

From LOTOS to LNT

Hubert Garavel Frédéric Lang Wendelin Serwe

INRIA, Grenoble, France
Univ. Grenoble Alpes, LIG, F-38000 Grenoble, France
CNRS, LIG, F-38000 Grenoble, France
E-mail: {hubert.garavel, frederic.lang, wendelin.serwe}@inria.fr
<http://convecs.inria.fr>

Abstract

We revisit the early publications of Ed Brinksma devoted, on the one hand, to the definition of the formal description technique LOTOS (ISO International Standard 8807:1989) for specifying communication protocols and distributed systems, and, on the other hand, to two proposals (Extended LOTOS and Modular LOTOS) for making LOTOS a simpler and more expressive language. We examine how this scientific agenda has been dealt with during the last decades. We review the successive enhancements of LOTOS that led to the definition of three languages: E-LOTOS (ISO International Standard 15437:2001), then LOTOS NT, and finally LNT. We present the software implementations (compilers and translators) developed for these new languages and report about their use in various application domains.

Keywords: abstract data type, algebraic specification, concurrency theory, E-LOTOS, formal description technique, formal method, formal specification, LOTOS, LNT, process algebra, process calculus, specification language.

1 Introduction

The present article was written in honor of Ed Brinksma and included in a collective *Festschrift* book offered to him at the occasion of his 60th birthday.

The first part of Ed Brinksma's research career has been devoted to the design of formal methods for the specification of communication protocols and distributed systems, the LOTOS language being the common theme and vital lead for the scientific contributions. This first part approximately extends over twelve years, between 1984 (as dated by the conference article [14]) and 1995 (as dated by the book chapter [9]). It was directly succeeded, with some chronological overlap, by a second part centered on conformance testing for protocols (with a first paper [16] published in 1991) and a third part centered on real time and performance evaluation (with early papers, e.g., [15] published in 1995).

In the present article, we focus on this first part, in which we distinguish two different threads of work: (i) the definition of the LOTOS language, which culminated with its adoption by ISO (International Standard 8807:1989) and (ii) the elaboration of two proposals for enhancing LOTOS, by introducing valuable features not present in the standard, either because they were not ready on time when it was adopted or because they did not reach international consensus.

The present article is organized as follows. Section 2 recalls the contributions of Ed Brinksma to the definition of LOTOS and gives a brief account of the impact of this language in academia and industry. The two next sections review two early languages proposed by Ed Brinksma for enhancing LOTOS, namely Extended LOTOS (Sect. 3) and Modular LOTOS (Sect. 4). The three next sections present three more recent languages that, between 1993 and now, have been proposed to supersede LOTOS, namely E-LOTOS (Sect. 5), LOTOS NT (Sect. 6), and LNT (Sect. 7), with some discussion about the actual impact of these languages. Finally, Section 8 gives a few concluding remarks.

2 LOTOS

Among all publications of Ed Brinksma related to the definition of LOTOS, we highlight three key contributions, each of a different nature and scope:

- Obviously, the ISO Draft International Standard defining LOTOS [59] occupies a place of choice. Even if earlier drafts of LOTOS had circulated before (e.g., Ed Brinksma’s first tutorial on LOTOS [10] given in 1985) and even if experiments with LOTOS had already been done at some universities (e.g., the model-checking verification of protocols in 1986 [35, 34]), this Draft International Standard published in 1987 was the first complete, coherent definition of LOTOS made available to the international community. Two years after, this document reached its final status by being approved as the ISO International Standard 8807:1989 [61].

The definition of LOTOS was a collective achievement done within an ISO committee (project 97.21.20.2) under the leadership of Ed Brinksma, who was the editor in charge of producing the standard. Tommaso Bolognesi, Günter Karjoth, Luigi Logrippo, Jan de Meer, Elie Najm, Juan Quemada, Pippo Scollo, Alaister Tocher, Jan Tretmans, and Chris Vissers participated, among others, in this committee.

The resulting LOTOS language was an audacious combination of the most recent innovations in formal methods at that time. To describe and manipulate data structures, the LOTOS committee selected abstract data types — more precisely, a dialect of the algebraic language ACT ONE [26] [27] [24]. To describe the behaviour of concurrent processes, the committee retained the key ideas of process algebra, blending into a single language the best features of several calculi, namely CCS [82], TCSP [17], and Circal [81]. LOTOS also brought original ideas, such as its “disable” operator, which models nondeterministic disruption (e.g., crashes and failures), and its “enable” operator, which allows value-passing sequential continuation after the termination of a group of parallel processes.

The definition of LOTOS provided in the ISO standard was fully formal, much in line with the longstanding Dutch tradition of computer-language definitions. The syntax was given as a BNF grammar; the static semantics was specified as a set of mathematical constraints and functions defined by induction over syntactic constructs; the semantics of data types was expressed as a many-sorted term algebra obtained by quotienting the algebra generated by a derivation system; finally the behavioural semantics of processes was defined operationally using a set of structured operational semantics rules. This formal definition was followed by annexes providing informal explanations and complementary information.

- Jointly written with Tommaso Bolognesi, Ed Brinskma’s tutorial on LOTOS [6] is also a highly cited publication. Written in a lively style and illustrated with a wealth of examples, this tutorial targets the end users of LOTOS. It is orthogonal and complementary to the (somewhat dry) ISO standard definition, primarily oriented towards language implementers and semanticists.
- Another insightful contribution is Ed Brinskma’s 1989 paper on constraint-oriented specification [12]. It is well-known that the decomposition of a computer system into concurrent/parallel tasks may take two forms: it is either *physical* if the decomposition closely reflects the actual distribution of tasks over processors, or *logical* otherwise, if the decomposition is rather intended to provide the system with a modular structure that does not necessarily correspond to its actual topology. Ed Brinskma develops the latter approach in the framework of the LOTOS multiway rendezvous, which enables two or more processes to synchronize, negotiate, and exchange data values during one atomic event. The paper formulates the fundamental intuition of *parallel composition as conjunction*, meaning that the multiway rendezvous achieves the logical conjunction of all the individual constraints expressed by a set of processes running concurrently. This idea enables a certain degree of “declarative” programming (namely, constraint solving) to be introduced in the framework of a fundamentally “operational” (i.e., constructive, imperative) language such as LOTOS. The usefulness of the approach is demonstrated on realistic examples of communication protocols [12], but it is also relevant to other application domains, e.g., hardware circuits ([42] shows how the complex arbitration protocol of the SCSI-2 bus can be concisely modelled using an eight-party LOTOS rendezvous) or robotics ([48] illustrates how a software controller for an entire manufacturing plant can be obtained as the parallel composition of many simple controllers, one for each device or degree of freedom of a device in the plant).

Retrospectively, the international effort invested in LOTOS was successful in several respects:

- Although LOTOS is a committee-designed language based on two very different concepts (algebraic data types and process calculi), it achieves a suitable compromise and a fair integration between its various elements. All its language constructs (perhaps with the exception of the **choice** and

par operators on gate lists) derive from concrete needs and are useful in practice.

- LOTOS is clearly more abstract and higher level than the two other standardized languages it was competing with (namely, Estelle [60] and SDL [18]), and proved that a specification language could be formal and executable at the same time.
- The design of LOTOS made it clear that process calculi were not only mathematical notations for studying concurrency theory, but that they could be turned into computer languages used to model real-life systems. LOTOS was indeed the first process calculus in which large specifications of complex systems (e.g., protocols and services of OSI and ISDN networks) were produced. Later, it was shown that the high abstraction level of LOTOS makes it also suitable to other application domains, e.g., multiprocessor architectures and asynchronous circuits.
- The LOTOS community put strong emphasis on software tools, often in the framework of European projects such as SEDOS, LOTOSphere, SPECS, EUCALYPTUS-1 and -2, etc. Today, most of these tools are no longer available, but the CADP toolbox¹ [46] is still actively maintained. Also, many ideas present in early LOTOS tools would certainly benefit from modern developments in symbolic execution and verification technology.

On the negative side, one can point out two main shortcomings of LOTOS:

- Despite its status of international standard, LOTOS did not manage to unite the academic community working on process calculi. Not only the preexisting algebras/calculi ACP, CCS, and CSP remained, but new languages appeared, e.g., μ CRL. This resulted in fragmented efforts and a lack of critical mass that became apparent in the mid-90s.
- LOTOS also failed to gain wide industrial acceptance, mostly due to its so-called “steep learning curve”. Because it is an abstract, expressive, and flexible language based on concepts absent from mainstream languages, LOTOS is best used by high-level experts rather than average software programmers: this is unfortunately a fatal flaw as far as dissemination is concerned.

3 Extended LOTOS

As soon as the definition of LOTOS was frozen as an ISO standard, it appeared that the language was not fully satisfactory and that some of its features could be redesigned in a better way. Ed Brinksma’s role as the editor of the LOTOS standard did not prevent him from suggesting enhancements to LOTOS.

His first contribution in this respect is his PhD thesis [11], defended in 1988, which proposes a language named “Extended LOTOS” that significantly differs

¹<http://cadp.inria.fr>

from LOTOS. Concerning data specifications, Extended LOTOS keeps the abstract data types of LOTOS, but adds better support for modules. Concerning behavioural specifications (namely, concurrent processes), Extended LOTOS brings deeper changes:

- It introduces a notion of *action product* inspired from SCCS [83], whereas LOTOS only has simple actions.
- Extended LOTOS attempts at unifying in one single operator both forms of sequential composition (action prefix and “enable”) that exist in LOTOS.
- Extended LOTOS breaks with the algebraic style of LOTOS and other process calculi by replacing unary and binary operators with n -ary constructs having a fully bracketed syntax, e.g., “**sel** B_1 [] B_2 [] ... [] B_n **endsel**” for nondeterministic choice or “**par** B_1 || B_2 || ... || B_n **endpar**” for parallel composition.
- Extended LOTOS proposes other desirable features, among which a **par** operator ranging over a finite domain of values.

Although Extended LOTOS has never been actually implemented, these ideas had the merit to point out the main shortcomings of LOTOS and made it clear that the language, despite its status of international standard, still deserved major enhancements.

4 Modular LOTOS

Published three years later, a deliverable (edited by Ed Brinksma) of the LOTOSphere project [13] adopts a point of view orthogonal to that of Extended LOTOS: leaving aside all ideas for improving the behaviour part of LOTOS, this deliverable focuses on enhancements to the data part of LOTOS, in which usability problems have been identified as most crucial, and proposes a new language called “Modular LOTOS”, two synthetic presentations of which can also be found in [9] and [90]. Modular LOTOS suggests the following enhancements:

- Distinction between *constructors* and *functions*, whereas LOTOS made no difference between these two forms of operations;
- Introduction of *partial functions*, whereas LOTOS only allowed totally defined operations;
- Support for *built-in types* (e.g., natural numbers, integer numbers, strings) and *generic data structures* (e.g., lists, sets, arrays, etc.);
- Introduction of *modules* gathering data and/or behaviour definitions, namely, types, constructors, functions, and processes;
- Introduction of module interfaces (called *descriptions*) that can be used to hide certain definitions contained in modules;

- Introduction of *renaming* to avoid name clashes between different modules;
- Support for *generic modules* parameterized by descriptions.

To our knowledge, Modular LOTOS has never been implemented, although key ideas (namely, distinction between constructors and functions, partial functions, and splitting of large LOTOS specifications into multiple files) were already supported in the *CÆSAR.ADT* compiler for LOTOS [36]. At this point, Ed Brinskma shifted his research interests to other topics, but the LOTOS reform movement he had initiated expanded rapidly.

5 E-LOTOS

Between 1993 and 2001, an ISO committee gathered under the lead of Juan Quemada to revise the LOTOS standard. Arnaud Février, Hubert Garavel, Alan Jeffrey, Guy Leduc, Luc Léonard, Luigi Logrippo, José Mañas, Elie Najm, Mihaela Sighireanu, and Jacques Sincennes participated in this committee as regular contributors, with the help of more than twenty occasional contributors [89].

At the beginning, the proposed changes were modest, trying to repair rather than replace LOTOS; as time passed, it appeared that more radical enhancements were desirable. This work eventually resulted in a new language named E-LOTOS (for “Enhanced LOTOS”) approved as ISO/IEC International Standard 15437:2001 [62]. Tutorials on E-LOTOS can be found in [101], [58, 57] and [71]. Compared to LOTOS, E-LOTOS brings deep changes that aim at greater expressiveness and/or better user-friendliness:

- Concerning the data types, E-LOTOS goes far beyond the ideas suggested for Modular LOTOS. Rather than enhancing ACT-ONE, E-LOTOS removes it, replacing abstract data types with a functional language — an approach also explored in [5], which proposes a concurrent language combining a process calculus (CCS) and a functional language (ML). E-LOTOS goes even further by giving its functional language an imperative flavour: in particular, E-LOTOS variables can be assigned and E-LOTOS functions can have output (i.e., call by result) parameters to return multiple results, which, in conventional functional languages, is usually done by returning tuple values.
- E-LOTOS data types can be records (with named or unnamed fields) or (possibly recursive) types defined by a list of constructors. E-LOTOS also provides predefined types (Booleans, naturals, integers, rationals, floating-point reals, characters, and strings) and abbreviations for declaring enumerated types, records, sets, and lists.
- Contrary to LOTOS, in which the data and behaviour parts are two entirely different sub-languages, E-LOTOS tries to unify functions and processes; functions can be seen as particular cases of processes that only do local calculations before terminating, do not perform any observable or invisible action, and do not let time elapse. Consequently, functions

and processes share a number of common constructs, among which: variable assignments, **if-then-else** conditionals, **case** with pattern matching, **while** loops, **for** loops, etc.

- In both its data and behaviour parts, E-LOTOS introduces a unique sequential composition operator, which unifies the action-prefix and “enable” operators present in the behaviour part of LOTOS.
- E-LOTOS provides support for exception handling. In the data part of E-LOTOS, exceptions bring a convenient solution to the need for partial functions. In the behaviour part, exceptions allow some involved communication protocols to be described compositionally — see [50] for an advocacy paper on exceptions in process calculi.
- Gates (i.e., communication ports) are explicitly typed in E-LOTOS, whereas they are untyped in LOTOS — see [37] for an introduction to gate typing, which leads to more readable specifications, detects communication mismatches at compile time rather than at run time, and provides a simple solution to the “structured events” issue in the constraint-oriented style.
- To express *quantitative time* aspects, the behaviour part of E-LOTOS allows to specify constraints on the duration of actions and/or the instant(s) at which they may occur. Such features are required to describe isochronous protocols and real-time systems precisely, and many timed extensions of LOTOS have been proposed, e.g., ET-LOTOS [72, 73] and RT-LOTOS [22].
- The behaviour part of E-LOTOS introduces a n -ary parallel operator [52] that generalizes the three binary parallel composition operators of LOTOS. This new operator is easier to use, more readable, and enables m -among- n synchronization (in particular, the 2-among- n synchronization of CCS).
- The behaviour part of E-LOTOS also introduces new operators, such as *rename* (which allows to change the name of observable actions and exceptions, to merge or split gates, and to add or remove offers from actions) and *suspend/resume* (which generalizes the “disable” operator of LOTOS by allowing resumable interrupts to be modelled).
- Finally, E-LOTOS provides *modules* that may contain types, functions, and/or processes. Modules can be imported and exported; they have *interfaces* for information hiding and can be *generic*.

Due to its new features resulting from multiple, sometimes conflicting influences, and despite the unification between functions and processes, E-LOTOS is a complex language, with an impressive number of semantic rules. The E-LOTOS standard has 120 pages (+80 pages of annexes), while the LOTOS standard has only 70 pages (+70 pages of annexes). It is therefore unclear whether E-LOTOS brings a satisfactory answer to the “steep learning curve” issue with LOTOS.

This probably explains why E-LOTOS only had a marginal impact in practice. Very few case studies have been done using E-LOTOS; one can mention [96, 99]

(which compares LOTOS and E-LOTOS on a common example), [25], [21], and [92, 93]. To our knowledge, E-LOTOS has never been implemented in software tools (except perhaps [25] or [75]) nor taught in university classes.

In some sense, the shift from LOTOS to E-LOTOS is reminiscent of the shift from Algol 60 to Algol 68: a simple, elegant, yet limited language was replaced by a larger, more expressive, formally defined language, which, because of its growth in complexity, failed to build a sufficient momentum of interest among its potential users.

6 LOTOS NT

6.1 Design of LOTOS NT

In 1997, when it became manifest that E-LOTOS was getting too large and too complex, INRIA Grenoble started investigating a fallback solution. This led to the design of LOTOS NT (where “NT” stands for “New Technology”), a simplified dialect of E-LOTOS that could be feasibly implemented and provide an actual replacement solution for LOTOS.

It was decided to not introduce in LOTOS NT some questionable features that significantly contribute to the complexity of E-LOTOS, among which: type synonyms, anonymous tuples (i.e., the possibility, borrowed from ML, that any list of values put between parentheses creates a new value having a valid, yet undeclared tuple type), extensible records, type equality relation based on structure equivalence (rather than name equivalence), subtyping relation based on record subtyping, etc.

For the same reasons, two features present in E-LOTOS but absent from LOTOS, the suspend-resume operator and the support for quantitative time, were not introduced in LOTOS NT, as it was felt that the potential applications of such features were already covered by competing formalisms such as timed automata [2] and were not worth the effort/impact ratio.

The formal definition of LOTOS NT (syntax, static semantics, and dynamic semantics) was given in [94]. Rationale for the design of LOTOS NT (and of E-LOTOS as well, since LOTOS NT influenced the latest evolutions of E-LOTOS) can be found in [51].

6.2 Implementation of LOTOS NT

To implement this language, a compiler named TRAIAN² [95] has been developed at INRIA Grenoble since 1997. It is built using the SYNTAX [7] and FNC-2 [64] compiler-generation tools designed at INRIA Rocquencourt. Unfortunately, FNC-2 ceased to be maintained in 1999, which prevented TRAIAN from being completed; as a consequence, TRAIAN only handles the data part of LOTOS NT (i.e., types and functions) but not the behaviour part (i.e., processes and channels).

²<http://vasy.inria.fr/traian>

As it is, TRAIAN performs lexical analysis, syntactic analysis, abstract syntax tree construction, static semantics analysis of LOTOS NT data specifications, and translates these into C programs, which can in turn be compiled and executed. TRAIAN has been regularly maintained and enhanced: ten releases have been issued since 1998, the latest version of TRAIAN (dated 2016) containing 55,500 lines of FNC-2 and C code.

6.3 Applications of LOTOS NT

Although TRAIAN only supports a fragment of LOTOS NT, it has useful applications in compiler construction. Our approach [45] consists in using the SYNTAX compiler generator for the lexical and syntactic analyses, together with LOTOS NT for semantical aspects, in particular the definition, construction, and traversals of abstract trees. Some involved parts of the compiler can be written directly in C if necessary, but most of the compiler is usually written in LOTOS NT, which is then translated into C code by TRAIAN.

The combined use of SYNTAX, LOTOS NT, and TRAIAN proves to be satisfactory, as regards both the rapidity of development and the quality of resulting compilers. So far, twelve compilers have been developed at INRIA Grenoble using this approach: AAL [76], ATLANTIF [97], CHP2LOTOS [47], CTRL2BLK [77], EVALUATOR 4.0 [80], EXP.OPEN 2.0 [68], FSP2LOTOS [69], GRL2LNT [63], LNT2LOTOS [19], NTIF [44], PIC2LNT [78], and SVL [43, 67, 23].

7 LNT

7.1 Design of LNT

Because of the limitations of TRAIAN, LOTOS NT does not provide a replacement solution for LOTOS. The need for a better language based on process calculi remains [39] [40], even if all prior attempts have failed to provide a usable solution.

In 2005, a new opportunity was found to progress this agenda: the Bull company was interested in using the CADP toolbox to formally verify multiprocessor architectures, but was reluctant to use LOTOS as a modelling language, mostly due to the verbosity of the LOTOS data part. To ease the writing of large specifications by Bull, still using the existing CADP tools, INRIA Grenoble undertook the development of a translator to convert LOTOS NT data types and functions into LOTOS ones. This made it possible to produce specifications combining a data part written in LOTOS NT (more concise and less error-prone than LOTOS) and a behaviour part written in LOTOS. The translator converted such composite specifications into plain LOTOS ones, which then could be analyzed by the CADP tools.

A first version of this translator was delivered to Bull in July 2005. Since then, the translator has been constantly improved and extended to handle new LOTOS NT features. In 2007, support for the behaviour part of LOTOS NT was added; this progressively removed the need for composite specifications, as

it became possible to write entire specifications in LOTOS NT, with no LOTOS code at all.

Due to the rapid evolution of this translator, its input language gradually diverged from the original LOTOS NT implemented in TRAIAN, which remained quite stable in comparison. To avoid ambiguities, it was decided in 2014 to give this input language a new name (“LNT”), while reserving the name “LOTOS NT” for the language accepted by TRAIAN — such a distinction was not made in papers published before Spring 2014, in which LOTOS NT and LNT were used as synonyms.

The current definition of LNT is given in [19]. In a nutshell, LNT combines, in a single language with an Ada-like syntax designed to favour readability, selected features borrowed from imperative languages, functional languages, and value-passing process calculi:

- An LNT specification is a set of *modules*, each of which may import other modules and define *types*, *functions*, *channels*, and/or *processes*.
- A *type* is either predefined (namely, **bool**, **nat**, **int**, **real**, **char**, and **string**), defined by specifying the *free constructors* that generate its domain of values, or defined using the *type combinators* **array**, **list**, **range**, **set**, **sorted list**, **sorted set**, and **where** (the latter enabling predicate subtyping).
- A *function* is either predefined (namely, logical, arithmetical, and relational operations on predefined types), automatically generated for some user-defined type (such as free constructors, but also equality, order relations, field accessors and selectors, etc., which are generated if the user requests them), or have a handwritten definition provided by the user.
- A *channel* is a gate type that, following the ideas of [37], specifies the types of values to be sent or received during interactions on a given gate. There exist two special channels: **none**, which expresses that no value can be sent or received (this is useful for pure synchronization and exceptions), and **any**, which permits all values to be sent or received (this allows gates to be “untyped”, as in LOTOS, thus ensuring backward compatibility).
- A *process* is a program fragment that, as in LOTOS and other process calculi, executes and communicates with its environment by sending and/or receiving values on a set of gates.

Globally, LNT has four different kinds of routines, of increasing complexity:

- A *constructor* has only **in** parameters (call by value), no explicit definition, and does not raise exceptions.
- A *pure function* has only **in** parameters, an implicit or explicit definition, and may raise exceptions if needed (this provides for partially-defined functions).
- A *procedural function* (or *procedure*, for short) may have **in**, **out** (call by result), or **in-out** (call by value-result) parameters; unlike constructors

and pure functions, it does not necessarily return a result; it usually has an explicit definition and may raise exceptions.

- A *process* may also have **in**, **out**, or **in-out** parameters; it has an explicit definition, may raise exceptions [50], and interacts with its environment by means of gates. The key difference between processes and other routines is that the execution of processes can be nondeterministic and let time elapse (the execution semantics is that of process calculi and labelled transition systems) whereas the three former kinds of routines execute deterministically and atomically (the execution semantics is that of functional languages).

LNT possesses three main concepts for denoting computation:

- An *expression* corresponds to the usual notion of expression in imperative programming languages. It is an algebraic term built using constants, variables, and calls to constructors and pure functions. The evaluation of each expression is deterministic (it always returns the same result or raises the same exception), atomic (it is expected to terminate and take a negligible amount of time), and free from side effects (it does not modify variables).
- An *instruction* corresponds to the usual notion of statement in imperative programming languages. Instructions serve to explicitly define the bodies of LNT functions. Basic instructions include: **null** (which does nothing), assignment to a variable or an array element, **return** of a function result, **raise** of an exception, procedure call, **assert**, etc. Instructions can be combined using structured-programming constructs, such as sequential composition, **if-then-else** conditionals, **case** with pattern matching, **for** and **while** loops, loops with **break** clauses, and declarations of variables with a limited scope. Because instructions manipulate and modify a store, the semantics of LNT relies on static analysis to prohibit all situations where uninitialized variables could be used; this way, instructions have an imperative-programming syntax and a functional-programming semantics. Like the evaluation of expressions, the execution of instructions is deterministic and atomic.
- A *behaviour* is the LNT equivalent of a LOTOS “behaviour expression”. Behaviours serve to define the bodies of LNT processes. Behaviours can be seen as a superset of instructions since they contain all instructions (except **return**) but also include additional constructs specific to process calculi: **stop** (deadlock), communication on a gate (possibly with sending and/or receiving values), assignment of a non-deterministic value to a variable, process call, forever loop without **break** clause, non-deterministic choice (which is n -ary, rather than binary), parallel composition (which is n -ary and *graphical* [52], i.e., explicitly describes the communications/synchronizations between concurrent behaviours), gate hiding, and disruption (i.e., the “disable” operators of LOTOS). Unlike instructions, the execution of behaviours is nondeterministic, non-atomic, and may never terminate.

In the design of LNT, one main attainment is the integration in a single language of two very different models of computations: imperative/functional languages and process calculi. This was not the case with LOTOS, nor with its competitors Estelle and SDL, both of which clumsily amalgamate state machines with another formalism for data computation. Such a unification, which had been tried without success in E-LOTOS, is now effective in LNT. A key issue was the status of sequential composition [41], which led to question and discard some well-established habits of process calculi, especially the action-prefix operator of CCS and LOTOS, and the use of dynamic semantics in place of static semantics.

7.2 Implementation of LNT

Contrary to the four aforementioned languages (Extended-LOTOS, Modular LOTOS, E-LOTOS, and LOTOS NT), the definition and implementation of which were planned as two successive steps (the latter being never undertaken or never completed), LNT was designed in a radically different way, using an “agile” approach. Every new language feature was first implemented and assessed on a growing base of non-regression tests before being adopted for LNT.

Initially designed as a standalone tool, the translator from LNT to LOTOS became an integral part of the CADP toolbox in 2010. Actually, this translator is not one single tool, but comprises three complementary tools:

- LPP³ (“LNT PreProcessor”) is a small translator (2000 lines of C and Lex code) that expands the user-friendly LNT notations for literal constants (numbers, characters, strings, etc.) into algebraic terms making use of predefined LOTOS sorts and operations defined in custom libraries.
- LNT2LOTOS⁴ is a rather complex translator developed using the aforementioned SYNTAX/TRAIAN technology (3800 lines of SYNTAX code, 35,500 lines of LOTOS NT, and 2900 lines of C code). LNT2LOTOS translates an LNT specification into LOTOS code, possibly augmented with some little C code fragments.
- LNT.OPEN⁵ is a small utility (400 lines of shell script) that provides a top-level entry point for processing LNT specifications with the CADP tools and, more specifically, with the CÆSAR.ADT [36] and CÆSAR [49] compilers for LOTOS, and the OPEN/CÆSAR framework [38] for simulation, verification, and testing. Taking as input an LNT specification and an OPEN/CÆSAR application program, LNT.OPEN first translates (the various modules composing) the LNT specification into LOTOS by calling LPP and LNT2LOTOS, then compiles the generated LOTOS specification by calling CÆSAR.ADT and CÆSAR, and finally invokes the OPEN/CÆSAR application program to explore and analyze the corresponding state space on the fly.

Without exposing in full detail the algorithms implemented in LNT2LOTOS, these are some key principles underlying the translation:

³<http://cadp.inria.fr/man/lpp.html>

⁴<http://cadp.inria.fr/man/lnt2lotos.html>

⁵<http://cadp.inria.fr/man/lnt.open.html>

- The main guideline is to keep the translation as simple as possible, so as to swiftly upgrade the translator each time the definition of LNT evolves. Consequently, duplication of work between the translator and the LOTOS compilers is avoided, meaning that the translator does not implement certain static semantics checks at the LNT level if the same checks are later performed at the LOTOS level. In particular, the translator makes no attempt to infer and check the types of LNT value expressions, deferring these tasks to the LOTOS compilers operating on the generated code.
- The channels used for typing LNT gates raise a specific problem because, on the one hand, the LNT2LOTOS translator is not intended to perform type checking and, on the other hand, LNT gates, which are typed, are translated into LOTOS gates, which are untyped, so that type-checking errors at the LNT level cannot be detected at the LOTOS level. To address this problem, LNT2LOTOS generates, for each LNT channel C , one or several overloaded LOTOS constant functions f_C , which take as parameters the expected typed values specified for C and always return true. For each LNT action involving some gate G whose channel is C , a LOTOS Boolean guard is generated, which invokes function f_C with the input or output offers of the action, thus expressing in LOTOS the type-checking constraints arising from the definition of C . If the action is not well-typed at the LNT level, the corresponding guard will provoke at the LOTOS level a type-checking error at compile time; otherwise, the guard will evaluate to true at run time.
- The LNT2LOTOS translator performs, on LNT functions and processes, static analyses not done at the LOTOS level; for instance, it rejects (or warns about) unused variables, variables used without being assigned before, variables assigned but never used, variables shared between concurrent processes, etc. Such checks are either required by the LNT semantics (see [41] for a discussion) or suitable to ensure that LNT specifications remain as simple as possible, so as to increase readability and efficiency of verification.
- The predefined types of LNT (**bool**, **nat**, etc.) are implemented using base libraries written in LOTOS and C code. The user-defined types (built using free constructors or type combinators) are translated into LOTOS abstract data types (possibly with some additional C code meant for efficiency), together with their associated functions (equality, order relations, field accessors and selectors, etc.).
- Although, in LNT, user-defined functions and processes have the same functional/imperative style and share many constructs (e.g., assignments, **assert**, **raise**, **if-then-else**, pattern-matching **case**, **for** and **while** loops, loops with **break**, etc.), the two algorithms that translate, respectively, these functions and processes into LOTOS are very different, due to the fundamental asymmetry, in the target language, between the data part (based on abstract data types) and the behaviour part (based on process calculi).
- Our algorithm for translating LNT functions generalizes the one proposed in [88], which translates into Horn clauses a small subset of C func-

tions with only integer and list types. Our algorithm translates LNT functions into LOTOS (non-constructor) operations, which are defined using algebraic equations considered as conditional term-rewrite rules. The translation takes advantage of the rewrite strategy implemented in the CÆSAR.ADT compiler, which assumes a decreasing priority between equations. Notice that each LNT function having **out** parameters (call by result) or **in-out** parameters (call by value-result) translates to several LOTOS functions. The translation of certain LNT constructs (**assert**, **case**, and loops) also generates auxiliary LOTOS functions.

- Our algorithm for translating LNT processes takes its roots in our prior works on the translation to LOTOS of three modelling languages for hardware and software systems: the CHP2LOTOS translator [47] for CHP, the FLAC translator [4] for Fiacre, and the FSP2LOTOS translator [69] for FSP. The algorithm is involved, but four main points are worth being highlighted.

(i) Certain behavioural LNT constructs directly map to equivalent LOTOS ones. For instance, each LNT gate translates to a corresponding LOTOS gate and each LNT process translates to a corresponding LOTOS process. The algorithm benefits from the fact that both LOTOS and LNT have an *action-based* (rather than *state-based*) semantics and share a common semantic model (namely, labelled transition systems). Thanks to action-based semantics, the translation can freely introduce auxiliary LOTOS processes and variables, still preserving the semantic model (which would not be possible with state-based semantics); in particular, execution traces are identical at the LOTOS and LNT levels, which avoids the usual need for a reverse translation of diagnostics from the target to the source level.

(ii) Certain behavioural LNT constructs are too powerful to be expressed using only the behaviour part of LOTOS; for such constructs, the data part of LOTOS must also be used, by generating auxiliary sorts, operations, and algebraic equations. For instance, the **case** construct present in LNT processes is translated using both the behaviour part of LOTOS (nondeterministic choice and Boolean guards are used to express the selection between the various **case** branches) and the data part of LOTOS (equations, considered as rewrite rules, are used to express pattern matching, which is not supported by the behaviour part of LOTOS).

(iii) An involved part of the algorithm translates the LNT parallel composition operator, which is n -ary, into an algebraic combination of LOTOS parallel composition operators, which are binary. Such a translation does not always succeed, meaning that certain network topologies specified in LNT cannot be expressed in LOTOS [52]; however, we never faced this problem in real-life case studies. Also, it was not possible to introduce in LNT the concept of n -among- m synchronization proposed in [52], because it is not supported in LOTOS; such a limitation is more annoying in practice, e.g., for the specification of Web services, which quite often require 2-among- m synchronization.

(iv) Another involved part of the algorithm translates the LNT sequential composition operator (which is unique, symmetric, atomic, and lets all values of variables assigned on its left-hand side flow implicitly to its

right-hand side [41]) into one of the two LOTOS sequential composition operators, either the action-prefix operator (which is asymmetric, atomic, and lets variable values flow implicitly from its left- to its right-hand side) or the “enable” operator (which is symmetric, non-atomic as it generates a τ -transition, and forbids variable values to flow from its left- to its right-hand side except if these variables are explicitly listed in an **accept** clause). Following the principles set for the CHP2LOTOS translator [47], we chose to generate action prefix as much as possible and “enable” only when unavoidable, which produces better LOTOS code at the expense of a more involved translation. To fight state-space explosion and preserve strong equivalence between the LNT and LOTOS specifications, we slightly deviated from LOTOS semantics by adding a special pragma “(*! atomic *)” that instructs the LOTOS compiler not to generate a τ -transition when implementing the “enable” operator. There are many other algorithmic subtleties, such as the creation of auxiliary “continuation” processes for those LNT behaviours following loops and conditionals (i.e., **case**, **if**, and **select**), the translation of parallel composition occurring on the left-hand side of sequential composition where each parallel branch computes the values of different variables, the translation of **out** and **in-out** parameters of LNT processes into **exit** results returned by LOTOS processes, the need to respect the strict typing rules set by LOTOS “functionality” constraints, the optimization of tail-recursive process instantiations, etc.

In addition to the above tools, which ultimately translate an LNT specification into a sequential C program, there also exists a compiler named DLC (“Distributed LNT Compiler”) [32, 33] that translates an LNT specification into a set of C programs executing concurrently and communicating through TCP sockets; to produce such a distributed implementation, the DLC compiler exploits the concurrent architecture defined by the parallel composition operators present in the LNT specification.

7.3 Applications of LNT

The usability of the LNT language gradually increased with the progress of its translator to LOTOS. As of mid-2009, this translator was sufficiently complete and robust to allow a total shift from LOTOS to LNT at INRIA Grenoble, where no LOTOS code has been manually written since then, LNT being now the preferred high-level language for modelling concurrent systems and analyzing them using the CADP tools.

At Grenoble INP and Université Grenoble-Alpes, LNT has also replaced LOTOS to teach master students the fundamentals of concurrency theory. We observed that LNT enables students to better focus on high-level concepts, rather than getting lost in low-level details of LOTOS syntax and static semantics.

The LNT language and its tools have been used for many case studies, at INRIA Grenoble and in other academic or industrial labs as well (we only mention those not affiliated with the authors’ institutions):

- *Avionics*: verification of an equipment failure management protocol⁶ and of a ground-plane communication protocol⁷ [53, 98] provided by Airbus;
- *Cloud computing*: verification of self-configuration protocols⁸ [28] (Orange Labs), of the Synergy reconfiguration protocol for component-based systems⁹ [8], and of dynamic management protocol for cloud applications¹⁰ [1];
- *Distributed algorithms*: verification and performance evaluation of mutual exclusion protocols¹¹ [79], verification of multiway synchronization protocols¹² [31, 29, 33], specification and rapid prototyping of Stanford's RAFT distributed consensus algorithm¹³ [30, 33], and performance evaluation of concurrent data structures¹⁴ [102] (RWTH Aachen, Germany and Chinese Academy of Sciences, Beijing, China);
- *Hardware design*: formal analysis and co-simulation of a dynamic task dispatcher¹⁵ [70] (STMicroelectronics), formal analysis of ARM's ACE cache coherency protocol¹⁶ [65, 66] (STMicroelectronics), verification and rapid prototyping of an asynchronous model of the Data Encryption Standard¹⁷ [91], verification of a fault-tolerant routing algorithm for a network-on-chip¹⁸ [103] (University of Utah, USA);
- *Human-computer interaction*: specification and validation of graphical user interfaces for a prototype control room of a nuclear power plant¹⁹ [85] and of plastic user interfaces exploiting domain ontologies²⁰ [20] (Toulouse, France);
- *Industrial systems*: model-based testing of the CANopen field bus and EnergyBus architecture²¹ [54] (Saarland University, Germany), formal specification and rapid prototyping of a software controller for a metal processing plant²² [48].

Another indication of the practical usefulness of LNT is given by its use as a target language in a growing number of translators, which implement various languages by translating them to LNT. Indeed, LNT suitably replaces LOTOS for automatically-generated code as well as for handwritten code, since the translation to LNT is much easier than the translation to LOTOS, and because

⁶<http://cadp.inria.fr/case-studies/09-k-failure-management.html>

⁷<http://cadp.inria.fr/case-studies/09-h-tftp.html>

⁸<http://cadp.inria.fr/case-studies/11-i-selfconfig.html>

⁹<http://cadp.inria.fr/case-studies/11-h-synergy.html>

¹⁰<http://cadp.inria.fr/case-studies/13-g-dynamic-management.html>

¹¹<http://cadp.inria.fr/case-studies/10-f-mutex.html>

¹²<http://cadp.inria.fr/case-studies/13-d-multiway.html>

¹³<http://cadp.inria.fr/case-studies/15-g-raft.html>

¹⁴<http://cadp.inria.fr/case-studies/16-b-concurrent.html>

¹⁵<http://cadp.inria.fr/case-studies/11-g-dtd.html>

¹⁶<http://cadp.inria.fr/case-studies/13-e-ace.html>

¹⁷<http://cadp.inria.fr/case-studies/15-f-des.html>

¹⁸<http://cadp.inria.fr/case-studies/13-f-utahnoc.html>

¹⁹<http://cadp.inria.fr/case-studies/14-d-hmi.html>

²⁰<http://cadp.inria.fr/case-studies/15-d-plastic-user-interfaces.html>

²¹<http://cadp.inria.fr/case-studies/14-c-energybus.html>

²²<http://cadp.inria.fr/case-studies/17-a-production-cell.html>

it is now preferable to let the LNT2LOTOS translator take in charge all algorithmic subtleties required to produce valid and efficient LOTOS code. We are aware of the following tools (again, we do not mention the authors' institutions):

- The BPMN2Py/Py2LNT translators²³ [56, 86] for analyzing choreographies of Web services specified in WS-CDL (Université de Nantes, France);
- The CMT translator [74] for the BPEL/WSDL specification languages for Web services (Tsinghua University, Beijing, China and MIT, Cambridge, MA, USA), and another, more complete algorithm for translating BPEL/WDSL/Xpath/XML Schema to LNT [98];
- The DFTCalc tool²⁴ [3, 55] for Dynamic Fault Trees (University of Twente, The Netherlands);
- The EB³2LNT translator²⁵ [100] for the EB³ specification language for information systems (Université Paris Est, France);
- The GRL2LNT translator²⁶ [63] for the GRL specification language for GALS (*Globally Asynchronous Locally Synchronous*) systems;
- The OCARINA tool²⁷ [84] for the AADL architecture description language (ISAE, Toulouse, France and University of Sfax, Tunisia);
- The PIC2LNT translator²⁸ [78] for the applied π -calculus (an extension of the π -calculus with typed data values);
- the VBPMN translator²⁹ [87] for the BPMN language for describing business processes (Université Paris Ouest, France).

8 Conclusion

Computer systems handling asynchronous concurrency are inherently complex and cannot be reliably designed without adequate specification languages supported by sound analysis tools. Ed Brinksma contributed to this agenda in two significant ways: (i) by leading the definition and standardization of the LOTOS language, which exposed the key ideas of process calculi to a large audience and sparked considerable interest in academia and industry; (ii) by sending a clear signal that LOTOS, despite its qualities, was not the end of the road and that further enhancements were possible and desirable.

The present paper provided a retrospective account of the evolution of the LOTOS-based family of specification languages, starting from LOTOS itself, reviewing the successive proposals for enhancing LOTOS (Extended LOTOS, Modular LOTOS, E-LOTOS, and LOTOS NT), and ending with LNT, the

²³<http://cadp.inria.fr/software/12-e-choreography.html>

²⁴<http://cadp.inria.fr/software/12-i-dftcalc.html>

²⁵<http://cadp.inria.fr/software/13-a-eb3.html>

²⁶<http://cadp.inria.fr/software/14-c-grl.html>

²⁷<http://cadp.inria.fr/software/15-b-ocarina.html>

²⁸<http://cadp.inria.fr/software/13-d-pic2lnt.html>

²⁹<http://cadp.inria.fr/software/16-a-vbpnm.html>

most recent descendent, which preserves the most valuable ideas of process calculi but entirely modifies the shape of the language to make it compatible with mainstream programming languages. The feedback acquired by using LNT for the design of complex industrial systems suggests that LNT provides a viable and effective replacement for LOTOS. Quoting STMicroelectronics engineers: “Although modeling the [Dynamic Task Dispatcher] in a classical formal specification language, such as LOTOS, is theoretically possible, using LNT made the development of a formal model practically feasible” [70].

Concerning future work, we can highlight two main research directions:

- The LNT language is not yet frozen and can still be further enhanced. For instance, the unification of exceptions across functions and processes is almost complete. We now consider equipping processes with optional **return** behaviours, so that functions become a strict subset of processes. We also plan to introduce, beyond assertions that already exist in LNT, pre-conditions, post-conditions, and loop invariants that would allow the application of mainstream theorem provers and static analyzers to LNT specifications.
- The current implementation of LNT by translation to LOTOS is justified by the reuse of existing LOTOS tools. It is intellectually challenging, but sometimes overly complex: for instance, LNT functions, written in a functional/imperative style, are first translated to LOTOS algebraic equations, and then compiled back to imperative C code. A native implementation of LNT would certainly be simpler and more efficient; it would also overcome certain LOTOS limitations that currently prevent useful constructs, such as the **trap** operator for exception catching [50] and the n -among- m synchronization pattern in parallel composition [52], from being added to LNT.

References

- [1] Rim Abid, Gwen Salaün, Francesco Bongiovanni, and Noël de Palma. Verification of a Dynamic Management Protocol for Cloud Applications. In Dang Van Hung and Mizuhito Ogawa, editors, *Proceedings of the 11th International Symposium on Automated Technology for Verification and Analysis (ATVA'13), Hanoi, Vietnam*, volume 8172 of *Lecture Notes in Computer Science*, pages 178–192. Springer, 2013.
- [2] Rajeev Alur and David L. Dill. A Theory of Timed Automata. *Theoretical Computer Science*, 126(2):183–235, April 1994.
- [3] Florian Arnold, Axel Belinfante, Freark van der Berg, Dennis Guck, and Marielle Stoelinga. DFTCalc: A Tool for Efficient Fault Tree Analysis. In Friedemann Bitsch, Jérémie Guiochet, and Mohamed Kaâniche, editors, *Proceedings of the 32nd International Conference on Computer Safety, Reliability, and Security (SAFECOMP'13), Toulouse, France*, volume 8153 of *Lecture Notes in Computer Science*, pages 293–301. Springer, 2013.

-
- [4] Bernard Berthomieu, Jean-Paul Bodeveix, Patrick Farail, Mamoun Filali, Hubert Garavel, Pierre Gauffillet, Frédéric Lang, and François Vernadat. FIACRE: An Intermediate Language for Model Verification in the TOPCASED Environment. In Jean-Claude Laprie, editor, *Proceedings of the 4th European Congress on Embedded Real-Time Software (ERTS'08)*, Toulouse, France, January 2008.
- [5] Bernard Berthomieu and Thierry Le Sergent. Programming with Behaviors in an ML Framework – The Syntax and Semantics of LCS. In Donald Sannella, editor, *Proceedings of the 5th European Symposium on Programming Languages and Systems (ESOP'94)*, Edinburgh, U.K, volume 788 of *Lecture Notes in Computer Science*, pages 89–104. Springer, April 1994.
- [6] Tommaso Bolognesi and Ed Brinksma. Introduction to the ISO Specification Language LOTOS. *Computer Networks and ISDN Systems*, 14(1):25–59, January 1988.
- [7] Pierre Boullier and Martin Jourdan. A New Error Repair and Recovery Scheme for Lexical and Syntactic Analysis. *Science of Computer Programming*, 9(3):271–286, 1987.
- [8] Fabienne Boyer, Olivier Gruber, and Gwen Salaün. Specifying and Verifying the SYNERGY Reconfiguration Protocol with LOTOS NT and CADP. In Michael Butler and Wolfram Schulte, editors, *Proceedings of the 17th International Symposium on Formal Methods (FM'11)*, Limerick, Ireland, volume 6664 of *Lecture Notes in Computer Science*, pages 103–117. Springer, June 2011.
- [9] E. Brinksma and G. Leih. Enhancements of LOTOS. In T. Bolognesi, J. van de Lagemaat, and C. Vissers, editors, *LOTOSphere: Software Development with LOTOS*, pages 453–466. Kluwer Academic Publishers, 1995.
- [10] Ed Brinksma. A Tutorial on LOTOS. In Michel Diaz, editor, *Proceedings of the 5th IFIP International Workshop on Protocol Specification, Testing and Verification (PSTV'85)*, Moissac, France, pages 171–194. North-Holland, June 1985.
- [11] Ed Brinksma. *On the Design of Extended LOTOS – A Specification Language for Open Distributed Systems*. PhD thesis, University of Twente, November 1988.
- [12] Ed Brinksma. Constraint-Oriented Specification in a Constructive Formal Description Technique. In J. W. de Bakker, Willem P. de Roever, and Grzegorz Rozenberg, editors, *Proceedings of the REX Workshop on Stepwise Refinement of Distributed Systems – Models, Formalisms, Correctness – Mook, The Netherlands*, volume 430 of *Lecture Notes in Computer Science*, pages 130–152. Springer, May 1989.
- [13] Ed Brinksma. Task 1.4 Deliverable on Language Enhancements, LOTOSphere (ESPRIT Projet 2304) Document ref. Lo/WP1/T1.4/N0016/V3, 146 pages, April 1992.

- [14] Ed Brinksma and Günter Karjoth. A Specification of the OSI Transport Service in LOTOS. In Yechiam Yemini, Robert E. Strom, and Shaula Yemini, editors, *Proceedings of the 4th IFIP International Workshop on Protocol Specification, Testing and Verification, Skytop Lodge, PA, USA*, pages 227–251. North-Holland, June 1984.
- [15] Ed Brinksma, Joost-Pieter Katoen, Rom Langerak, and Diego Latella. A Stochastic Causality-Based Process Algebra. *The Computer Journal*, 38(7):552–565, 1995.
- [16] Ed Brinksma, Jan Tretmans, and Louis Verhaard. A Framework for Test Selection. In Bengt Jonsson, Joachim Parrow, and Björn Pehrson, editors, *Proceedings of the IFIP WG6.1 9th International Symposium on Protocol Specification, Testing and Verification, Stockholm, Sweden*, pages 233–248. North-Holland, June 1991.
- [17] Stephen D. Brookes, C. A. R. Hoare, and A. W. Roscoe. A Theory of Communicating Sequential Processes. *Journal of the ACM*, 31(3):560–599, July 1984.
- [18] CCITT. Specification and Description Language. Recommendation Z.100, International Consultative Committee for Telephony and Telegraphy, Geneva, March 1988.
- [19] David Champelovier, Xavier Clerc, Hubert Garavel, Yves Guerte, Christine McKinty, Vincent Powazny, Frédéric Lang, Wendelin Serwe, and Gideon Smeding. Reference Manual of the LNT to LOTOS Translator (Version 6.7). INRIA, Grenoble, France, July 2017.
- [20] Abdelkrim Chebieb and Yamine Aït Ameer. Formal Verification of Plastic User Interfaces Exploiting Domain Ontologies. In Huang Zhiqiu and Sun Jun, editors, *Proceedings of the International Symposium on Theoretical Aspects of Software Engineering (TASE'15), Nanjing, China*, pages 79–86. IEEE Computer Society, 2015.
- [21] Robert G. Clark and Ana Moreira. Use of E-LOTOS in Adding Formality to UML. *Journal of Universal Computer Science*, 6(11):1071–1087, November 2000.
- [22] Jean-Pierre Courtiat, Celso A. S. Santos, Christophe Lohr, and B. Outtaj. Experience with RT-LOTOS, a Temporal Extension of the LOTOS Formal Description Technique. *Computer Communications*, 23(12):1104–1123, 2000.
- [23] Pepijn Crouzen and Frédéric Lang. Smart Reduction. In Dimitra Giannakopoulou and Fernando Orejas, editors, *Proceedings of Fundamental Approaches to Software Engineering (FASE'11), Saarbrücken, Germany*, volume 6603 of *Lecture Notes in Computer Science*, pages 111–126. Springer, March 2011.
- [24] Jan de Meer, Rudolf Roth, and Son Vuong. Introduction to Algebraic Specifications Based on the Language ACT ONE. *Computer Networks and ISDN Systems*, 23(5):363–392, 1992.

- [25] Wanderley Lopes de Souza, Paulo Nazareno Maia Sampaio, Maria Janilce B. Almeida, Luciano Paschoal Gaspary, Lisandro Zambenedetti Granville, Jean-Marie Farines, Roberto Milton Scheffel, Roberto Willrich, Murilo S. de Camargo, and Marcelo Domingos. Design of Distributed Multimedia Applications (DAMD). In Dieter Hutter, Werner Stephan, Paolo Traverso, and Markus Ullmann, editors, *Applied Formal Methods – Proceedings of the International Workshop on Current Trends in Applied Formal Methods (FM-Trends’98)*, Boppard, Germany, pages 77–91. Springer, October 1999.
- [26] Hartmut Ehrig, Werner Fey, and Horst Hansen. An Algebraic Specification Language with Two Levels of Semantics. Bericht No 83-03, Fachbereich 20-Informatik, Technische Universität Berlin, 1983.
- [27] Hartmut Ehrig and Bernd Mahr. *Fundamentals of Algebraic Specification 1 – Equations and Initial Semantics*, volume 6 of *EATCS Monographs on Theoretical Computer Science*. Springer, 1985.
- [28] Xavier Etchevers, Gwen Salaün, Fabienne Boyer, Thierry Coupaye, and Noël de Palma. Reliable Self-Deployment of Distributed Cloud Applications. *Software Practice and Experience*, 47(1):3–20, 2017.
- [29] Hugues Evrard. *Génération automatique d’implémentation distribuée à partir de modèles formels de processus concurrents asynchrones*. Thèse de Doctorat, Université de Grenoble, July 2015.
- [30] Hugues Evrard. DLC: Compiling a Concurrent System Formal Specification to a Distributed Implementation. In Marsha Chechik and Jean-François Raskin, editors, *Proceedings of the 22nd International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS’16)*, Eindhoven, The Netherlands, volume 9636 of *Lecture Notes in Computer Science*, pages 553–559. Springer, 2016.
- [31] Hugues Evrard and Frédéric Lang. Formal Verification of Distributed Branching Multiway Synchronization Protocols. In Dirk Beyer and Michele Boreale, editors, *Proceedings of the IFIP Joint International Conference on Formal Techniques for Distributed Systems (FORTE/FMOODS’13)*, Florence, Italy, volume 7892 of *Lecture Notes in Computer Science*, pages 146–160. Springer, June 2013.
- [32] Hugues Evrard and Frédéric Lang. Automatic Distributed Code Generation from Formal Models of Asynchronous Concurrent Processes. In Marco Aldinucci, Masoud Daneshmand, Ville Leppänen, and Johan Lilius, editors, *Proceedings of the 23rd Euromicro International Conference on Parallel, Distributed and Network-based Processing – Special Session on Formal Approaches to Parallel and Distributed Systems (PDP/4PAD’15)*, Turku, Finland, pages 459–466. IEEE Computer Society Press, March 2015.
- [33] Hugues Evrard and Frédéric Lang. Automatic Distributed Code Generation from Formal Models of Asynchronous Processes Interacting by Multiway Rendezvous. *Journal of Logical and Algebraic Methods in Programming*, 88:121–153, 2017.

- [34] Hubert Garavel. Utilisation du système CESAR pour la vérification de protocoles spécifiés en LOTOS. Rapport SPECTRE C2, Laboratoire de Génie Informatique – Institut IMAG, Grenoble, December 1986.
- [35] Hubert Garavel. Vérification de programmes LOTOS à l’aide du système QUASAR. Master’s thesis, Institut National Polytechnique de Grenoble, September 1986.
- [36] Hubert Garavel. Compilation of LOTOS Abstract Data Types. In Son T. Vuong, editor, *Proceedings of the 2nd International Conference on Formal Description Techniques FORTE’89 (Vancouver B.C., Canada)*, pages 147–162. North-Holland, December 1989.
- [37] Hubert Garavel. On the Introduction of Gate Typing in E-LOTOS. In Piotr Dembinski and Marek Sredniawa, editors, *Proceedings of the 15th IFIP International Workshop on Protocol Specification, Testing and Verification (PSTV’95), Warsaw, Poland*, pages 283–298. Chapman & Hall, June 1995.
- [38] Hubert Garavel. OPEN/CÆSAR: An Open Software Architecture for Verification, Simulation, and Testing. In Bernhard Steffen, editor, *Proceedings of the First International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS’98), Lisbon, Portugal*, volume 1384 of *Lecture Notes in Computer Science*, pages 68–84, Berlin, March 1998. Springer. Full version available as INRIA Research Report RR-3352.
- [39] Hubert Garavel. Défense et illustration des algèbres de processus. In Zoubir Mammeri, editor, *Actes de l’Ecole d’été Temps Réel ETR 2003 (Toulouse, France)*. Institut de Recherche en Informatique de Toulouse, September 2003.
- [40] Hubert Garavel. Reflections on the Future of Concurrency Theory in General and Process Calculi in Particular. In Catuscia Palamidessi and Frank D. Valencia, editors, *Proceedings of the LIX Colloquium on Emerging Trends in Concurrency Theory, Ecole Polytechnique de Paris, France, November 13–15, 2006*, volume 209 of *Electronic Notes in Theoretical Computer Science*, pages 149–164. Elsevier Science Publishers, April 2008. Also available as INRIA Research Report RR-6368.
- [41] Hubert Garavel. Revisiting Sequential Composition in Process Calculi. *Journal of Logical and Algebraic Methods in Programming*, 84(6):742–762, November 2015.
- [42] Hubert Garavel and Holger Hermanns. On Combining Functional Verification and Performance Evaluation using CADP. In Lars-Henrik Eriksson and Peter A. Lindsay, editors, *Proceedings of the 11th International Symposium of Formal Methods Europe FME’2002 (Copenhagen, Denmark)*, volume 2391 of *Lecture Notes in Computer Science*, pages 410–429. Springer, July 2002. Full version available as INRIA Research Report 4492.

- [43] Hubert Garavel and Frédéric Lang. SVL: a Scripting Language for Compositional Verification. In Myungchul Kim, Byoungmoon Chin, Sungwon Kang, and Danhyung Lee, editors, *Proceedings of the 21st IFIP WG 6.1 International Conference on Formal Techniques for Networked and Distributed Systems (FORTE'01), Cheju Island, Korea*, pages 377–392. Kluwer Academic Publishers, August 2001. Full version available as INRIA Research Report RR-4223.
- [44] Hubert Garavel and Frédéric Lang. NTIF: A General Symbolic Model for Communicating Sequential Processes with Data. In Doron Peled and Moshe Vardi, editors, *Proceedings of the 22nd IFIP WG 6.1 International Conference on Formal Techniques for Networked and Distributed Systems (FORTE'02), Houston, TX, USA*, volume 2529 of *Lecture Notes in Computer Science*, pages 276–291. Springer, November 2002. Full version available as INRIA Research Report RR-4666.
- [45] Hubert Garavel, Frédéric Lang, and Radu Mateescu. Compiler Construction using LOTOS NT. In Nigel Horspool, editor, *Proceedings of the 11th International Conference on Compiler Construction (CC'02), Grenoble, France*, volume 2304 of *Lecture Notes in Computer Science*, pages 9–13. Springer, April 2002.
- [46] Hubert Garavel, Frédéric Lang, Radu Mateescu, and Wendelin Serwe. CADP 2011: A Toolbox for the Construction and Analysis of Distributed Processes. *Springer International Journal on Software Tools for Technology Transfer (STTT)*, 15(2):89–107, April 2013.
- [47] Hubert Garavel, Gwen Salaün, and Wendelin Serwe. On the Semantics of Communicating Hardware Processes and their Translation into LOTOS for the Verification of Asynchronous Circuits with CADP. *Science of Computer Programming*, 74(3):100–127, January 2009.
- [48] Hubert Garavel and Wendelin Serwe. The Unheralded Value of the Multiway Rendezvous: Illustration with the Production Cell Benchmark. In Holger Hermanns and Peter Höfner, editors, *Proceedings of the 2nd Workshop on Models for Formal Analysis of Real Systems (MARS'17), Uppsala, Sweden*, volume 244 of *Electronic Proceedings in Theoretical Computer Science*, pages 230–270, April 2017.
- [49] Hubert Garavel and Joseph Sifakis. Compilation and Verification of LOTOS Specifications. In L. Logrippo, R. L. Probert, and H. Ural, editors, *Proceedings of the 10th IFIP International Symposium on Protocol Specification, Testing and Verification (PSTV'90), Ottawa, Canada*, pages 379–394. North-Holland, June 1990.
- [50] Hubert Garavel and Mihaela Sighireanu. On the Introduction of Exceptions in LOTOS. In Reinhard Gotzhein and Jan Brederke, editors, *Proceedings of the IFIP Joint International Conference on Formal Description Techniques for Distributed Systems and Communication Protocols, and Protocol Specification, Testing, and Verification (FORTE/PSTV'96), Kaiserslautern, Germany*, pages 469–484. Chapman & Hall, October 1996.

- [51] Hubert Garavel and Mihaela Sighireanu. Towards a Second Generation of Formal Description Techniques – Rationale for the Design of E-LOTOS. In Jan-Friso Groote, Bas Luttik, and Jos van Wamel, editors, *Proceedings of the 3rd International Workshop on Formal Methods for Industrial Critical Systems (FMICS'98), Amsterdam, The Netherlands*, pages 187–230, Amsterdam, May 1998. CWI. Invited lecture.
- [52] Hubert Garavel and Mihaela Sighireanu. A Graphical Parallel Composition Operator for Process Algebras. In Jianping Wu, Qiang Gao, and Samuel T. Chanson, editors, *Proceedings of the IFIP Joint International Conference on Formal Description Techniques for Distributed Systems and Communication Protocols, and Protocol Specification, Testing, and Verification (FORTE/PSTV'99), Beijing, China*, pages 185–202. Kluwer Academic Publishers, October 1999.
- [53] Hubert Garavel and Damien Thivolle. Verification of GALS Systems by Combining Synchronous Languages and Process Calculi. In Corina Pasareanu, editor, *Model Checking Software – Proceedings of the 16th International SPIN Workshop on Model Checking of Software (SPIN'09), Grenoble, France*, volume 5578 of *Lecture Notes in Computer Science*, pages 241–260. Springer, June 2009.
- [54] Alexander Graf-Brill, Holger Hermanns, and Hubert Garavel. A Model-based Certification Framework for the EnergyBus Standard. In Erika Abraham and Catuscia Palamidessi, editors, *Proceedings of the 34th IFIP International Conference on Formal Techniques for Distributed Objects, Components and Systems (FORTE'15), Berlin, Germany*, volume 8461 of *Lecture Notes in Computer Science*, pages 84–99. Springer, June 2014.
- [55] Dennis Guck, Jip Spel, and Marielle Stoelinga. DFTCalc: Reliability Centered Maintenance via Fault Tree Analysis. In Michael Butler, Sylvain Conchon, and Fatiha Zaïdi, editors, *Proceedings of the 17th International Conference on Formal Engineering Methods (ICFEM'15), Paris, France*, volume 9407 of *Lecture Notes in Computer Science*, pages 304–311. Springer, 2015.
- [56] Matthias Güdemann, Gwen Salaün, and Meriem Ouederni. Counterexample Guided Synthesis of Monitors for Realizability Enforcement. In Supratik Chakraborty and Madhavan Mukund, editors, *Proceedings of the 10th International Symposium on Automated Technology for Verification and Analysis (ATVA'12), Thiruvananthapuram, India*, volume 7561 of *Lecture Notes in Computer Science*, pages 238–253. Springer, 2012.
- [57] Gabriel Huecas, Luis Llana-Díaz, Juan Quemada, Tomás Robles, and Alberto Verdejo. Process Calculi: E-LOTOS. In Howard Bowman and John Derrick, editors, *Formal Methods for Distributed Processing: A Survey of Object-Oriented Approaches*, pages 77–104. Cambridge University Press, 2001.
- [58] Gabriel Huecas, Luis Llana-Díaz, Tomás Robles, and Alberto Verdejo. E-LOTOS: An Overview. In Marco Ajmone Marsan, Juan Quemada, Tomás Robles, and Manuel Silva, editors, *Proceedings of the Workshop on*

- Formal Methods and Telecommunications (WFMT'99)*, Zaragoza, Spain, pages 94–102. Prensas Universitarias de Zaragoza, September 1999.
- [59] ISO/IEC. LOTOS – A Formal Description Technique Based on the Temporal Ordering of Observational Behaviour. Draft International Standard 8807, International Organization for Standardization – Information Processing Systems – Open Systems Interconnection, Geneva, July 1987.
- [60] ISO/IEC. ESTELLE – A Formal Description Technique Based on an Extended State Transition Model. International Standard 9074, International Organization for Standardization – Information Processing Systems – Open Systems Interconnection, Geneva, September 1988.
- [61] ISO/IEC. LOTOS – A Formal Description Technique Based on the Temporal Ordering of Observational Behaviour. International Standard 8807, International Organization for Standardization – Information Processing Systems – Open Systems Interconnection, Geneva, September 1989.
- [62] ISO/IEC. Enhancements to LOTOS (E-LOTOS). International Standard 15437:2001, International Organization for Standardization – Information Technology, Geneva, September 2001.
- [63] Fatma Jebali, Frédéric Lang, and Radu Mateescu. Formal Modelling and Verification of GALS systems using GRL and CADP. *Formal Aspects of Computing*, 28(5):767–804, 2016.
- [64] Martin Jourdan and Didier Parigot. Application Development with the FNC-2 Attribute Grammar System. In Dieter K. Hammer, editor, *Proceedings of the 3rd International Workshop on Compiler Construction (CC'90)*, Schwerin, Germany, volume 477 of *Lecture Notes in Computer Science*, pages 11–25. Springer, October 1990.
- [65] Abderahman Kriouile and Wendelin Serwe. Formal Analysis of the ACE Specification for Cache Coherent Systems-on-Chip. In Charles Pecheur and Michael Dierkes, editors, *Proceedings of the 18th International Workshop on Formal Methods for Industrial Critical Systems (FMICS'13)*, Madrid, Spain, volume 8187 of *Lecture Notes in Computer Science*, pages 108–122. Springer, September 2013.
- [66] Abderahman Kriouile and Wendelin Serwe. Using a Formal Model to Improve Verification of a Cache-Coherent System-on-Chip. In Christel Baier and Cesare Tinelli, editors, *Proceedings of the 21st International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS'15)*, London, United Kingdom, volume 9035 of *Lecture Notes in Computer Science*, pages 708–722. Springer, 2015.
- [67] Frédéric Lang. Compositional Verification using SVL Scripts. In Joost-Pieter Katoen and Perdita Stevens, editors, *Proceedings of the 8th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS'02)*, Grenoble, France, volume 2280 of *Lecture Notes in Computer Science*, pages 465–469. Springer, April 2002.

- [68] Frédéric Lang. EXP.OPEN 2.0: A Flexible Tool Integrating Partial Order, Compositional, and On-the-fly Verification Methods. In Judi Romijn, Graeme Smith, and Jaco van de Pol, editors, *Proceedings of the 5th International Conference on Integrated Formal Methods (IFM'05), Eindhoven, The Netherlands*, volume 3771 of *Lecture Notes in Computer Science*, pages 70–88. Springer, November 2005. Full version available as INRIA Research Report RR-5673.
- [69] Frédéric Lang, Gwen Salaün, Rémi Hérilier, Jeff Kramer, and Jeff Magee. Translating FSP into LOTOS and Networks of Automata. *Formal Aspects of Computing*, 22(6):681–711, November 2010.
- [70] Etienne Lantreibeq and Wendelin Serwe. Formal Analysis of a Hardware Dynamic Task Dispatcher with CADP. *Science of Computer Programming*, 80(Part A):130–149, February 2014.
- [71] Guy Leduc, Alan Jeffrey, and Mihaela Sighireanu. Introduction à E-LOTOS. In Ana Cavalli, editor, *Ingénierie des protocoles et qualité de service*, collection IC2, chapter 6, pages 213–253. Hermès, Paris, France, 2001.
- [72] Luc Léonard and Guy Leduc. An Introduction to ET-LOTOS for the Description of Time-Sensitive Systems. *Computer Networks and ISDN Systems*, 29(3):271–292, 1997.
- [73] Luc Léonard and Guy Leduc. A Formal Definition of Time in LOTOS. *Formal Aspects of Computing*, 10(3):248–266, 1998.
- [74] Xitong Li, Stuart Madnick, Hongwei Zhu, and Yushun Fan. Improving Data Quality for Web Services Composition. In *Proceedings of the 7th International Workshop on Quality in Databases (QDB'09), Lyon, France*, August 2009.
- [75] Francisco Isidro Massetto, Wanderley Lopes de Souza, and Sérgio Donizetti Zorzo. Simulator for E-LOTOS Specifications. In *Proceedings of the 35th Annual Simulation Symposium (SS'02), San Diego, California, USA*, pages 389–394. IEEE Computer Society, April 2002.
- [76] Radu Mateescu. A Generic Framework for Model Checking Software Architectures. In Juan Carlos Augusto and Ulrich Ultes-Nitsche, editors, *Proceedings of the 2nd International Workshop on Verification and Validation of Enterprise Information Systems (VVEIS'04), Porto, Portugal*. INSTICC Press, April 2004. Keynote presentation.
- [77] Radu Mateescu, Pedro Tiago Monteiro, Estelle Dumas, and Hidde de Jong. Computation Tree Regular Logic for Genetic Regulatory Networks. In Sung Deok Cha, Jin-Young Choi, Moonzoo Kim, Insup Lee, and Mahesh Viswanathan, editors, *Proceedings of the 6th International Symposium on Automated Technology for Verification and Analysis (ATVA'08), Seoul, South Korea*, volume 5311 of *Lecture Notes in Computer Science*, pages 48–63. Springer, October 2008. Full version available as INRIA Research Report RR-6521.

- [78] Radu Mateescu and Gwen Salaün. PIC2LNT: Model Transformation for Model Checking and Applied Pi-Calculus. In Nir Piterman and Scott A. Smolka, editors, *Proceedings of the 19th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS'13), Rome, Italy*, volume 7795 of *Lecture Notes in Computer Science*, pages 192–198. Springer, 2013.
- [79] Radu Mateescu and Wendelin Serwe. Model Checking and Performance Evaluation with CADP Illustrated on Shared-Memory Mutual Exclusion Protocols. *Science of Computer Programming*, 78(7):843–861, July 2013.
- [80] Radu Mateescu and Damien Thivolle. A Model Checking Language for Concurrent Value-Passing Systems. In Jorge Cuellar, Tom Maibaum, and Kaisa Sere, editors, *Proceedings of the 15th International Symposium on Formal Methods (FM'08), Turku, Finland*, volume 5014 of *Lecture Notes in Computer Science*, pages 148–164. Springer, May 2008.
- [81] G. J. Milne. CIRCAL and the Representation of Communication, Concurrency, and Time. *ACM Transactions on Programming Languages and Systems*, 7(2):270–298, April 1985.
- [82] Robin Milner. *A Calculus of Communicating Systems*, volume 92 of *Lecture Notes in Computer Science*. Springer, 1980.
- [83] Robin Milner. Calculi for Synchrony and Asynchrony. *Theoretical Computer Science*, 25:267–310, 1983.
- [84] Hana Mkaouar, Bechir Zalila, Jérôme Hugues, and Mohamed Jmaiel. From AADL Model to LNT Specification. In Juan A. de la Puente and Tullio Vardanega, editors, *Proceedings of the 20th ADA-Europe International Conference on Reliable Software Technologies (Ada-Europe'15), Madrid, Spain*, volume 9111 of *Lecture Notes in Computer Science*, pages 146–161. Springer, 2015.
- [85] Raquel Oliveira, Sophie Dupuy-Chessa, Gaelle Calvary, and Danielle Dadolle. Using Formal Models to Cross Check an Implementation. In Kris Luyten and Philippe Palanque, editors, *Proceedings of the 8th ACM SIGCHI Symposium on Engineering Interactive Computing Systems (EICS'16), Brussels, Belgium*, pages 126–137. ACM, June 2016.
- [86] Pascal Poizat and Gwen Salaün. Checking the Realizability of BPMN 2.0 Choreographies. In *Proceedings of the 27th Symposium On Applied Computing (SAC'12), Riva del Garda, Italy*. ACM Press, March 2012.
- [87] Pascal Poizat, Gwen Salaün, and Ajay Krishna. Checking Business Process Evolution. In Olga Kouchnarenko and Ramtin Khosravi, editors, *Revised Selected Papers of the 13th International Conference on Formal Aspects of Component Software (FACS'16), Besançon, France*, volume 10231 of *Lecture Notes in Computer Science*, pages 36–53. Springer, October 2016.
- [88] Olivier Ponsini, Caroline Fédèle, and Emmanuel Kounalis. Rewriting of Imperative Programs into Logical Equations. *Science of Computer Programming*, 56(3):363–401, May–June 2005.

- [89] Juan Quemada. E-LOTOS Has Born. E-mail announcement available from <ftp://ftp.inrialpes.fr/pub/vasy/publications/elotos/announce-97.txt>, February 1997.
- [90] Rudolf Roth, Jan de Meer, and Silke Storp. Data Specifications in Modular LOTOS. In T. Bolognesi, J. van de Lagemaat, and C. Vissers, editors, *LOTOSphere: Software Development with LOTOS*, pages 467–479. Kluwer Academic Publishers, 1995.
- [91] Wendelin Serwe. Formal Specification and Verification of Fully Asynchronous Implementations of the Data Encryption Standard. In Rob van Glabbeek, Jan Friso Groote, and Peter Höfner, editors, *Proceedings of the International Workshop on Models for Formal Analysis of Real Systems (MARS'15), Suva, Fiji*, volume 196 of *Electronic Proceedings in Theoretical Computer Science*. Open Publishing Association, 2015.
- [92] Carron Shankland and Alberto Verdejo. Time, E-LOTOS, and the FireWire. In Marco Ajmone Marsan, Juan Quemada, Tomás Robles, and Manuel Silva, editors, *Proceedings of the Workshop on Formal Methods and Telecommunications (WFMT'99), Zaragoza, Spain*, pages 103–119. Prensas Universitarias de Zaragoza, September 1999.
- [93] Carron Shankland and Alberto Verdejo. A Case Study in Abstraction Using E-LOTOS and the FireWire. *Computer Networks*, 37(3/4):481–502, 2001.
- [94] Mihaela Sighireanu. *Contribution à la définition et à l'implémentation du langage "Extended LOTOS"*. Thèse de Doctorat, Université Joseph Fourier (Grenoble), January 1999.
- [95] Mihaela Sighireanu, Alban Catry, David Champelovier, Hubert Garavel, Frédéric Lang, Guillaume Schaeffer, Wendelin Serwe, and Jan Stoecker. LOTOS NT User's Manual (Version 2.8). INRIA/CONVECS, Grenoble, France, <ftp://ftp.inrialpes.fr/pub/vasy/traian/manual.pdf>, 109 pages, February 2016.
- [96] Mihaela Sighireanu and Kenneth J. Turner. Requirement Capture, Formal Description and Verification of an Invoicing System. Research Report RR-3575, INRIA, Grenoble, December 1998.
- [97] Jan Stoecker, Frédéric Lang, and Hubert Garavel. Parallel Processes with Real-Time and Data: The ATLANTIF Intermediate Format. In Michael Leuschel and Heike Wehrheim, editors, *Proceedings of the 7th International Conference on Integrated Formal Methods (IFM'09), Düsseldorf, Germany*, volume 5423 of *Lecture Notes in Computer Science*, pages 88–102. Springer, February 2009. Full version available as INRIA Research Report RR-6950.
- [98] Damien Thivolle. *Langages modernes pour la vérification des systèmes asynchrones*. Thèse de Doctorat, Université Joseph Fourier (Grenoble, France) and Universitatea Politehnica din Bucuresti (Bucharest, Romania), April 2011.

-
- [99] Kenneth J. Turner and Mihaela Sighireanu. ELOTOS: (Enhanced) Language Of Temporal Ordering Specification. In Marc Frappier and Henri Habrias, editors, *Software Specification Methods: An Overview Using a Case Study*, pages 166–190. Springer, 2001.
 - [100] Dimitrios Vekris, Frédéric Lang, Catalin Dima, and Radu Mateescu. Verification of EB3 Specifications Using CADP. *Formal Aspects of Computing*, 28(1):145–178, 2016.
 - [101] Alberto Verdejo. E-LOTOS: Tutorial and Semantics. Master’s thesis, Departamento de Sistemas Informáticos y Programación, Universidad Complutense de Madrid, Spain, June 1999.
 - [102] Hao Wu, Xiaoxiao Yang, and Joost-Pieter Katoen. Performance Evaluation of Concurrent Data Structures. In Martin Fränzle, Deepak Kapur, and Naijun Zhan, editors, *Proceedings of the Symposium on Dependable Software Engineering (SETTA’16), Beijing, China*, volume 9984 of *Lecture Notes in Computer Science*, pages 38–49. Springer, 2016.
 - [103] Zhen Zhang, Wendelin Serwe, Jian Wu, Tomohiro Yoneda Hao Zheng, and Chris Myers. An Improved Fault-Tolerant Routing Algorithm for a Network-on-Chip Derived with Formal Analysis. *Science of Computer Programming*, 118:24–39, 2016.