reliability and Fault Tree Analysis: Theoretical and Applied Aspects of System Reliability and Safety Assessment*, pp.213–236, SIAM Press, Philadelphia, PA, USA.
- Friedenthal, S., Moore, A. and Steiner, R. (2011) *A Practical Guide to SysML: the Systems Modeling Language*, The MK/OMG Press, Morgan Kaufmann, San Francisco, CA 94104, USA.
- Fritzson, P. (2015) *Principles of Object-Oriented Modeling and Simulation with Modelica 3.3: a Cyber-Physical Approach*, Wiley-IEEE Press, Hoboken, New Jersey 07030-5774, USA.
- Fuhrmann, H.A.L. (2011) *On the Pragmatics of Graphical Modeling*, Kiel Computer Science Series, Book on Demand, Norderstedt, Germany.

- Güdemann, M. and Ortmeier, F. (2010) 'A framework for qualitative and quantitative model-based safety analysis', in *Proceedings of the IEEE 12th High Assurance System Engineering Symposium (HASE 2010)*, IEEE, San Jose, CA, USA, pp.132–141.
- Gamma, E., Helm, R., Johnson, R. and Vlissides, J. (1994) *Design Patterns – Elements of Reusable Object-Oriented Software*, Addison-Wesley Professional Computing Series, Addison-Wesley, Boston, MA 02116, USA, October.
- Halbwachs, N., Caspi, P., Raymond, P. and Pilaud, D. (1991) 'The synchronous dataflow programming language lustre', *Proceedings of the IEEE*, Vol. 79, No. 9, pp.1305–1320.
- International iec standard iec61508 (2010) *International iec standard iec61508 – functional safety of electrical/electronic/programmable safety-related systems (e/e/pe, or e/e/pes)*, Standard, International Electrotechnical Commission, Geneva, Switzerland, April.
- ISO 26262 (2012) *ISO 26262:2011(en): Road Vehicles – Functional Safety*, International Standardization Organization, Geneva, Switzerland.
- ISO/TR 12489 (2013) *ISO/TR 12489:2013 Petroleum, Petrochemical and Natural Gas Industries – Reliability Modelling and Calculation of Safety Systems*, November, International Organization for Standardization, Geneva, Switzerland.
- Kehren, C., Seguin, C., Bieber, P., Castel, C., Bougnol, C., Heckmann, J-P. and Metge, S. (2004) 'Architecture patterns for safe design', in *Electronic Proceedings of 1st Complex and Safe Systems Engineering Conference (CS2E 2004)*, AAAF, Arcachon, France.
- Klee, H. and Allen, R. (2011) *Simulation of Dynamic Systems with MATLAB and Simulink*, CRC Press, Boca Raton, FL 33431, USA, February.
- Kloul, L. and Rauzy, A. (2017) 'Production trees: a new modeling methodology for production availability analyses', *Reliability Engineering and System Safety*, November, Vol. 167, pp.561–571.
- Krob, D. (2017) 'CESAM: CESAMES systems architecting method: a pocket guide', *CESAMES* [online]<http://www.cesames.net> (accessed January 2017).
- Kumamoto, H. and Henley, E.J. (1996) *Probabilistic Risk Assessment and Management for Engineers and Scientists*, IEEE Press, Piscataway, New Jersey, USA.
- Kwiatkowska, M., Norman, G. and Parker, D. (2011) 'Prism 4.0: verification of probabilistic real-time systems', in *Proceedings 23rd International Conference on Computer Aided Verification (CAV '11)*, Vol. 6806 of LNCS, Springer, New York, USA, pp.585–591.
- Lisnianski, A. and Levitin, G. (2003) *Multi-State System Reliability, Quality, Reliability and Engineering Statistics*, World Scientific, London, England.
- Madre, J-C., Coudert, O., Fraïssé, H. and Bouissou, M. (1994) 'Application of a new logically complete ATMS to digraph and network-connectivity analysis', in *Proceedings of the Annual Reliability and Maintainability Symposium, ARMS '94*, IEEE, Anaheim, California, pp.118–123.
- Maier, M.W. (2009) *The Art of Systems Architecting*, CRC Press, Boca Raton, FL 33431, USA, ISBN: 978-1420079135.
- Malinowski, G. (2009) 'Many-valued logic and its philosophy', in Gabbay, D.M. and Woods, J. (Eds.): *Handbook of the History of Logic, The Many Valued and Non-monotonic Turn in Logic*, Vol. 8, pp.13–94, Elsevier, Oxford, UK.
- Mauborgne, P., Deniaud, S., Levrat, E., Bonjour, E., Micaëlli, J-P. and Loise, D. (2016) 'Operational and system hazard analysis in a safe systems requirement engineering process – application to automotive industry', *Safety Science*, August, Vol. 87, pp.256–268.
- Mhenni, F., Choley, J-Y., Nguyen, N. and Frazza, C. (2016) 'Flight control system modeling with sysml to support validation, qualification and certification', *IFAC-PapersOnLine*, Vol. 49, No. 3, pp.453–458.

- Mosleh, A., Rasmuson, D.M. and Marshall, F.M. (1998) *Guidelines on Modeling Common-cause Failures in pra*, Technical Report NUREG/CR-5485, US Nuclear Regulatory Commission.
- Natvig, B. (2010) *Multistate Systems Reliability Theory with Applications*, Wiley, Hoboken, New Jersey, USA.
- Noble, J., Taivalaari, A. and Moore, I. (1999) *Prototype-Based Programming: Concepts, Languages and Applications*, Springer-Verlag, Berlin and Heidelberg, Germany.
- Papadopoulos, Y., Walker, M., Parker, D., Rüde, E., Hamann, R., Uhlig, A., Grätz, U. and Liend, R. (2011) 'An approach to optimization of fault tolerant architectures using hip-hops', *Journal of Engineering Failure Analysis*, March, Vol. 18, No. 2, pp.590–608.
- Point, G. and Rauzy, A. (1999) 'AltaRica: constraint automata as a description language', *Journal Européen des Systèmes Automatisés*, Vol. 33, Nos. 8–9, pp.1033–1052.
- Prosvirnova, T. and Rauzy, A. (2015) 'Automated generation of minimal cutsets from AltaRica 3.0 models', *International Journal of Critical Computer-Based Systems*, Vol. 6, No. 1, pp.50–79.
- Prosvirnova, T., Batteux, M., Brameret, P.-A., Cherfi, A., Friedlhuber, T., Roussel, J.-M. and Rauzy, A. (2013) 'The AltaRica 3.0 project for model-based safety assessment', in *Proceedings of 4th IFAC Workshop on Dependable Control of Discrete Systems, DCDS '2013*, International Federation of Automatic Control, York, Great Britain, September, pp.127–132.
- Rasmussen, N.C. (1975) *Reactor Safety Study, an Assessment of Accident Risks in US Commercial Nuclear Power Plants*, Technical Report WASH 1400, NUREG-75/014, US Nuclear Regulatory Commission, Rockville, MD, USA, October.
- Rausand, M. and Høyland, A. (2004) *System Reliability Theory: Models, Statistical Methods, and Applications*, 2nd ed., Wiley-Blackwell, Hoboken, New Jersey, USA, January 2004.
- Rauzy, A. and Yang, L. (2018) 'Finite degradation structures', *Journal of Applied Logic*, article submitted.
- Rauzy, A. (2003) 'A new methodology to handle Boolean models with loops', *IEEE Transactions on Reliability*, Vol. 52, No. 1, pp.96–105.
- Rauzy, A. (2004) 'An experimental study on six algorithms to compute transient solutions of large Markov systems', *Reliability Engineering and System Safety*, October, Vol. 86, No. 1, pp.105–115.
- Rauzy, A. (2008) 'Guarded transition systems: a new states/events formalism for reliability studies', *Journal of Risk and Reliability*, Vol. 222, No. 4, pp.495–505.
- Rauzy, A. (2012) 'Anatomy of an efficient fault tree assessment engine', in Virolainen, R. (Ed.): *Proceedings of International Joint Conference PSAM '11/ESREL '12*, Helsinki, Finland, June.
- Rauzy, A. (2015) 'Towards a sound semantics for dynamic fault trees', *Reliability Engineering and System Safety*, October, Vol. 142, pp.184–191.
- Shier, D.R. (1991) *Network Reliability and Algebraic Structures*, Oxford Science Publications, Oxford, England.
- Signoret, J.-P., Dutuit, Y., Cacheux, J.-P., Folleau, C., Collas, S. and Thomas, P. (2013) 'Make your petri nets understandable: Reliability block diagrams driven Petri nets', *Reliability Engineering and System Safety*, Vol. 113, pp.61–75.
- Ziehms, H. (1974) *Reliability Analysis of Phased Missions*, PhD thesis, Naval Postgraduate School.
- Vaurio, J.K. (2001) 'Making systems with mutually exclusive events analyzable by standard fault tree analysis tools', *Reliability Engineering and System Safety*, Vol. 74, No. 1, pp.75–80.

- Yakymets, N., Julho, Y.M. and Lanusse, A. (2014) ‘Sophia framework for model-based safety analysis’, in *Actes du congrès Lambda-Mu 19 (actes électroniques)*, Institut pour la Maîtrise des Risques, Dijon, France, October.
- Zhang, Y., Barros, A. and Rauzy, A. (2018) ‘Finite degradation structures’, *Journal of Risk and Reliability*, article submitted.
- Zio, E. (2013) *The Monte Carlo Simulation Method for System Reliability and Risk Analysis*, Springer Series in Reliability Engineering, Springer, London, England.

Note

- 1 This equation echoes the title of the famous book *Data Structures + Algorithms = Programs*.