

RESEARCH CENTRE

Inria Lyon Center

IN PARTNERSHIP WITH:

Université Claude Bernard (Lyon 1),
Ecole normale supérieure de Lyon, CNRS

2022

ACTIVITY REPORT

Project-Team

CASH

Compilation and Analyses for Software and Hardware

IN COLLABORATION WITH: Laboratoire de l'Informatique du
Parallélisme (LIP)

DOMAIN

**Algorithmics, Programming, Software
and Architecture**

THEME

Architecture, Languages and Compilation

The Inria logo is a stylized, cursive script in red, positioned in the bottom right corner of the page.

Contents

Project-Team CASH	1
1 Team members, visitors, external collaborators	2
2 Overall objectives	2
3 Research program	3
3.1 Research direction 1: Parallel and Dataflow Programming Models	3
3.1.1 Expected Impact	5
3.1.2 Scientific Program	5
3.2 Research direction 2: Expressive, Scalable and Certified Static Analyses	5
3.2.1 Expected impact	6
3.2.2 Scientific Program	6
3.3 Research direction 3: Optimizing Program Transformations	7
3.3.1 Expected impact	9
3.3.2 Scientific Program	9
3.4 Research direction 4: Simulation and Hardware	10
3.4.1 Expected Impact	11
3.4.2 Scientific Program	11
4 Application domains	12
5 Social and environmental responsibility	12
5.1 Footprint of research activities	12
5.2 Impact of research results	13
6 Highlights of the year	13
6.1 Articles	13
7 New software and platforms	13
7.1 New software	13
7.1.1 DCC	13
7.1.2 PoCo	14
7.1.3 Encore with dataflow explicit futures	14
7.1.4 fkcc	14
7.1.5 Vellvm	15
7.1.6 vaphor	15
7.1.7 Data Abstraction	15
7.1.8 S4BXI	15
7.1.9 llvm-pass	16
7.1.10 ribbit	16
7.1.11 calv	16
7.1.12 adtr	16
7.1.13 dowsing	17
7.1.14 odoc	17
7.1.15 ERCtool	17
7.1.16 PoLA	18
7.1.17 Actors-OCaml	18
7.1.18 ctrees	18

8	New results	18
8.1	Research direction 1: Parallel and Dataflow Programming Models	18
8.1.1	Flexible Synchronization for Parallel Computations.	18
8.1.2	An Optimised Flow for Futures: From Theory to Practice.	19
8.1.3	Locally abstract globally concrete semantics	19
8.1.4	Deterministic parallel programs	19
8.1.5	PNets: Parametrized networks of automata	20
8.1.6	A Survey on Verified Reconfiguration	20
8.1.7	A Survey on Parallelism and Determinacy	20
8.1.8	Verified Compilation Infrastructure for Concurrent Programs	21
8.1.9	Heterogeneous monadic computations in Coq	21
8.1.10	Actors and algebraic effects	21
8.2	Research direction 2: Expressive, Scalable and Certified Analyses	22
8.2.1	Verification of electric properties on transistor-level descriptions of circuits, using formal methods	22
8.2.2	Data Abstraction: A General Framework to Handle Program Verification of Data Structures	22
8.2.3	Search functions by types	22
8.2.4	A new module system for OCaml	23
8.2.5	Formalising Futures and Promises in Viper	23
8.3	Research direction 3: Optimizing Program Transformations	23
8.3.1	Optimizing compilation for tree-like structures	23
8.3.2	Parallel rewriting operations on Trees	24
8.3.3	Memory optimizations for Algebraic Data Types	24
8.3.4	Vellvm: Verified LLVM	24
8.3.5	Verified Abstract Interpreters as Monadic Interpreters	25
8.3.6	Lightweight Array Contraction by Trace-Based Polyhedral Analysis	25
8.3.7	A Polyhedral Approach for Scalar Promotion	26
8.3.8	Affine Multibanking for High-Level Synthesis	26
8.3.9	Scalable Parallelization for High-Level Synthesis	26
8.3.10	Compiler Support for Dynamic Task Scheduling	27
8.3.11	Dynamic Analysis for Sparse Kernels	27
8.4	Research direction 4: Simulation and Hardware	27
8.4.1	S4BXI: the MPI-ready Portals 4 Simulator	27
9	Bilateral contracts and grants with industry	27
9.1	CIFRE Ph.D of Julien Emmanuel with Bull/Atos, hosted by Inria. 2020-2023.	27
9.2	Partnership with the Aniah startup on circuit verification	28
9.3	CAVOC Project with Inria/Nomadic Labs	28
10	Partnerships and cooperations	28
10.1	International initiatives	28
10.1.1	Inria Exploratory Actions	28
10.1.2	Inria associate team not involved in an IIL or an international program	29
10.2	International research visitors	29
10.2.1	Visits of international scientists	29
10.2.2	Visits to international teams	30
10.3	European initiatives	30
10.3.1	Other european programs/initiatives	30
11	Dissemination	30
11.1	Promoting scientific activities	30
11.1.1	Scientific events: organisation	30
11.1.2	Scientific events: selection	30
11.1.3	Journal	30

11.1.4	Invited talks	31
11.1.5	Leadership within the scientific community	31
11.1.6	Scientific expertise	31
11.1.7	Research administration	31
11.2	Teaching - Supervision - Juries	31
11.2.1	Teaching	31
11.2.2	Supervision	32
11.2.3	Defended Ph.D	32
11.2.4	Juries	33
11.3	Popularization	33
11.3.1	Internal or external Inria responsibilities	33
11.3.2	Education	33
12	Scientific production	33
12.1	Major publications	33
12.2	Publications of the year	33
12.3	Other	35
12.4	Cited publications	35

Project-Team CASH

Creation of the Project-Team: 2019 June 01

Keywords

Computer sciences and digital sciences

- A2.1. – Programming Languages
 - A2.1.1. – Semantics of programming languages
 - A2.1.2. – Imperative programming
 - A2.1.4. – Functional programming
 - A2.1.6. – Concurrent programming
 - A2.1.7. – Distributed programming
 - A2.1.10. – Domain-specific languages
 - A2.1.11. – Proof languages
- A2.2. – Compilation
 - A2.2.1. – Static analysis
 - A2.2.2. – Memory models
 - A2.2.3. – Memory management
 - A2.2.4. – Parallel architectures
 - A2.2.5. – Run-time systems
 - A2.2.6. – GPGPU, FPGA...
 - A2.2.8. – Code generation
- A2.3.1. – Embedded systems
- A2.4. – Formal method for verification, reliability, certification
 - A2.4.1. – Analysis
 - A2.4.2. – Model-checking
 - A2.4.3. – Proofs
- A2.5.3. – Empirical Software Engineering
- A2.5.4. – Software Maintenance & Evolution
- A7.2.1. – Decision procedures
- A7.2.3. – Interactive Theorem Proving

Other research topics and application domains

- B9.5.1. – Computer science

1 Team members, visitors, external collaborators

Research Scientists

- Christophe Alias [INRIA, Researcher, HDR]
- Ludovic Henrio [CNRS, Researcher, HDR]
- Gabriel Radanne [INRIA, ISFP]
- Yannick Zakowski [INRIA, Researcher]

Faculty Members

- Matthieu Moy [Team leader, UNIV LYON I, Associate Professor, HDR]
- Laure Gonnord [GRENOBLE INP, Professor, HDR]

Post-Doctoral Fellows

- Bruno Ferres [ENS DE LYON, from Mar 2022 until Aug 2022]
- Bruno Ferres [INRIA, from Sep 2022]

PhD Students

- Thaïs Baudon [ENS DE LYON]
- Nicolas Chappe [ENS DE LYON]
- Julien Emmanuel [BULL]
- Amaury Maille [UNIV LYON I, until Sep 2022]
- Oussama Oulkaid [ANIAH, CIFRE, from Oct 2022]
- Hugo Thievenaz [INRIA]

Administrative Assistant

- Solene Audoux [INRIA]

2 Overall objectives

Research objectives. The overall objective of the CASH team is to take advantage of the characteristics of the specific hardware (generic hardware, hardware accelerators, or reconfigurable chips) to *compile energy efficient software and hardware*. To reach this goal, the CASH team provides efficient analyses and optimizing compilation frameworks for dataflow programming models. These contributions are combined with two other research directions. First, research on foundations of programming language and program analysis provides a theoretical basis for our work. Second, parallel and scalable simulation of hardware systems, combined with high-level synthesis tools, result in an end-to-end workflow for circuit design.

The scientific focus of CASH is on compute kernels and assembly of kernels, and the first goal is to improve their efficient compilation. However the team also works in collaboration with application developers, to understand better the overall need in HPC and design optimizations that are effective in the context of the targeted applications. Small computation kernels (tens of lines of code) that can be analyzed and optimized aggressively, medium-size kernels (hundreds of lines of code) that require modular analysis, and assembly of compute kernels (either as classical imperative programs or written directly in a dataflow language).

Our objective is to allow developers to design their own kernels, and benefit from good performance in terms of speed and energy efficiency without having to deal with fine-grained optimizations by hand. Consequently, our objective is first to improve the performance and energy consumption for HPC applications, while providing programming tools that can be used by developers and are at a convenient level of abstraction.

Obviously, large applications are not limited to assembly of compute kernels. Our languages and formalism definitions and analyses must also be able to deal with general programs. Our targets also include generalist programs with complex behaviors such as recursive programs operating on arrays, lists and trees; worklist algorithms (we often use the polyhedral model, a powerful theory to optimize loop nests, but it does not support data structures such as lists). Analysis on these programs should be able to detect non licit memory accesses, memory consumption, hotspots, . . . , and to prove functional properties.

Our Approach and methodology. We target a balance between theory and practice: problems extracted from industrial requirements often yield theoretical problems.

On the practical side, the CASH team targets applied research, in the sense that most research topics are driven by actual needs, discussed either through industrial partnership or extracted from available benchmarks.

The theoretical aspects ensure the coherency and the correctness of our approach. We rely on a *precise definition of the manipulated languages and their semantics*. The formalization of the different representations of the programs and of the analyses allow us to show that these different tasks will be performed with the same understanding of the program semantics.

Our approach is to cross-fertilize between several communities. For example, the abstract interpretation community provides a sound theoretical framework and very powerful analysis, but these are rarely applied in the context of optimizing compilation. Similarly, the hardware simulation community usually considers compilers as black-boxes and does not interact with researchers in compilation.

While a global approach links CASH activities and members, we do not plan to have a single unified toolchain where all contributions would be implemented. For example, contributions in the domain of static analysis of sequential programs may be implemented in the LLVM tool, results on dataflow models are applied both in the SigmaC compiler and in the DCC HLS tool, . . . This also implies that different activities of CASH target different application domains and potential end-users.

Research directions. The main objectives of the cash team are to provide *scalable and expressive static analysis* and *optimizing parallel compilers*. These directions rely on programming languages and representation of programs in which *parallelism and dataflow* play a crucial role. A central research direction aims at the study of parallelism and dataflow aspects in programming languages, both from a practical perspective (syntax or structure), and from a theoretical point of view (semantics). The CASH team also has *simulation activities* that are both applied internally in CASH, to simulate intermediate representations, and for embedded systems.

3 Research program

3.1 Research direction 1: Parallel and Dataflow Programming Models

In the last decades, several frameworks have emerged to design efficient compiler algorithms. The efficiency of all the optimizations performed in compilers strongly relies on effective *static analyses* and *intermediate representations*. Dataflow models are a natural intermediate representation for hardware compilers (HLS) and more generally for parallelizing compilers. Indeed, dataflow models capture task-level parallelism and can be mapped naturally to parallel architectures. In a way, a dataflow model is a partition of the computation into processes and a partition of the flow dependences into channels. This partitioning prepares resource allocation (which processor/hardware to use) and medium-grain communications.

The main goal of the CASH team is to provide efficient analyses and the optimizing compilation frameworks for dataflow programming models. The results of the team relies on programming languages

and representation of programs in which parallelism and dataflow play a crucial role. This first research direction aims at defining these dataflow languages and intermediate representations, both from a practical perspective (syntax or structure), and from a theoretical point of view (semantics). This first research direction thus defines the models on which the other directions will rely. It is important to note that we do not restrict ourselves to a strict definition of dataflow languages: more generally, we are interested in the parallel languages in which dataflow synchronization plays a significant role.

Intermediate dataflow model. The intermediate dataflow model is a representation of the program that is adapted for optimization and scheduling. It is obtained from the analysis of a (parallel or sequential) program and should at some point be used for compilation. The dataflow model must specify precisely its semantics and parallelism granularity. It must also be analyzable with polyhedral techniques, where powerful concepts exist to design compiler analysis, e.g., scheduling or resource allocation. Polyhedral Process Networks [0] extended with a module system could be a good starting point. But then, how to fit non-polyhedral parts of the program? A solution is to hide non-polyhedral parts into processes with a proper polyhedral abstraction. This organization between polyhedral and non-polyhedral processes will be a key aspect of our medium-grain dataflow model. The design of our intermediate dataflow model and the precise definition of its semantics will constitute a reliable basis to formally define and ensure the correctness of algorithms proposed by CASH: compilation, optimizations and analyses.

Dataflow programming languages. Dataflow paradigm has also been explored quite intensively in programming languages. Indeed, there exists a large panel of dataflow languages, whose characteristics differ notably, the major point of variability being the scheduling of agents and their communications. There is indeed a continuum from the synchronous dataflow languages like Lustre [0] or Streamit [0], where the scheduling is fully static, and general communicating networks like KPNs [0] or RVC-Cal [0] where a dedicated runtime is responsible for scheduling tasks dynamically, when they *can* be executed. These languages share some similarities with actor languages that go even further in the decoupling of processes by considering them as independent reactive entities. Another objective of the CASH team is to study dataflow programming *languages*, their semantics, their expressiveness, and their compilation. The specificity of the CASH team is that these languages will be designed taking into consideration the compilation using polyhedral techniques. In particular, we will explore which dataflow constructs are better adapted for our static analysis, compilation, and scheduling techniques. In practice we want to propose high-level primitives to express data dependency, this way the programmer can express parallelism in a dataflow way instead of the classical communication-oriented dependencies. The higher-level more declarative point of view makes programming easier but also give more optimization opportunities. These primitives will be inspired by the existing works in the polyhedral model framework, as well as dataflow languages, but also in the actors and active object languages [0] that nowadays introduce more and more dataflow primitives to enable data-driven interactions between agents, particularly with *futures* [0, 0].

Formal semantics

Proving the correctness of an analysis or of a program transformation requires a formal semantics of the language considered. Depending on the context, our formalizations may take the form of paper definitions, or of a mechanization inside of a proof assistant. While more time consuming, the latter may ensure in the adequate context some additional trust in the proofs established, as well as a tighter connection to an executable artifact. We have been recently studying in particular the formalization of concurrent and parallel paradigms, under weak memory models notably, by building on top of the interaction tree [0] approach developed for the Coq proof assistant.

Programming models and program transformations.

So far, the programming models designed in this direction allow to express parallelism in novel ways, but don't leverage the optimising compiler transformation introduced in direction 3. Indeed, optimising compilers only provide control over their behavior through extra-language annotations called "pragmas". Since those annotations are outside the language, they do not benefit from abstraction and modularity, and are often brittle. We plan to provide better integration between the optimisation passes of compiler inside the language itself through the use of meta-programming, by presenting optimisations as first class objects which can be applied, composed and manipulated in the language. A first step of this long term project is to investigate how to express loop transformations (developed by polyhedral model approaches) using existing meta-programming constructs.

3.1.1 Expected Impact

The impact of this research direction is both the usability of our representation for static analyses and optimizations performed in Sections 3.2 and 3.3, and the usability of its semantics to prove the correctness of these analyses.

3.1.2 Scientific Program

Medium-term activities. We plan to extend the existing results to widen the expressiveness of our intermediate representation and design new parallelism constructs. We will also work on the semantics of dataflow languages:

- Propose new stream programming models and a clean semantics where all kinds of parallelisms are expressed explicitly, and where all activities from code design to compilation and scheduling can be clearly expressed.
- Identify a core language that is rich enough to be representative of the dataflow languages we are interested in, but abstract and small enough to enable formal reasoning and proofs of correctness for our analyses and optimizations.

Long-term activities. In a longer-term vision, the work on semantics, while remaining driven by the applications, would lead to more mature results, for instance:

- Design more expressive dataflow languages and intermediate representations which would at the same time be expressive enough to capture all the features we want for aggressive HPC optimizations, and sufficiently restrictive to be (at least partially) statically analyzable at a reasonable cost.
- Define a module system for our medium-grain dataflow language. A program will then be divided into modules that can follow different compilation schemes and execution models but still communicate together. This will allow us to encapsulate a program that does not fit the polyhedral model into a polyhedral one and vice versa. Also, this will allow a compositional analysis and compilation, as opposed to global analysis which is limited in scalability.

3.2 Research direction 2: Expressive, Scalable and Certified Static Analyses

The design and implementation of efficient compilers becomes more difficult each day, as they need to bridge the gap between *complex languages* and *complex architectures*. Application developers use languages that bring them close to the problem that they need to solve which explains the importance of high-level programming languages. However, high-level programming languages tend to become more distant from the hardware which they are meant to command.

In this research direction, we propose to design expressive and scalable static analyses for compilers. This topic is closely linked to Sections 3.1 and 3.3 since the design of an efficient intermediate representation is made while regarding the analyses it enables. The intermediate representation should be expressive enough to embed maximal information; however if the representation is too complex the design of scalable analyses will be harder.

The analyses we plan to design in this activity will of course be mainly driven by the HPC dataflow optimizations we mentioned in the preceding sections; however we will also target other kinds of analyses applicable to more general purpose programs. We will thus consider two main directions:

- Extend the applicability of the polyhedral model, in order to deal with HPC applications that do not fit totally in this category. More specifically, we plan to work on more complex control and also on complex data structures, like sparse matrices, which are heavily used in HPC.
- Design of specialized static analyses for memory diagnostic and optimization inside general purpose compilers.

For both activities, we plan to cross fertilize ideas coming from the abstract interpretation community as well as language design, dataflow semantics, and WCET estimation techniques.

Correct by construction analyses. The design of well-defined semantics for the chosen programming language and intermediate representation will allow us to show the correctness of our analyses. The precise study of the semantics of Section 3.1 will allow us to adapt the analysis to the characteristics of the language, and prove that such an adaptation is well founded. This approach will be applicable both on the source language and on the intermediate representation.

We are interested both in paper proofs and verified proofs using a proof assistant such as Coq. Formally verified analysis crucially rely on a formal semantics of the programming language the analysis operates on: Yannick Zakowski precisely developed recently a new formal semantics in Coq for the sequential fragment of LLVM IR [0], the intermediate representation at the heart of the LLVM compilation infrastructure.

The semantics of Vellvm, which technically relies on Interaction Trees [0], enjoys crucial properties of compositionality and modularity. By leveraging these meta-theoretic properties of the semantics of the language, we believe that the additional objective of formal correctness can be compatible with the objectives of expressivity and scalability of the analyses we wish to develop for LLVM in particular.

The design of formal semantics allows formulating well-foundedness criteria relatively to the language semantics, that we can use to design our analyses, and then to study which extensions of the languages can be envisioned and analyzed safely, and which extensions (if any) are difficult to analyze and should be avoided. Here the correct identification of a core language for our formal studies (see Section 3.1) will play a crucial role as the core language should feature all the characteristics that might make the analysis difficult or incorrect.

Scalable abstract domains. We already have experience in designing low-cost semi relational abstract domains for pointers [0, 0], as well as tailoring static analyses for specialized applications in compilation [0, 0], Synchronous Dataflow scheduling [0], and extending the polyhedral model to irregular applications [0]. We also have experience in the design of various static verification techniques adapted to different programming paradigms.

Modularity of programming languages Modularity is an essential property of modern programming languages, allowing to assemble pieces of software in a high level and composable fashion. We aim to develop new module systems and tools for large scale ecosystems. A first aspect of this work is to pursue the collaboration with Didier Remy (Inria Cambium) and Jacques Garrigue (University of Nagoya) on designing module systems for ML languages. Gabriel Radanne is working on the formalization and implementation of a new rich module system which can serve as foundation for further experiment on the OCaml module system. A second aspect is to improve the ease of use of large ecosystems. We also work on tools to assist software developers, such as a tool to search functions by types, in a way that scales to complete ecosystems.

3.2.1 Expected impact

The impact of this work is the significantly widened applicability of various tools/compilers related to parallelization: allow optimizations for a larger class of programs, and allow low-cost analysis that scale to very large programs.

We target both analysis for optimization and analysis to detect, or prove the absence of bugs.

3.2.2 Scientific Program

Medium-term activities. In the context of Paul Iannetta's Phd thesis, we have proposed a *semantic rephrasing* of the polyhedral model and proposed first steps toward an effective "polyhedral-like compilation" for algebraic datastructures like trees. In medium term, we want to extend the applicability of this new model for arbitrary layouts. The most challenging ones are sparse matrices. This activity still relies on a formalization of the optimization activities (dependency computation, scheduling, compilation) in a more general Abstract-Interpretation based framework in order to make the approximations explicit.

At the same time, we plan to continue to work on scaling static analyses for general purpose programs, in the spirit of Maroua Maalej's PhD [0], whose contribution is a sequence of memory analyses inside production compilers. We already began a collaboration with Caroline Collange (PACAP team of IRISA

Laboratory) on the design of static analyses to optimize copies from the global memory of a GPU to the block kernels (to increase locality). In particular, we have the objective to design specialized analyses but with an explicit notion of cost/precision compromise, in the spirit of the paper [0] that tries to formalize the cost/precision compromise of interprocedural analyses with respect to a “context sensitivity parameter”.

Long-term activities. In a longer-term vision, the work on scalable static analyses, whether or not directed from the dataflow activities, will be pursued in the direction of large general-purpose programs.

An ambitious challenge is to find a generic way of adapting existing (relational) abstract domains within the Single Static Information [0] framework so as to improve their scalability. With this framework, we would be able to design static analyses, in the spirit of the seminal paper [0] which gave a theoretical scheme for classical abstract interpretation analyses.

We also plan to work on the interface between the analyses and their optimization clients inside production compilers.

3.3 Research direction 3: Optimizing Program Transformations

In this part, we propose to design the compiler analyses and optimizations for the *medium-grain* dataflow model defined in section 3.1. We also propose to exploit these techniques to improve the compilation of dataflow languages based on actors. Hence our activity is split into the following parts:

- Translating a sequential program into a medium-grain dataflow model. The programmer cannot be expected to rewrite the legacy HPC code, which is usually relatively large. Hence, compiler techniques must be invented to do the translation.
- Transforming and scheduling our medium-grain dataflow model to meet some classic optimization criteria, such as throughput, local memory requirements, or I/O traffic.
- Combining agents and polyhedral kernels in dataflow languages. We propose to apply the techniques above to optimize the processes in actor-based dataflow languages and combine them with the parallelism existing in the languages.

We plan to rely extensively on the polyhedral model to define our compiler analysis. The polyhedral model was originally designed to analyze imperative programs. Analysis (such as scheduling or buffer allocation) must be redefined in light of dataflow semantics.

Translating a sequential program into a medium-grain dataflow model. The programs considered are compute-intensive parts from HPC applications, typically big HPC kernels of several hundreds of lines of C code. In particular, we expect to analyze the process code (actors) from the dataflow programs. On short ACL (Affine Control Loop) programs, direct solutions exist [0] and rely directly on array dataflow analysis [0]. On bigger ACL programs, this analysis no longer scales. We plan to address this issue by *modularizing* array dataflow analysis. Indeed, by splitting the program into processes, the complexity is mechanically reduced. This is a general observation, which was exploited in the past to compute schedules [0]. When the program is no longer ACL, a clear distinction must be made between polyhedral parts and non-polyhedral parts. Hence, our medium-grain dataflow language must distinguish between polyhedral process networks, and non-polyhedral code fragments. This structure raises new challenges: How to abstract away non-polyhedral parts while keeping the polyhedrality of the dataflow program? Which trade-off(s) between precision and scalability are effective?

Medium-grain data transfers minimization. When the system consists of a single computing unit connected to a slow memory, the roofline model [0] defines the optimal ratio of computation per data transfer (*operational intensity*). The operational intensity is then translated to a partition of the computation (loop tiling) into *reuse units*: inside a reuse unit, data are transferred locally; between reuse units, data are transferred through the slow memory. On a *fine-grain* dataflow model, reuse units are exposed with loop tiling; this is the case for example in Data-aware Process Network (DPN) [0]. The following questions are however still open: How does that translate on *medium-grain* dataflow models? And fundamentally what does it mean to *tile* a dataflow model?

Combining agents and polyhedral kernels in dataflow languages. In addition to the approach developed above, we propose to explore the compilation of dataflow programming languages. In fact, among

the applications targeted by the project, some of them are already thought or specified as dataflow actors (video compression, machine-learning algorithms,...).

So far, parallelization techniques for such applications have focused on taking advantage of the decomposition into agents, potentially duplicating some agents to have several instances that work on different data items in parallel [0]. In the presence of big agents, the programmer is left with the splitting (or merging) of these agents by-hand if she wants to further parallelize her program (or at least give this opportunity to the runtime, which in general only sees agents as non-malleable entities). In the presence of arrays and loop-nests, or, more generally, some kind of regularity in the agent's code, however, we believe that the programmer would benefit from automatic parallelization techniques such as those proposed in the previous paragraphs. To achieve the goal of a totally integrated approach where programmers write the applications they have in mind (application flow in agents where the agents' code express potential parallelism), and then it is up to the system (compiler, runtime) to propose adequate optimizations, we propose to build on solid formal definition of the language semantics (thus the formal specification of parallelism occurring at the agent level) to provide hierarchical solutions to the problem of compilation and scheduling of such applications.

Certified compilation We will develop a research direction around the formal proof of compilation passes, and of optimizing program transformations in particular. Although realistic formally verified optimizing compilers are roughly 15 years old, three limitations to the current state of the art are apparent.

First, loop optimizations have been very sparsely tackled, their proof rising difficult semantic issues. We intend on one side to leverage the compositionality of Interaction-Tree-based semantics as used in Vellvm to improve the situation. An orthogonal axis we wish to explore is the formalization in Coq of the Polyhedral Model, as pioneered in 2021 by Courant and Leroy [0].

Second, parallelism and concurrency have been almost ignored by the verified compilation community. This problem is a major long term endeavor for we first need to develop the appropriate semantic tools. Ludovic Henrio and Yannick Zakowski will work with a master student, Ambre Suhamy, to explore the use of Interaction Trees to model various paradigms for concurrency, paving the long term way to an extension of Vellvm to concurrency.

Third, these proofs are very brittle for they rely on concrete implementation of memory models rather than axiomatizations of those. Ludovic Henrio and Yannick Zakowski will work with a master student, Alban Reynaud, to develop semantic tools to reason formally up-to arbitrary algebras in Coq. One of the core objectives of this project is to prove optimizations at a higher level of abstraction, so that these proofs remain valid by construction under changes in the memory model.

The compiler analyses proposed above do not target a specific platform. In this part, we propose to leverage these analysis to develop source-level optimizations for high-level synthesis (HLS).

High-level synthesis consists in compiling a kernel written in a high-level language (typically in C) into a circuit. As for any compiler, an HLS tool consists in a *front-end* which translates the input kernel into an *intermediate representation*. This intermediate representation captures the control/flow dependences between computation units, generally in a hierarchical fashion. Then, the *back-end* maps this intermediate representation to a circuit (e.g. FPGA configuration). We believe that HLS tools must be thought as fine-grain automatic parallelizers. In classic HLS tools, the parallelism is expressed and exploited at the back-end level during the scheduling and the resource allocation of arithmetic operations. We believe that it would be far more profitable to derive the parallelism at the front-end level.

Hence, CASH will focus on the *front-end* pass and the *intermediate representation*. Low-level *back-end* techniques are not in the scope of CASH. Specifically, CASH will leverage the dataflow representation developed in Section 3.1 and the compilation techniques developed in Section 3.3 to develop a relevant intermediate representation for HLS and the corresponding front-end compilation algorithms.

Our results will be evaluated by using existing HLS tools (e.g., Intel HLS compiler, Xilinx Vivado HLS). We will implement our compiler as a source-to-source transformation in front of HLS tools. With this approach, HLS tools are considered as a “back-end black box”. The CASH scheme is thus: (i) *front-end*: produce the CASH dataflow representation from the input C kernel. Then, (ii) turn this dataflow representation to a C program with pragmas for an HLS tool. This step must convey the characteristics of the dataflow representation found by step (i) (e.g. dataflow execution, fifo synchronisation, channel size). This source-to-source approach will allow us to get a full source-to-FPGA flow demonstrating the benefits of our tools while relying on existing tools for low-level optimizations. Step (i) will start from the DCC tool developed by Christophe Alias, which already produces a dataflow intermediate representation:

the Data-aware Process Networks (DPN) [0]. Hence, the very first step is then to choose an HLS tool and to investigate which input should be fed to the HLS tool so it “respects” the parallelism and the resource allocation suggested by the DPN. From this basis, we plan to investigate the points described thereafter.

Roofline model and dataflow-level resource evaluation. Operational intensity must be tuned according to the roofline model. The roofline model [0] must be redefined in light of FPGA constraints. Indeed, the peak performance is no longer constant: it depends on the operational intensity itself. The more operational intensity we need, the more local memory we use, the less parallelization we get (since FPGA resources are limited), and finally the less performance we get! Hence, multiple iterations may be needed before reaching an efficient implementation. To accelerate the design process, we propose to iterate at the dataflow program level, which implies a fast resource evaluation at the dataflow level.

Reducing FPGA resources. Each parallel unit must use as little resources as possible to maximize parallel duplication, hence the final performance. This requires to factorize the control and the channels. Both can be achieved with source-to-source optimizations at dataflow level. The main issue with outputs from polyhedral optimization is large piecewise affine functions that require a wide silicon surface on the FPGA to be computed. Actually we do not need to compute a closed form (expression that can be evaluated in bounded time on the FPGA) *statically*. We believe that the circuit can be compacted if we allow control parts to be evaluated dynamically. Finally, though dataflow architectures are a natural candidate, adjustments are required to fit FPGA constraints (2D circuit, few memory blocks). Ideas from systolic arrays [0] can be borrowed to re-use the same piece of data multiple times, despite the limitation to regular kernels and the lack of I/O flexibility. A trade-off must be found between pure dataflow and systolic communications.

Improving circuit throughput. Since we target streaming applications, the throughput must be optimized. To achieve such an optimization, we need to address the following questions. How to derive an optimal upper bound on the throughput for polyhedral process network? Which dataflow transformations should be performed to reach it? The limiting factors are well known: I/O (decoding of burst data), communications through addressable channels, and latencies of the arithmetic operators. Finally, it is also necessary to find the right methodology to measure the throughput statically and/or dynamically.

3.3.1 Expected impact

In general, splitting a program into simpler processes simplifies the problem. This observation leads to the following points:

- By abstracting away irregular parts in processes, we expect to structure the long-term problem of handling irregular applications in the polyhedral model. The long-term impact is to widen the applicability of the polyhedral model to irregular kernels.
- Splitting a program into processes reduces the problem size. Hence, it becomes possible to scale traditionally expensive polyhedral analysis such as scheduling or tiling to quote a few.

As for the third research direction, the short term impact is the possibility to combine efficiently classical dataflow programming with compiler polyhedral-based optimizations. We will first propose ad-hoc solutions coming from our HPC application expertise, but supported by strong theoretical results that prove their correctness and their applicability in practice. In the longer term, our work will allow specifying, designing, analyzing, and compiling HPC dataflow applications in a unified way. We target semi-automatic approaches where pertinent feedback is given to the developer during the development process.

3.3.2 Scientific Program

Short-term and ongoing activities. We plan to evaluate the impact of state-of-the-art polyhedral source-to-source transformations on HLS for FPGA. Our results on polyhedral HLS (DPN [0, 0]) could also be a good starting point for this purpose. We will give a particular focus to memory layout transformations, easier to implement as a source level transformation. Then, we will tackle control optimizations through the adaptation of loop tiling to HLS constraints.

Medium-term activities. The results of the preceding paragraph are partial and have been obtained with a simple experimental approach only using off-the-shelf tools. We are thus encouraged to pursue research on combining expertise from dataflow programming languages and polyhedral compilation. Our long term objective is to go towards a formal framework to express, compile, and run dataflow applications with intrinsic instruction or pipeline parallelism.

We plan to investigate in the following directions:

- Investigate how polyhedral analysis extends on modular dataflow programs. For instance, how to modularize polyhedral scheduling analysis on our dataflow programs?
- Develop a proof of concept and validate it on linear algebra kernels (SVD, Gram-Schmidt, etc.).
- Explore various areas of applications from classical dataflow examples, like radio and video processing, to more recent applications in deep learning algorithmic. This will enable us to identify some potential (intra and extra) agent optimization patterns that could be leveraged into new language idioms.

Also, we plan to explore how polyhedral transformations might *scale on larger applications*, typically those found in deep-learning algorithms. We will investigate how the regularity of polyhedral kernels can be exploited to infer general affine transformations from a few offline execution traces. This is the main goal of the *PolyTrace exploratory action*, started on 2021 in collaboration with Waseda University. We will first target offline memory allocation, an important transformation used in HLS and generally in automatic parallelization.

Finally, we plan to explore how on-the-fly evaluation can reduce the complexity of the control. A good starting point is the control required for the load process (which fetch data from the distant memory). If we want to avoid multiple load of the same data, the FSM (Finite State Machine) that describes it is usually very complex. We believe that dynamic construction of the load set (set of data to load from the main memory) will use less silicon than an FSM with large piecewise affine functions computed statically.

Long-term activities. Current work focus on purely polyhedral applications. Irregular parts are not handled. Also, a notion of tiling is required so the communications of the dataflow program with the outside world can be tuned with respect to the local memory size. Hence, we plan to investigate the following points:

- Assess simple polyhedral/non polyhedral partitioning: How non-polyhedral parts can be hidden in processes/channels? How to abstract the dataflow dependencies between processes? What would be the impact on analyses? We target programs with irregular control (e.g., while loop, early exits) and regular data (arrays with affine accesses).
- Design tiling schemes for modular dataflow programs: What does it mean to tile a dataflow program? Which compiler algorithms to use?
- Implement a mature compiler infrastructure from the front-end to code generation for a reasonable subset of the representation.

Also, we plan to systematize the definition of scalable polyhedral compilers using extrapolation from offline traces. Both theoretical and applied research are required to reach this goal. The research strategy consists in studying several instances (memory allocation, scheduling, etc). Then, in producing the theoretical ingredients to reach a general methodology of conception.

3.4 Research direction 4: Simulation and Hardware

Complex systems such as systems-on-a-chip or HPC computer with FPGA accelerator comprise both hardware and software parts, tightly coupled together. In particular, the software cannot be executed without the hardware, or at least a simulator of the hardware.

Because of the increasing complexity of both software and hardware, traditional simulation techniques (Register Transfer Level, RTL) are too slow to allow full system simulation in reasonable time. New techniques such as Transaction Level Modeling (TLM) [0] in SystemC [0] have been introduced

and widely adopted in the industry. Internally, SystemC uses discrete-event simulation, with efficient context-switch using cooperative scheduling. TLM abstracts away communication details, and allows modules to communicate using function calls. We are particularly interested in the loosely timed coding style where the timing of the platform is not modeled precisely, and which allows the fastest simulations. This allowed gaining several orders of magnitude of simulation speed. However, SystemC/TLM is also reaching its limits in terms of performance, in particular due to its lack of parallelism.

Work on SystemC/TLM parallel execution is both an application of other work on parallelism in the team and a tool complementary to HLS presented in Sections 3.1 (dataflow models and programs) and 3.3 (application to FPGA). Indeed, some of the parallelization techniques we develop in CASH could apply to SystemC/TLM programs. Conversely, a complete design-flow based on HLS needs fast system-level simulation: the full-system usually contains both hardware parts designed using HLS, handwritten hardware components, and software.

We also work on simulation of the DPN intermediate representation. Simulation is a very important tool to help validate and debug a complete compiler chain. Without simulation, validating the front-end of the compiler requires running the full back-end and checking the generated circuit. Simulation can avoid the execution time of the backend and provide better debugging tools.

Automatic parallelization has shown to be hard, if at all possible, on loosely timed models [0]. We focus on semi-automatic approaches where the programmer only needs to make minor modifications of programs to get significant speedups. We already obtained results in the joint PhD (with Tanguy Sassolas) of Gabriel Busnot with CEA-LIST. The research targets parallelizing SystemC heterogeneous simulations, extending SScale [0], which is very efficient to simulate parallel homogeneous platforms such as multi-core chips. We removed the need for manual address annotations, which did not work when the software does non-trivial memory management (virtual memory using a memory management unit, dynamic allocation), since the address ranges cannot be known statically. We can now parallelize simulation running with a full software stack including Linux.

We are also working with Bull/Atos on HPC interconnect simulation, using SimGrid [0]. Our goal is to allow simulating an application that normally runs on a large number of nodes on a single computer, and obtain relevant performance metrics.

3.4.1 Expected Impact

The short term impact is the possibility to improve simulation speed with a reasonable additional programming effort. The amount of additional programming effort will thus be evaluated in the short term.

In the longer term, our work will allow scaling up simulations both in terms of models and execution platforms. Models are needed not only for individual Systems on a Chip, but also for sets of systems communicating together (e.g., the full model for a car which comprises several systems communicating together), and/or heterogeneous models. In terms of execution platform, we are studying both parallel and distributed simulations.

3.4.2 Scientific Program

Medium-term activities. We started working on the “heterogeneous” aspect of simulations with an approach allowing changing the level of details in a simulation at runtime.

Several research teams have proposed different approaches to deal with parallelism and heterogeneity. Each approach targets a specific abstraction level and coding style. While we do not hope for a universal solution, we believe that a better coordination of different actors of the domain could lead to a better integration of solutions. We could imagine, for example, a platform with one subsystem accelerated with SScale [0] from CEA-LIST, some compute-intensive parts delegated to sc-during [0] from Matthieu Moy, and a co-simulation with external physical solvers using SystemC-MDVP [0] from LIP6. We plan to work on the convergence of approaches, ideally both through point-to-point collaborations and with a collaborative project.

A common issue with heterogeneous simulation is the level of abstraction. Physical models only simulate one scenario and require concrete input values, while TLM models are usually abstract and not aware of precise physical values. One option we would like to investigate is a way to deal with loose

information, e.g. manipulate intervals of possible values instead of individual, concrete values. This would allow a simulation to be symbolic with respect to the physical values.

Long-term activities. In the long term, our vision is a simulation framework that will allow combining several simulators (not necessarily all SystemC-based), and allow running them in a parallel way. The Functional Mockup Interface (FMI) standard is a good basis to build upon, but the standard does not allow expressing timing and functional constraints needed for a full co-simulation to run properly.

4 Application domains

The CASH team targets HPC programs, at different levels. Small computation kernels (tens of lines of code) that can be analyzed and optimized aggressively, medium-size kernels (hundreds of lines of code) that require modular analysis, and assembly of compute kernels (either as classical imperative programs or written directly in a dataflow language).

The work on various application domains and categories of programs is driven by the same idea: exploring various topics is a way to converge on unifying representations and algorithms even for specific applications. All these applications share the same research challenge: find a way to integrate computations, data, mapping, and scheduling in a common analysis and compilation framework.

Typical HPC kernels include linear solvers, stencils, matrix factorizations, BLAS kernels, etc. Many kernels can be found in the Polybench/C benchmark suite [0]. The irregular versions can be found in [0]. Numerical kernels used in quantitative finance [0] are also good candidates, e.g., finite difference and Monte-Carlo simulation.

The medium-size applications we target are streaming algorithms [0], scientific workflows [0], and also the now very rich domain of deep learning applications [0]. We explore the possibilities of writing (see Section 3.1) and compiling (see Section 3.3) applications using a dataflow language. As a first step, we will target dataflow programs written in SigmaC [0] for which the fine grain parallelism is not taken into account. In parallel, we will also study the problem of deriving relevant (with respect to safety or optimization) properties on dataflow programs with array iterators.

The approach of CASH is based on compilation, and our objective is to allow developers to design their own kernels, and benefit from good performance in terms of speed and energy efficiency without having to deal with fine-grained optimizations by hand. Consequently, our objective is first to improve the performance and energy consumption for HPC applications, while providing programming tools that can be used by developers and are at a convenient level of abstraction.

Obviously, large applications are not limited to assembly of compute kernels. Our languages and formalism definitions and analyses must also be able to deal with general programs. Our targets also include generalist programs with complex behaviors such as recursive programs operating on arrays, lists and trees; worklist algorithms (lists are not handled within the polyhedral domain). Analysis on these programs should be able to detect non licit memory accesses, memory consumption, hotspots, ..., and to prove functional properties.

The simulation activities are both applied internally in CASH, to simulate intermediate representations, and for embedded systems. We are interested in Transaction-Level Models (TLM) of Systems-on-a-Chip (SoCs) including processors and hardware accelerators. TLM provides an abstract but executable model of the chip, with enough details to run the embedded software. We are particularly interested in models written in a loosely timed coding style. We plan to extend these to heterogeneous simulations including a SystemC/TLM part to model the numerical part of the chip, and other simulators to model physical parts of the system.

5 Social and environmental responsibility

5.1 Footprint of research activities

Although we do not have a precise measure of our carbon (and other environmental) footprint, the two main sources of impact of computer-science research activities are usually transport (plane) and digital equipment (lifecycle of computers and other electronic devices).

Many members of the CASH team are already in an approach of reducing their international travel, and hopefully the new solutions we had to set up to continue our activities during the COVID crisis will allow us to continue our research with a sustainable amount of travel, and using other forms of remote collaborations when possible.

As far as digital equipment is concerned, we try to extend the lifetime of our machines as much as possible.

5.2 Impact of research results

Many aspects of our research are meant to provide tools to make programs more efficient, in particular more power-efficient. It is very hard, however, to assess the actual impact of such research. In many cases, improvements in power-efficiency lead to a rebound effect which may weaken the benefit of the improvement, or even lead to an increase in total consumption (backfire).

CASH provides tools for developers, but does not develop end-user applications. We believe the social impact of our research depends more on the way developers will use our tools than on the way we conduct our research. We do have a responsibility on the application domains we promote, though.

Ludovic Henrio followed the "Atelier Sciences Environnements Sociétés Inria 2021" (atelier Sens) organized by Eric Tannier in June 2021. Then, for the voluntary Cash members, he has animated an atelier Sens during the Cash seminar in October 2021.

6 Highlights of the year

6.1 Articles

- We published a survey on parallelism and determinacy in ACM Computing Survey [0]. It reviews the existing programming models that ensures some form of determinacy while allowing for parallel computations.
- Our works on choice trees, an extension of interaction trees for representing non-deterministic computations with effects in Coq, has been accepted at POPL 2023 [0].

7 New software and platforms

7.1 New software

7.1.1 DCC

Name: DPN C Compiler

Keywords: Polyhedral compilation, Automatic parallelization, High-level synthesis

Functional Description: Dcc (Data-aware process network C Compiler) compiles a regular C kernel to a data-aware process network (DPN), a dataflow intermediate representation suitable for high-level synthesis in the context of high-performance computing. Dcc has been registered at the APP ("Agence de protection des programmes") and transferred to the XtremLogic start-up under an Inria license.

News of the Year: This year, Dcc was enhanced with user-guided loop tiling. Given a user-specified tiling template, a correct loop tiling with minimal latency is inferred.

Publication: [hal-03143777](https://hal.archives-ouvertes.fr/hal-03143777)

Contact: Christophe Alias

Participants: Christophe Alias, Alexandru Plesco

7.1.2 PoCo

Name: Polyhedral Compilation Library

Keywords: Polyhedral compilation, Automatic parallelization

Functional Description: PoCo (Polyhedral Compilation Library) is framework to develop program analysis and optimizations in the polyhedral model. PoCo features polyhedral building blocks as well as state-of-the-art polyhedral program analysis. PoCo has been registered at the APP (“agence de protection des programmes”) and transferred to the XtremLogic start-up under an Inria licence.

News of the Year: This year, GLPK was interfaced to the symbolic engine. Also, the Farkas engine was improved to handle more complex affine constraints.

Contact: Christophe Alias

Participant: Christophe Alias

7.1.3 Encore with dataflow explicit futures

Keywords: Language, Optimizing compiler, Source-to-source compiler, Compilers

Functional Description: Fork of the Encore language compiler, with a new "Flow" construct implementing data-flow explicit futures.

URL: <https://gitlab.inria.fr/lhenrio/encorewithdatafuts>

Contact: Ludovic Henrio

7.1.4 fkcc

Name: The Farkas Calculator

Keywords: DSL, Farkas Lemma, Polyhedral compilation

Scientific Description: fkcc is a scripting tool to prototype program analyses and transformations exploiting the affine form of Farkas lemma. Our language is general enough to prototype in a few lines sophisticated termination and scheduling algorithms. The tool is freely available and may be tried online via a web interface. We believe that fkcc is the missing chain to accelerate the development of program analyses and transformations exploiting the affine form of Farkas lemma.

Functional Description: fkcc is a scripting tool to prototype program analyses and transformations exploiting the affine form of Farkas lemma. Our language is general enough to prototype in a few lines sophisticated termination and scheduling algorithms. The tool is freely available and may be tried online via a web interface. We believe that fkcc is the missing chain to accelerate the development of program analyses and transformations exploiting the affine form of Farkas lemma.

Release Contributions: - Script language - Polyhedral constructors - Farkas summation solver

News of the Year: This year, Christophe improved the lexico-optimization engine by using GLPK.

URL: <http://foobar.ens-lyon.fr/fkcc/>

Publication: hal-03106000

Contact: Christophe Alias

Participant: Christophe Alias

7.1.5 Vellvm

Keywords: Coq, Semantic, Compilation, Proof assistant, Proof

Scientific Description: A modern formalization in the Coq proof assistant of the sequential fragment of LLVM IR. The semantics, based on the Interaction Trees library, presents several rare properties for mechanized development of this scale: it is compositional, modular, and extracts to a certified executable interpreter. A rich equational theory of the language is provided, and several verified tools based on this semantics are in development.

Functional Description: Formalization in the Coq proof assistant of a subset of the LLVM compilation infrastructure.

URL: <https://github.com/vellvm/vellvm>

Contact: Yannick Zakowski

Participants: Yannick Zakowski, Steve Zdancewic, Calvin Beck, Irene Yoon

Partner: University of Pennsylvania

7.1.6 vaphor

Name: Verification of Programs with Horn Clauses

Keyword: Program verification

Functional Description: Program to horn clauses horn clauses with arrays abstraction

Contact: Laure Gonnord

Partner: Vérimag

7.1.7 Data Abstraction

Name: Data Abstraction

Keywords: Static analysis, Program verification, Propositional logic

Functional Description: The tool is an element of a static program (or other) verification process which is done in three steps:

1. Transform the verification problem into Horn clauses, perhaps using MiniJavaConverter or SeaHorn
2. Simplify the Horn clauses using data abstraction (this tool).
3. Solve the Horn clauses using a Horn solver such as Z3

Contact: Laure Gonnord

Partner: Vérimag

7.1.8 S4BXI

Keywords: Simulation, HPC, Network simulator

Functional Description: S4BXI is a simulator of the Portals4 network API. It is written using SimGrid's S4U interface, which provides a fast flow-model. More specifically, this simulator is tuned to model as best as possible Bull's hardware implementation of portals (BXI interconnect)

URL: <https://s4bxi.julien-emmanuel.com/docs/>

Contact: Julien Emmanuel

Partner: Bull - Atos Technologies

7.1.9 llvm-pass

Name: LLVM Pass analyzer

Keyword: Compilation

Functional Description: tool suite to work on llvm passes

URL: <https://github.com/llvmpass/llvm-passes>

Contact: Laure Gonnord

7.1.10 ribbit

Keywords: Compilation, Pattern matching, Algebraic Data Types

Functional Description: Ribbit is a compiler for pattern languages with algebraic data types which is parameterized by the memory representation of types. Given a memory representation, it generates efficient and correct code for pattern matching clauses.

URL: <https://gitlab.inria.fr/cash/ribbit>

Contact: Gabriel Radanne

7.1.11 calv

Name: AVL calculator

Keywords: Data structures, OpenMP

Functional Description: calv is a calculator which is used to run different implementations of AVL trees, and compare their relative performances.

URL: <https://gitlab.inria.fr/paiannet/calv>

Contact: Paul Iannetta

7.1.12 adtr

Name: ADT Rewriting language

Keywords: Compilation, Static typing, Algebraic Data Types, Term Rewriting Systems

Functional Description: ADTs are generally represented by nested pointers, for each constructors of the algebraic data type. Furthermore, they are generally manipulated persistently, by allocating new constructors.

ADTr allow representing ADTs in a flat way while compiling a pattern match-like construction as a rewrite on the memory representation. The goal is to then use this representation to optimize the rewriting and exploit parallelism.

URL: <https://github.com/Drup/adtr>

Publication: hal-03355377

Contact: Gabriel Radanne

Participants: Gabriel Radanne, Paul Iannetta, Laure Gonnord

7.1.13 dowsing

Keywords: Static typing, Ocaml

Functional Description: Dowsing is a tool to search function by types. Given a simple OCaml type, it will quickly find all functions whose types are compatible.

Dowsing works by building a database containing all the specified libraries. New libraries can be added to the database. It then builds an index which allow to quickly answer to requests.

URL: <https://github.com/Drup/dowsing/>

Publication: [hal-03355381](#)

Contact: Gabriel Radanne

Participants: Gabriel Radanne, Laure Gonnord

7.1.14 odoc

Keyword: Ocaml

Functional Description: OCaml is a statically typed programming language with wide-spread use in both academia and industry. Odoc is a tool to generate documentation of OCaml libraries, either as HTML websites for online distribution or to create PDF manuals and man pages.

News of the Year: This year, Gabriel Radanne rewrote a significant portion of odoc to provide improved HTML output, make it possible to produce other document formats, and introduce the ability to produce man pages. Florian Angeletti then implemented PDF output and integrated the usage of odoc in the official OCaml distribution. Concurrently, Jon Ludlam and Leo White rewrote the resolution mechanism of odoc. This all led to a joint presentation at the OCaml workshop and an upcoming new major release.

URL: <https://github.com/ocaml/odoc/>

Contact: Gabriel Radanne

Participants: Jon Ludlam, Gabriel Radanne, Florian Angeletti, Leo White

7.1.15 ERCtool

Name: Electrical Rule Checking Tool

Keywords: Verification, Model Checking, Electrical circuit, Transistor, Program verification, Formal methods

Functional Description: ERCtool is developed as part of a collaboration with the Aniah company, specialized in the verification of electric properties on circuits. A specificity of ERCtool is that it allows the analysis of a circuit with multiple power-supplies, and allows getting formal guarantees on the absence of error.

Contact: Matthieu Moy

Partner: Aniah

7.1.16 PoLA

Name: PoLA: a Polyhedral Liveness Analyser

Keywords: Polyhedral compilation, Array contraction

Functional Description: PoLA is a C++ tool that optimizes the footprint of C(++) programs of the polyhedral model by applying reduced mappings deduced from dynamic analysis of the program. More precisely, we apply a dataflow analysis on traces of a program, obtained either by execution or interpretation, and infer parametrized mappings for the arrays used for intermediate computations. This tool is part of the Polytrace project.

URL: <https://hthieven.gitlabpages.inria.fr/pola/>

Publications: [thievenaz:hal-03862219](#), [thievenaz:hal-03862218](#)

Contact: Christophe Alias

Participants: Hugo Thievenaz, Christophe Alias, Keiji Kimura

Partner: Waseda University

7.1.17 Actors-OCaml

Keywords: Concurrency, Ocaml

Functional Description: An actor library for OCaml

URL: <https://gitlab.inria.fr/mandrieu/actors-ocaml>

Contact: Gabriel Radanne

7.1.18 ctrees

Name: Choice Trees

Keywords: Coq, Concurrency, Formalisation, Semantics, Proof assistant

Functional Description: We develop so-called "ctrees", a data-structure in Coq suitable for modelling and reasoning about non-deterministic programming languages as an executable monadic interpreter. We link this new library to the Interaction Trees project: ctrees offer a valid target for interpretation of non-deterministic events.

URL: <https://github.com/vellvm/ctrees/>

Contact: Yannick Zakowski

8 New results

This section presents the scientific results obtained in the evaluation period. They are grouped according to the directions of our research program.

8.1 Research direction 1: Parallel and Dataflow Programming Models

8.1.1 Flexible Synchronization for Parallel Computations.

Participants: Ludovic Henrio, Matthieu Moy, Amaury Maillé.

Parallel applications make use of parallelism where work is shared between tasks; often, tasks need to exchange data stored in arrays and synchronize depending on the availability of these data. Fine-grained synchronizations, *e.g.* one synchronization for each element in the array, may lead to too many synchronizations while coarse-grained synchronizations, *e.g.* a single synchronization for the whole array, may prevent parallelism. This year, we worked on a parametrisation of FIFO queues that are used to stream information between processes. We proposed an extension of classical FIFO queues where the granularity of synchronization can vary at runtime, and the granularity can be chosen automatically based on an analytical model. We currently prioritize the redaction of the Ph.D manuscript of Amaury Maillé, but expect to publish these results later on.

8.1.2 An Optimised Flow for Futures: From Theory to Practice.

Participants: Nicolas Chappe, Ludovic Henrio, Amaury Maillé, Matthieu Moy, Hadrien Renaud.

The idea of Data-Flow Explicit Future is not new: it has been proposed by Ludovic Henrio in [0], and a partial implementation and theoretical study was done in [0]. An implementation of data-flow explicit futures in the Encore language was made but not published during the M2 internship of Amaury Maillé. We performed performance analysis and compared several implementation during the internship of Hadrien Renaud, and proposed an optimization together with its proof of correctness during the internship of Nicolas Chappe. These three internships were concluded by a journal publication [0], presented in 2022 at the <Programming> conference.

8.1.3 Locally abstract globally concrete semantics

Participants: Ludovic Henrio, Reiner Hähnle, Einar Broch Johnsen, Violet Ka I Pun, Crystal Chang Din, Lizeth Tapia Tarifa.

This research direction aims at designing a new way to write semantics for concurrent languages. The objective is to design semantics in a compositional way, where each primitive has a local behavior, and to adopt a style much closer to verification frameworks so that the design of an automatic verifier for the language is easier. The local semantics is expressed in a symbolic and abstract way, a global semantics gathers the abstract local traces and concretizes them. We have a reliable basis for the semantics of a simple language (a concurrent while language) and for a complex one (ABS), but the exact semantics and the methodology for writing it is still under development. After 2 meetings in 2019, this work has slowed down in 2020 and 2021, partly because of Covid restrictions but the visit of Reiner Hähnle in the Cash team allowed us to progress on the subject and to prepare a follow-up relating scheduling and LAGC.

In 2022 we finished the journal paper and submitted it to TOPLAS journal, we are currently submitting a revision of the article. The work on scheduling derived from LAGC semantics has been improved and illustrates how to specify a scheduler directly derived from LAGC semantics. The separation of concerns in the LAGC semantics between rules that generate computation states on one hand and the scheduling rules on the other, makes it possible to characterize fairness *constructively* at a semantic level and prove fairness of the scheduling at this level. The work on scheduling is under submission to the ECOOP conference.

This is a joint with Reiner Hähnle (TU Darmstadt), Einar Broch Johnsen, Crystal Chang Din, Lizeth Tapia Tarifa (Univ Oslo), Ka I Pun (Univ Oslo and Univ of applied science).

8.1.4 Deterministic parallel programs

Participants: Ludovic Henrio, Einar Broch Johnsen, Violet Ka I Pun, Yannick Zakowski.

This research direction takes place through visits and remote meetings between Ludovic Henrio and our Norwegian colleagues. First results were published in 2021 on a simple static criteria for deterministic behaviour of active objects. We are now extending this work to be able to ensure deterministic behaviour in more cases and to lay a theoretical background that will make our results more general and easier to adapt to different settings. We started formalising our results in Coq with Yannick Zakowski. Another visit from our colleagues to the CASH team is planned for 2023.

8.1.5 PNets: Parametrized networks of automata

Participants: Ludovic Henrio, Quentin Corradi, Eric Madelaine, Rabéa Ameur Boulifa.

pNets (parameterised networks of synchronised automata) are semantic objects for defining the semantics of composition operators and parallel systems. We have used pNets for the behavioral specification and verification of distributed components, and proved that open pNets (i.e. pNets with holes) were a good formalism to reason on operators and parameterized systems. This year, we have the following new results:

- We published the journal article that formalises all the results concerning the bisimulation theory for open pNets. The article has been published in the JLAMP journal in 2023 [0].
- We worked again on the refinement theory for open pNets.

An article that follows the internship of Quentin Corradi on refinement for open pNets is being written.

8.1.6 A Survey on Verified Reconfiguration

Participants: Ludovic Henrio, Helene Coullon, Frederic Loulergue, Simon Robillard.

We have conducted a survey on the use of formal methods to ensure safety of reconfiguration of distributed system, that is to say the runtime adaptation of a deployed distributed software system. The survey article is written together with H el ene Coullon and Simon Robillard (IMT Atlantique, Inria, LS2N, UBL), and Fr ed eric Loulergue (Northern Arizona University). H el ene Coullon is the coordinator and the article is under major revision. A second revision is being reviewed in the journal ACM Computing Survey.

8.1.7 A Survey on Parallelism and Determinacy

Participants: Ludovic Henrio, Laure Gonnord, Lionel Morel, Gabriel Radanne.

We have started to investigate on the solutions that exist to ensure complete or partial determinacy in parallel programs. The objective of this work is to provide a survey based on the different kinds of solutions that exist to ensure determinism or at least limit data-races in concurrent execution of programs. The study will cover language-based, compilation-based and also runtime-based solutions. We started the bibliographic studies in 2019. After a last revision, the survey has been accepted in September 2022 [0].

8.1.8 Verified Compilation Infrastructure for Concurrent Programs

Participants: Nicolas Chappe, Ludovic Henrio, Matthieu Moy, Yannick Zakowski.

The objective of this research direction is to provide semantic and reasoning tools for the formalization of concurrent programs and the verification of compilers for concurrent languages. In particular, we want to apply these results to the design of verified optimizing compilers for parallel high-level languages. We wish to proceed in the spirit of the approach advocated in Vellvm [0]: compositional, modular, executable monadic interpreters based on Interaction Trees [0] are used to specify the semantics of the language, in contrast with more traditional transition systems. Proving correct optimizations for such concurrent languages naturally requires new proof techniques that we need to design as well. This research direction is at its early stages, but we have promising early results. This year has seen major achievements pushed to fruition:

- The "ctrees" project mentioned last year has been completed with great success. Indeed, a fully fledged Coq library for modeling and reasoning about nondeterministic computations through monadic interpreters has been designed and implemented 7.1.18. Its utility has been assessed through two main case studies, ccs to demonstrate the use over message passing languages, as well as a toy language with cooperative scheduling to demonstrate the modeling of shared memory. Developed in collaboration with Steve Zdancewic and Paul He from the University of Pennsylvania, this work has been accepted to the conference POPL23 [0].
- Nicolas Chappe, additionally to major contributions to the core of the "ctrees" project, has used the library to design a framework to model uniformly various memory model plugged into a common language. He has demonstrated how to formally establish simulation relations between different models, and plans on investigating executable schedulers in this context. We project this work to crystallize into a paper by the end of spring.

8.1.9 Heterogeneous monadic computations in Coq

Participants: Jean Abou Samra, Martin Bodin, Yannick Zakowski.

Libraries such as the interaction trees advocate for an approach to modeling computations based on the free monad: a computation is first characterize by the operations it can perform. This statement can be taken to various degrees: for instance when manipulating the states, all components may be stated as broadly able to interact with the memory through read and writes, or each computation may independently specify the subset of locations they may read or write. Martin Bodin, Inria researcher in the SPADES team, and myself have started investigating how one can push in practice the fine specification approach. Preliminary results, obtained with the help of a student, has been accepted for publication at the national conference JFLA'23[zakowski:hal-03886975].

8.1.10 Actors and algebraic effects

Participants: Martin Andrieux, Ludovic Henrio, Gabriel Radanne.

This works aims to understand the link between two constructions. Actors, on one hand, aim to provide high level language constructors for concurrency and parallelism. They have been implemented and successfully used in several industry-grade frameworks, such as Akka. Algebraic effects allow the precise modelling of operation with effects, while providing excellent composition properties. They have been used both as a fundamental primitive for theoretical study, but also used as effective building blocks

to create new complex control and effectful operators. The new version of OCaml with multicore support promotes the use of algebraic effects to implement new concurrency primitives.

In this line of work, we implement actors using algebraic effects, both to better understand the relation between these two construction, and to obtain a practical, efficient implementation of Actors for OCaml. We designed such embedding and implemented it as a proof of concept library 7.1.17 using multicore OCaml. Our current embedding support concurrency and parallelism, as long as verifies some properties to ensure safety of the high-level constructs. We are now working on formalizing the embedding, and aim to extend the implementation to a distributed setting.

8.2 Research direction 2: Expressive, Scalable and Certified Analyses

8.2.1 Verification of electric properties on transistor-level descriptions of circuits, using formal methods

Participants: Oussama Oulkaid, Bruno Ferres, Ludovic Henrio, Matthieu Moy, Gabriel Radanne.

We started discussions with the [Aniah](#) start-up in 2019, and started a formal partnership in 2022, with the recruitment of Bruno Ferres as a post-doc, and Oussama Oulkaid as a Ph.D (co-supervised by Aniah, Verimag, and LIP). We developed the tool ERCtool (Electrical Rules Checking Tool), a compiler from transistor-level circuit descriptions (CDL file format) to logical formula expressing the semantics of the circuit plus a property to verify, and use an SMT solver (Z3) to check the validity of the property. The tool was successfully used on a real-life case study, and we showed that our approach can reduce the number of false-alarms significantly compared to traditional approaches. We submitted a short presentation of these results to the “late breaking results” track of the DATE 2023 conference, and we expect to submit a more complete version in 2023.

In parallel with the technical work, we conducted a thorough review of existing work on the domain, and are currently writing a survey article.

8.2.2 Data Abstraction: A General Framework to Handle Program Verification of Data Structures

Participants: Julien Braine, Laure Gonnord, David Monniaux.

Proving properties on programs accessing data structures such as arrays often requires universally quantified invariants, e.g., “all elements below index i are nonzero”. In our paper accepted at SAS 2021, we propose a general data abstraction scheme operating on Horn formulas, into which we recast previously published abstractions. We show that our instantiation scheme is relatively complete: the generated purely scalar Horn clauses have a solution (inductive invariants) if and only if the original problem has one expressible by the abstraction. This work is the topic of Julien Braine’s PhD thesis [0], successfully defended in May 2022.

8.2.3 Search functions by types

Participants: Gabriel Radanne, Laure Gonnord, Clement Allain, Pauline Garelli.

Sometimes, we need a function so deeply that we have to go out and search for it. How do we find it? Sometimes, we have clues: a function which manipulates `list` is probably in the `List` module ... but not always! Sometimes, all we have is its functionality: doing the sum of a list of integers. Unfortunately, search by functionality is difficult.

Rittri proposed an approximation: use the *type* of the function as a key to search through libraries: in our case, `int list -> int`. To avoid stumbling over details such as the order of arguments, he proposed

to use matching *modulo type isomorphism* – a notion broader than syntactic equality. Unfortunately, algorithms for unification modulo type isomorphisms are extremely costly (at best NP). An exhaustive search would not be usable during programming in practice.

In this research direction, we investigate how to scale search by types. For this purpose, we developed new algorithm technique similar to indexes used in databases, but appropriate for keys following a rich language of types. Early result have been published in ML2021 [0] and we have developed a prototype, Dowsing 7.1.13, implementing these ideas. We have since developed a new indexing that exploits the partial order of “matching” to further speed-up the search [0].

8.2.4 A new module system for OCaml

Participants: Clement Blaudeau, Didier Remy, Gabriel Radanne.

ML modules are offer large-scale notions of composition and modularity. Provided as an additional layer on top of the core language, they have proven both vital to the working OCaml and SML programmers, and inspiring to other use-cases and languages. Unfortunately, their meta-theory remains difficult to comprehend, requiring heavy machinery to prove their soundness.

In this research direction, we study a translation from ML modules to $F\omega$ to provide a new comprehensive description of a generative subset of OCaml modules, embarking on a journey right from the source OCaml module system, up to $F\omega$, and back. We propose a “middle representation” called canonical that combines the best of both worlds. Our goal is to obtain type soundness, but also and more importantly, a deeper insight into the signature avoidance problem, along with ways to improve both the OCaml language and its typechecking algorithm.

We first looked at the “generative” subset of OCaml, which we published recently [0] and are now developing the applicative part, which would cover OCaml more accurately.

8.2.5 Formalising Futures and Promises in Viper

Participants: Ludovic Henrio, Cinzia Di Giusto, Loïc Germerie Guizouarn, Etienne Lozes.

Futures and promises are respectively a read-only and a write-once pointer to a place-holder in memory. They are used to transfer information between threads in the context of asynchronous concurrent programming. Nonetheless they can be error prone as data races may arise when references are shared and transferred. We aim at providing a formal tool to detect those errors. We propose a proof of concept by implementing the future/promise mechanism in Viper: a verification infrastructure, that provides a way to reason about resource ownership in programs. This work was published in JFLA 2022 [0].

8.3 Research direction 3: Optimizing Program Transformations

8.3.1 Optimizing compilation for tree-like structures

Participants: Paul Iannetta, Laure Gonnord, Gabriel Radanne.

Trees are used everywhere, yet their internal operations are nowhere as optimized as arrays are. This work is in the continuity of the research on cache-oblivious algorithms for trees. It investigate the properties of structural transformations on trees via two directions.

First, we studied how to develop structural low-level operations to efficiently manipulate trees in-place and how to optimized and parallelize such operations [0]. This yielded an optimized implementation to manipulate AVL trees 7.1.11.

Following this study, we developed a more general framework that aims to decompose and optimize arbitrary transformation on trees by expressing them as rewrite rules on Algebraic Data Types. In this approach, the traditional semantic of functional programs with pattern matching is reinterpreted as optimized, in-place, imperative loop nests. These results were published in GPCE2021 [0] and implemented in a prototype compiler 7.1.12. This work is the topic of Paul Iannetta's thesis, successfully defended May 2022 [0].

8.3.2 Parallel rewriting operations on Trees

Participants: Thaïs Baudon, Carsten Fuhs, Laure Gonnord.

During the M2 internship of Thaïs Baudon, partially funded by the CODAS's ANR project, we studied the problem of scheduling programs with complex data structures with the point of view of term rewriting. As related work, we show that although scheduling operations on efficient data structures apart from arrays should have a great impact on performance, there is nearly no work on this topic. However, some related work from the rewriting community gives us great insights about the link between termination and scheduling. Our contributions as a first step toward full efficient compilation of programs with inductive data structures are 1. first algorithms for this parallel evaluation, 2. a code generation algorithm to generate efficient parallel evaluators, 3. a prototype implementation and first experimental results

The novel parallel complexity notion was initially published in WST'21 [0] and further developed in LOPSTER'22 [0]. The confluence aspect has been published in IWC'22 [0].

8.3.3 Memory optimizations for Algebraic Data Types

Participants: Thaïs Baudon, Gabriel Radanne, Laure Gonnord.

In the last few decades, Algebraic Data Types (ADT) have emerged as an incredibly effective tool to model and manipulate data for programming. Additionally, ADTs could provide numerous advantages for optimizing compilers, as the rich declarative description could allow them to choose the memory representation of the types.

Initially, ADTs were mostly present in functional programming languages such as OCaml and Haskell. Such GC-managed functional languages generally use uniform memory representation which prohibit aggressive optimisations of the representation of ADTs. However, ADTs are now present in many different languages, notably Scala and Rust, which permit such optimizations.

The goal of this research direction is to investigate how to represent terms of Algebraic Data Types and how to compile pattern matching efficiently. We aim to develop a generic compilation framework which accomodate arbitrarily complex memory representation for terms, and to provide news ways to optimize the representation of ADTs. A prototyper compiler has been implemented 7.1.10. We have now developed a framework to describe such representations, which was published at AFADL [0], along with a set of tools to automatically synthesis such representation based on high-level specifications using constraint-solving [0].

8.3.4 Vellvm: Verified LLVM

Participants: Calvin Beck, Bastien Rousseau, Irene Yoon, Yannick Zakowski, Steve Zdancewic.

We develop, in collaboration with the University of Pennsylvania, a formally verified in Coq compilation infrastructure based on LLVM, dubbed Vellvm 7.1.5. Compared to other existing verified compilation framework, we define the semantics of the languages we consider as monadic interpreters built on top of

the Interaction Trees framework. This approach brings us benefits in terms of modularity, compositionality and executability, as well as leads to an equational mode of reasoning to establish refinements. The following major achievements have taken place this year:

- We have published the description of our new infrastructure at ICFP'21 [0].
- A major redefinition of the memory model, led by Calvin Beck and accounting for the necessary finite view of the memory imposed by the presence of pointer to integer casts in LLVM IR, has been started. Most of it is finished, we are working on validating the new model by proving optimizations corrects and will submit an article about it during the year 2022.
- Reasoning about the various monadic interpreters involved in such semantics is challenging. A major process of abstraction and generalization of the equational theory expected from these monads, and a zoo of implementations of these theories, has been developed — an effort led by Irene Yoon. We have published this year these results at ICFP'22 [0]
- Spiral is a compilation framework for the generalization of efficient low level code for numerical computations. HELIX is a formalization in Coq of part of this framework that Vadim Zaliva has developed during his PhD. We have proved correct a back end of the HELIX compilation chain down to Vellvm. The full chain will be described in a journal paper during the year 2022.
- As part of his internship, Bastien Rousseau has developed a set of verified combinators to build control flow graphs compositionally and anonymously. By providing in this way a DSL to write program transformations over LLVM IR, we obtain for free a set of well-formedness properties over the generated code. Furthermore, each combinator coming with a semantic characteristic equation, we obtain compositional reasoning principles. We plan on complementing this work with its dual: a DSL of patterns to deconstruct a graph into its construction as combinators.

8.3.5 Verified Abstract Interpreters as Monadic Interpreters

Participants: Laure Gonnord, Sébastien Michelland, Yannick Zakowski.

In the realm of verified compilation, one typically wants to verify the static analyzes used by the compiler. In existing works, the analysis is typically written as a fuel-based pure function in Coq and verified against the semantics described as a transition system. The goal of this research is to develop the tools and reasoning principles to transfer these ideas to a context where the semantics of the language is defined as a monadic interpreters built on Interaction Trees.

During his internship, Sébastien Michelland has developed a first promising prototype, establishing a highly modular framework to build and prove correct such analyses. He has instantiated his result on a toy Imp language, and is now aiming at instantiating it on a toy assembly language. We plan on submitting an article describing these first results, before considering the application of the technique to Vellvm.

8.3.6 Lightweight Array Contraction by Trace-Based Polyhedral Analysis

Participants: Hugo Thievenaz, Keiji Kimura, Christophe Alias.

In this work, we defend the iconoclast idea that polyhedral optimizations might be computed without expensive polyhedral operations, simply by applying a lightweight analysis on a few off-line execution traces. The main intuition being that, since polyhedral transformations are expressed as affine mappings, only a few points are required to infer the general mapping. Our hope is to compute those points from a few off-line execution traces. We focus on array contraction, a well known technique to reallocate temporary arrays thanks to affine mappings so the array size is reduced. We describe a trace selection

algorithm, a liveness algorithm from an execution trace, and another to compute the maximum number of variables alive alongside a dimension, from which we get our scalar modular mappings. We show that a simple interpolation allow to infer the modulo mapping.

Our results have been presented at C3PO'22 [0] and IMPACT'22 [0]. We are currently working on the inference of more general linear mappings, able to reduce further the memory footprint.

8.3.7 A Polyhedral Approach for Scalar Promotion

Participants: Alec Sadler, Christophe Alias, Hugo Thievenaz.

Memory accesses are a well known bottleneck whose impact might be mitigated by using properly the memory hierarchy until registers. In this work, we address array scalarization, a technique to turn temporary arrays into a collection of scalar variables to be allocated to registers. We revisit array scalarization in the light of the recent advances of the polyhedral model. We propose a general algorithm for array scalarization, ready to be plugged in a polyhedral compiler, among other passes. Our scalarization algorithm operates on the polyhedral intermediate representation. In particular, our scalarization algorithm is parametrized by the program schedule possibly computed by a previous compilation pass. We rely on schedule-directed array contraction and we propose a loop tiling algorithm able to reduce the footprint down to the available amount of registers on the target architecture. Experimental results confirm the effectiveness and the efficiency of our approach, in particular for optimizing tensor kernels.

Our first results have been presented at COMPAS'22 [0]. We are investigating the application of this compilation technique to high-level synthesis.

8.3.8 Affine Multibanking for High-Level Synthesis

Participants: Christophe Alias, Matthieu Moy.

High-level circuit compilers (high-level synthesis, HLS) are able to translate a C program to an FPGA circuit configuration. Unlike software parallelization, there is no operating system to place the computation and the memory at runtime. All the parallelization decisions must be done at compile time. In this work, we address the compilation of data placement under parallelism and resource constraints. We propose an HLS algorithm able to partition the data across memory banks, so parallel access will target distinct banks to avoid data transfer serialization. Our algorithm is able to minimize the number of banks and the maximal bank size.

Our results have been presented at IMPACT'22 [0]. We are working on extending the experimental validation to more general polyhedral kernels and design space exploration to find an appropriate trade-off circuit size/latency.

8.3.9 Scalable Parallelization for High-Level Synthesis

Participants: Christophe Alias.

We have proposed a complete compilation chain based on data-aware process networks (DPN), a dataflow intermediate representation suitable for HLS in the context of high-performance computing. DPN combines the benefits of a low-level dataflow representation-close to the final circuit-and affine iteration space tiling to explore the parallelization trade-offs (local memory size, communication volume, parallelization degree). We are currently working on a parallelization scheme based on parametrization and instantiation, in order to scale-up the parallelization process.

The final goal is to integrate our results to the XTREMLOGIC production compiler.

8.3.10 Compiler Support for Dynamic Task Scheduling

Participants: Christophe Alias, Samuel Thibault, Corentin Bouquet.

The recent task-based runtime systems usually rely on dynamic task scheduling, using mostly eager heuristics with a limited view over the application workload, thus limiting their optimization potential. The task granularity is also most often set to a single size, which does not fit the heterogeneity of nowadays' platforms. We propose to leverage the view that a compiler can have with polyhedral analysis, to help the runtime, i.e. bringing a *static* (compile-time) knowledge into the *dynamic* (run-time) execution.

In his internship, Corentin Bouquet shows how time-to-completion might precomputed thanks to polyhedral techniques and study the behaviour of the STARPU runtime with these informations.

8.3.11 Dynamic Analysis for Sparse Kernels

Participants: Christophe Alias, Gabriel Dehame.

Automatic parallelisation focuses on programming automatically parallel computers from a sequential specification. In the past decades, a general unified framework [0] was designed to solve that problem for *regular* loop kernels manipulating *dense* tensors (arrays). However, most kernels of interest manipulate *sparse* tensors.

In this internship, Gabriel investigated the inference *automatically, at runtime* of the tensors regions flowing between the loop kernels of a scientific application. Several techniques ranging from linear relation analysis to explicit resolution with affine relations were tried and evaluated.

8.4 Research direction 4: Simulation and Hardware

8.4.1 S4BXI: the MPI-ready Portals 4 Simulator

Participants: Julien Emmanuel, Matthieu Moy, Ludovic Henrio, Grégoire Pichon.

We present a simulator for High Performance Computing (HPC) interconnection networks. It models Portals 4, a standard low-level API for communication, and it allows running unmodified applications that use higher-level network APIs such as the Message Passing Interface (MPI). It is based on SimGrid, a framework used to build single-threaded simulators based on a cooperative actor model. Unlike existing tools like SMPI, we rely on an actual MPI implementation, hence our simulation takes into account MPI's implementation details in the performance. This paper also presents a case study using the BullSequana eXascale Interconnect (BXI) made by Atos, which highlights how such a simulator can help design space exploration (DSE) for new interconnects. In 2022 we consolidated the work with new experimental results, bugfixes, and improvement of the model. Our priority on the last semester was the redaction of the Ph.D manuscript that has been submitted in January 2023.

9 Bilateral contracts and grants with industry

9.1 CIFRE Ph.D of Julien Emmanuel with Bull/Atos, hosted by Inria. 2020-2023.

Participants: Matthieu Moy, Ludovic Henrio, Julien Emmanuel.

We reached the last year of the CIFRE Ph.D of Julien Emmanuel. The last year was dedicated to reinforcement of existing results, improving the accuracy of our performance model for the BXI interconnect developed at Atos. The Ph.D defense is scheduled for March 8th.

9.2 Partnertship with the Aniah startup on circuit verification

Participants: Bruno Ferres, Matthieu Moy, Ludovic Henrio, Gabriel Radanne.

The CASH team started discussion with the **Aniah** startup in 2019, to work on verification of electrical properties of circuits at transistor level. We recruited a post-doc (Bruno Ferres) in March 2022, and formalized the collaboration with a bilateral contract (Réf. Inria : 2021-1144), in parallel with a joint internship with LIP, Verimag laboratory and Aniah (Oussama Oulkaid), which led to a CIFRE Ph.D (LIP/Verimag/Aniah) started in October 1st 2022. The collaboration led to the development of the tool ERCtool, and to one article submission at the DATE conference.

9.3 CAVOC Project with Inria/Nomadic Labs

Participants: Guilhem Jaber, Gabriel Radanne, Laure Gonnord.

This project aims to develop a *sound and precise static analyzer* for OCaml, that can catch large classes of bugs represented by uncaught exceptions. It will deal with both user-defined exceptions, and built-in ones used to represent *error behaviors*, like the ones triggered by `failwith`, `assert`, or a match failure. Via “assert-failure” detection, it will thus be able to check that invariants annotated by users hold. The analyzer will reason *compositionally* on programs, in order to analyze them at the granularity of a function or of a module. It will be *sound* in a strong way: if an OCaml module is considered to be correct by the analyzer, then one will have the guarantee that no OCaml code interacting with this module can trigger uncaught exceptions coming from the code of this module. In order to be *precise*, it will take into account the abstraction properties provided by the type system and the module system of the language: local values, abstracted definition of types, parametric polymorphism. The goal being that most of the interactions taken into account correspond to typeable OCaml code.

This project is part of the partnership between Inria and Nomadic Labs, and lead by Guilhem Jaber, from the Inria Team Galinette.

10 Partnerships and cooperations

10.1 International initiatives

10.1.1 Inria Exploratory Actions

POLYTRACE

Title: Rephrasing Polyhedral Compilers with Trace-based analysis

Duration: 2021 -> 2024

Coordinator: Christophe Alias

Partners: • Prof. Keiji Kimura (Waseda University, Tokyo, Japan)

Summary: The polyhedral model provides powerful but expensive analysis for compiling HPC kernels. In this project, we propose to speed-up the compilation process by interpolating polyhedral transformations from executions traces. In particular, we will investigate how to select execution traces to ensure code coverage and how to extrapolate the results of lightweight trace analysis to

program transformations. We will focus on array contraction, a key polyhedral transformation to improve the footprint. The approach will be prototyped and validated on the benchmarks of the polyhedral community. This project will benefit from the common expertise on automatic parallelization and computer architecture from Waseda University and CASH.

10.1.2 Inria associate team not involved in an IIL or an international program

CAPESA

Title: Characterization of Software Evolution with Static Analyses

Duration: 2020 ->

Coordinator: Sébastien Mosser (mosser.sebastien@uqam.ca)

Partners:

- Université du Québec À Montréal (Canada)

Inria contact: Ludovic Henrio

Summary: In this project we propose to study code transformation in terms of "semantic diff". This notion will be defined thanks to code intermediate representations such as Abstract Syntax Trees (AST) or the control flow graph, not by textual representation. The objective is not only to compute but also to manipulate these "diffs" in several contexts: being able to reapply a diff on another program than the one it comes from, quantify the interference between two diffs, and more generally to study the composability of several diffs. The approach will be experimentally validated on problems coming from the domain of expertise of both teams of the cooperation: compiler pass analyses (expertise of CASH), and git commits (expertise of UQAM). The complementarity of the analysis and compilation approaches of the CASH team and the expertise on software engineering of the UQAM member will ensure the success and the originality of the project.

10.2 International research visitors

10.2.1 Visits of international scientists

Other international visits to the team

Reiner Hähnle

Status Professor

Institution of origin: University of Darmstadt

Country: Germany

Dates: 28 February 2022 - 4 March 2022

Context of the visit: Collaboration with Ludovic Henrio on semantics of languages for parallel programming.

Mobility program/type of mobility: research stay

Violet Ka I Pun and Einar Broch Johnsen

Institution of origin: University of Bergen and Univ of Oslo

Country: Norway

Dates: 3 days in September and one week in December

Context of the visit: Collaboration with Ludovic Henrio on semantics of languages for deterministic parallel programming.

Mobility program/type of mobility: research stay – PHC Aurora project

10.2.2 Visits to international teams

- In the context of the Polytrace exploratory action, Christophe Alias has visited Prof. Keiji Kimura (Waseda University, Tokyo, Japan). To initiate collaborations around HLS, Christophe Alias also visited Prof. Shinya Takamaeda (The University of Tokyo), Prof. Kenshu Seto (Tokyo City University) and Prof. Masanori Hariyama (Tohoku University). The visit lasted 5 weeks.
- In the context of the PHC Aurora: SLAS, Ludovic Henrio visited university of Bergen for one week.

10.3 European initiatives

10.3.1 Other european programs/initiatives

PHC Aurora: SLAS

Title: Synchronous Languages meet Asynchronous Semantics

Partner Institution(s):

- university of Bergen, Norway
- university of Oslo, Norway

Date/Duration: 2020 (extended to 2021 and 2022 due to Covid)

11 Dissemination

11.1 Promoting scientific activities

11.1.1 Scientific events: organisation

General chair, scientific chair Ludovic Henrio is member of the steering committee of ICE workshops.

11.1.2 Scientific events: selection

Member of the conference program committees

- Christophe Alias was a PC member for PPOPP'23 and IMPACT'22.
- Yannick Zakowski was a PC member for ICFP'22.
- Gabriel Radanne was a PC member for CC'22, ProWeb'22, VPL'22 and JFLA'23.
- Ludovic Henrio was PC member of FACS 2022, ICE 2022

Reviewer

- Christophe Alias was a reviewer for SAS'22.
- Matthieu Moy was a reviewer for DATE'22.
- Yannick Zakowski was in the external review committee for OOPSLA'23.
- Gabriel Radanne was an external reviewer for PPOPP'23 and TACAS'23

11.1.3 Journal

Reviewer - reviewing activities

- Christophe Alias was a reviewer for Parallel Computing, TACO (ACM) and TECS (ACM).
- Matthieu Moy was a reviewer for ACM TECS'22.
- Gabriel Radanne was a reviewer for JFP and TioT.
- Ludovic Henrio was a reviewer for JLAMP and IJPP.

11.1.4 Invited talks

- Christophe Alias gave a talk to the FPGA day ("*Calcul scientifique accéléré sur FPGA*") organised by the GdR 720 ISIS (Information, Signal, Image et ViSion).
- Christophe Alias gave a serie of doctoral lectures on Polyhedral Compilation at Waseda University (Tokyo, Japan).

11.1.5 Leadership within the scientific community

Ludovic Henrio is one of the responsilbes, with Kévin Martin of the GdT CLAP inside the GDR GPL.

Laure Gonnord is responsible for "Ecole des Jeunes Chercheurs en Programmation" in the GdR GPL.

11.1.6 Scientific expertise

- Christophe Alias is scientific advisor for the XTREMLOGIC startup.

11.1.7 Research administration

- Ludovic Henrio is member of the "commission recherche" of labex Milyon.

11.2 Teaching - Supervision - Juries

11.2.1 Teaching

Licence

- Christophe Alias: "Compilation", INSA CVL 3A, cours+TD, 27h ETD.
- Hugo Thievenaz: "Algorithmique programmation impérative, initiation", L1 UCBL, TD+TP, 49h ETD.
- Matthieu Moy: Responsable of the "licence d'informatique UCBL". "Programmation concurrente", L3 UCBL, 26h; "Projet Informatique", 9.5h.
- Gabriel Radanne: "Projet Informatique", 2 groupes.
- Nicolas Chappe: "Programmation", L3 informatique, ENS Lyon, TP, 32h ETD.
- Amaury Maillé: "Algorithmique et Programmation Orientée Objet", L3 UCBL, TD, 14h EQTD
- Amaury Maillé: "Algorithmique et Programmation Orientée Objet", L3 UCBL, TD+TP, 32.16h EQTD
- Amaury Maillé: "Programmation Logique", L2 UCBL, TP, 12h EQTD
- Amaury Maillé: "Algorithmique, programmation impérative, initiation", L1 UCBL, TD+TP, 1.67h EQTD
- Amaury Maillé: "Algorithmique, programmation et structures de données", L2 UCBL, TD+TP, 24h EQTD
- Amaury Maillé: "Architecture des ordinateurs", L2 UCBL, TP, 16h EQTD

Master 1

- Christophe Alias: "Optimisation des applications embarquées", INSA CVL 4A, cours+TD, 27h ETD.
- Christophe Alias: "Compilation", Préparation à l'agrégation d'informatique, ENS-Lyon, cours, 18h ETD.
- Hugo Thievenaz: "Compilation / traduction des programmes", M1 UCBL, TP, 15h ETD.
- Matthieu Moy: "Compilation et traduction des programmes", M1 UCBL, responsable, 31h; "Gestion de projet et génie logiciel", M1 UCBL, responsable, 32h; "Projet pour l'Orientation en Master", M1 UCBL, 2 students supervised.
- Gabriel Radanne, Ludovic Henrio and Nicolas Chappe: "Compilation and Analysis" (CAP), ENS-Lyon, Master d'Informatique Fondamentale, cours, 48h CM + 28h TD, 64h ETD.
- Gabriel Radanne: "Projet pour l'Orientation en Master", M1 UCBL, 2 students supervised.

Master 2

- Christophe Alias, Laure Gonnord, Yannick Zakowski: "Static Analysis and Optimizing Compilers", ENS-Lyon, cours+TD, 48h ETD.
- Yannick Zakowski: "Program Verification with Coinduction and Proof Assistants", ENS-Lyon, cours, 24h ETD.
- Christophe Alias: "Polyhedral High-Level Synthesis", Université de Bretagne Occidentale, cours, 7.5ETD.
- Ludovic Henrio: University of Nice Sophia Antipolis, M2 Ubinet. "an algorithmic approach to distributed systems"3h30 CM+TD
- Nicolas Chappe: "Systèmes et réseaux", Préparation à l'agrégation d'informatique, ENS Lyon, TP, 10h ETD.
- Thaïs Baudon: "Compilation", Préparation à l'agrégation d'informatique, ENS-Lyon, TP, 10h ETD.

11.2.2 Supervision

- Christophe Alias co-advises the PhD thesis of Hugo Thievenaz with Prof. Keiji Kimura (Waseda University).
- Ludovic Henrio and Yannick Zakowski co-advise the PhD thesis of Nicolas Chappe.
- Gabriel Radanne and Laure Gonnord co-advise the PhD thesis of Thaïs Baudon.
- Julien Emmanuel's PhD thesis is supervised by Matthieu Moy and Ludovic Henrio
- Amaury Maillé's PhD thesis is supervised by Ludovic Henrio and Matthieu Moy

11.2.3 Defended Ph.D

- Julien Braine: "The Data-abstraction Framework : abstracting unbounded data-structures in Horn clauses, the case of arrays" [0], directed by Laure Gonnord and David Monniaux.
- Paul Iannetta: "Compiling Trees : Combining Data Layouts and the Polyhedral Model" [0], directed by Laure Gonnord, Gabriel Radanne and Lionel Morel.

11.2.4 Juries

- Christophe Alias was a reviewer for the PhD thesis of Basile Clément at ENS Paris, supervised by Albert Cohen and Xavier Leroy.
- Christophe Alias was the computer science examiner for the *oral d'informatique du second concours d'entrée à l'ENS de Lyon*, session 2022.

11.3 Popularization

11.3.1 Internal or external Inria responsibilities

- Christophe Alias was a member of the hiring jury for CRCN and ISFP research positions at Inria Lyon.

11.3.2 Education

- Thaïs Baudon: Conception d'activités débranchées pour la vulgarisation de l'informatique et des mathématiques auprès du grand public, Maison des Mathématiques et de l'Informatique (MMI), 32h ETD.

12 Scientific production

12.1 Major publications

- [0] C. Alias and A. Plesco. 'Data-Aware Process Networks'. In: CC 2021 - 30th ACM SIGPLAN International Conference on Compiler Construction. Virtual, South Korea: ACM, 2nd Mar. 2021, pp. 1–11. DOI: [10.1145/3446804.3446847](https://doi.org/10.1145/3446804.3446847). URL: <https://hal.inria.fr/hal-03143777>.
- [0] G. Busnot, T. Sassolas, N. Ventroux and M. Moy. 'Standard-compliant parallel SystemC simulation of loosely-timed transaction level models: From baremetal to Linux-based applications support'. In: *Integration, the VLSI Journal* 79 (July 2021), pp. 23–40. DOI: [10.1016/j.vlsi.2020.12.006](https://doi.org/10.1016/j.vlsi.2020.12.006). URL: <https://hal.archives-ouvertes.fr/hal-03487607>.
- [0] N. Chappe, L. Henrio, A. Maillé, M. Moy and H. Renaud. 'An Optimised Flow for Futures: From Theory to Practice'. In: *The Art, Science, and Engineering of Programming* 6.1 (15th July 2021), pp. 1–41. DOI: [10.22152/programming-journal.org/2022/6/3](https://doi.org/10.22152/programming-journal.org/2022/6/3). URL: <https://hal.inria.fr/hal-03440766>.
- [0] K. Fernandez-Reyes, D. Clarke, L. Henrio, E. Broch Johnsen and T. Wrigstad. 'Godot: All the Benefits of Implicit and Explicit Futures'. In: ECOOP 2019 - 33rd European Conference on Object-Oriented Programming. Leibniz International Proceedings in Informatics (LIPIcs). London, United Kingdom, 2019, pp. 1–28. URL: <https://hal.archives-ouvertes.fr/hal-02302214>.
- [0] L. Gonnord, L. Henrio, L. Morel and G. Radanne. 'A Survey on Parallelism and Determinism'. In: *ACM Computing Surveys* (27th Sept. 2022). DOI: [10.1145/3564529](https://doi.org/10.1145/3564529). URL: <https://hal.inria.fr/hal-03828497>.
- [0] Y. Zakowski, C. Beck, I. Yoon, I. Zaichuk, V. Zaliva and S. Zdancewic. 'Modular, compositional, and executable formal semantics for LLVM IR'. In: *Proceedings of the ACM on Programming Languages* 5.ICFP (22nd Aug. 2021), pp. 1–30. DOI: [10.1145/3473572](https://doi.org/10.1145/3473572). URL: <https://hal.archives-ouvertes.fr/hal-03525711>.

12.2 Publications of the year

International journals

- [0] R. Ameur-Boulifa, L. Henrio and E. Madelaine. 'Compositional Equivalences Based on Open pNets'. In: *Journal of Logical and Algebraic Methods in Programming* (2022). URL: <https://hal.science/hal-03894031>.

- [0] N. Chappe, P. He, L. Henrio, Y. Zakowski and S. Zdancewic. ‘Choice Trees: Representing Nondeterministic, Recursive, and Impure Programs in Coq’. In: *Proceedings of the ACM on Programming Languages* (15th Jan. 2023). DOI: [10.1145/3571254](https://doi.org/10.1145/3571254). URL: <https://hal.science/hal-03886910>.
- [0] L. Gonnord, L. Henrio, L. Morel and G. Radanne. ‘A Survey on Parallelism and Determinism’. In: *ACM Computing Surveys* (27th Sept. 2022). DOI: [10.1145/3564529](https://doi.org/10.1145/3564529). URL: <https://hal.inria.fr/hal-03828497>.
- [0] I. Yoon, Y. Zakowski and S. Zdancewic. ‘Formal reasoning about layered monadic interpreters’. In: *Proceedings of the ACM on Programming Languages* 6.ICFP (29th Aug. 2022), pp. 254–282. DOI: [10.1145/3547630](https://doi.org/10.1145/3547630). URL: <https://hal.science/hal-03850324>.

International peer-reviewed conferences

- [0] T. Baudon, C. Fuhs and L. Gonnord. ‘Analysing Parallel Complexity of Term Rewriting’. In: LOPSTR 2022 - 32nd International Symposium on Logic-based Program Synthesis and Transformation. Tbilisi, Georgia, 21st Sept. 2022. URL: <https://hal.science/hal-03749679>.
- [0] T. Baudon, C. Fuhs and L. Gonnord. ‘On Confluence of Parallel-Innermost Term Rewriting’. In: <http://cl-informatik.uibk.ac.at/iwc/2022/proceedings.pdf>. IWC 2022 - 11th International Workshop on Confluence. Haifa, Israel, 2022. URL: <https://hal.inria.fr/hal-03710007>.
- [0] C. Blaudeau, D. Rémy and G. Radanne. ‘Retrofitting OCaml modules: Fixing signature avoidance in the generative case’. In: *Journées Francophones des Langages Applicatifs*. JFLA 2023 - 34èmes Journées Francophones des Langages Applicatifs. Praz-sur-Arly, France, 16th Jan. 2023, pp. 59–100. URL: <https://hal.inria.fr/hal-03936636>.
- [0] B. Ferres, O. Oulkaid, L. Henrio, M. Khosravian, M. Moy, G. Radanne and P. Raymond. ‘Electrical Rule Checking of Integrated Circuits using Satisfiability Modulo Theory’. In: <https://ieeexplore.ieee.org/xpl/conhome/9774496/proceeding>. DATE: Design, Automation and Test in Europe Conference. Anvers (Antwerpen), Belgium, 2022. URL: <https://hal.science/hal-04007446>.
- [0] I. Lasfar, C. Alias, M. Moy, R. Neveu and A. Carré. ‘Affine Multibanking for High-Level Synthesis’. In: IMPACT’22 - 12th International Workshop on Polyhedral Compilation Techniques. Budapest, Hungary, 20th June 2022. URL: <https://hal.inria.fr/hal-03862220>.
- [0] A. Sadler, C. Alias and H. Thievenaz. ‘A Polyhedral Approach for Scalar Promotion’. In: Conférence francophone d’informatique en Parallélisme, Architecture et Système (COMPAS’22). Amiens, France, 5th July 2022. URL: <https://hal.inria.fr/hal-03862223>.
- [0] H. Thievenaz, K. Kimura and C. Alias. ‘Lightweight Array Contraction by Trace-Based Polyhedral Analysis’. In: C3PO 2022 - International Workshop on Compiler-assisted Correctness Checking and Performance Optimization for HPC. Hamburg, Germany, 2nd June 2022. URL: <https://hal.inria.fr/hal-03862219>.
- [0] H. Thievenaz, K. Kimura and C. Alias. ‘Rephrasing Polyhedral Optimizations with Trace Analysis’. In: 12th International Workshop on Polyhedral Compilation Techniques (IMPACT’22). Budapest, Hungary, 20th June 2022. URL: <https://hal.inria.fr/hal-03862218>.

National peer-reviewed Conferences

- [0] J. Abou-Samra, Y. Zakowski and M. Bodin. ‘Effectful Programming across Heterogeneous Computations -Work in Progress’. In: *Journées Francophones des Langages Applicatifs*. JFLA 2023 - 34èmes Journées Francophones des Langages Applicatifs. Praz-sur-Arly, France, 31st Jan. 2023, pp. 7–23. URL: <https://hal.science/hal-03886975>.
- [0] T. Baudon, G. Radanne and L. Gonnord. ‘Knit&Frog: Pattern matching compilation for custom memory representations (doctoral session)’. In: AFADL 2022 - 21ème journées Approches Formelles dans l’Assistance au Développement de Logiciels. Vannes, France, 8th June 2022. URL: <https://hal.inria.fr/hal-03676356>.

- [0] C. Giusto, L. Guizouarn, L. Henrio and E. Lozes. ‘Formalising Futures and Promises in Viper’. In: *Journées Francophones des Langages Applicatifs*. JFLA 2022 - 33èmes Journées Francophones des Langages Applicatifs. Saint-Médard-d’Excideuil, France, 2nd Feb. 2022, pp. 165–183. URL: <https://hal.inria.fr/hal-03626843>.

Doctoral dissertations and habilitation theses

- [0] J. Braine. ‘The Data-abstraction Framework : abstracting unbounded data-structures in Horn clauses, the case of arrays’. Université de Lyon, 11th May 2022. URL: <https://theses.hal.science/tel-03771839>.
- [0] P. Iannetta. ‘Compiling Trees : Combining Data Layouts and the Polyhedral Model’. Université de Lyon, 2nd May 2022. URL: <https://theses.hal.science/tel-03771830>.

Reports & preprints

- [0] T. Baudon, L. Gonnord and G. Radanne. *Knit&Frog: Pattern matching compilation for custom memory representations*. RR-9473. Inria Lyon, 1st June 2022, p. 35. URL: <https://hal.inria.fr/hal-03684334>.

Other scientific publications

- [0] P. Garelli. ‘Type indexing in OCaml for function search in a large ecosystem’. Polytechnique, 20th June 2022. URL: <https://hal.inria.fr/hal-03940573>.

12.3 Other

Softwares

- [0] [SW] G. Radanne, L. Gonnord and T. Baudon, *Ribbit*, 18th Nov. 2022. LIC: MIT License. HAL: ([hal-03860442](https://hal.inria.fr/hal-03860442)), URL: <https://hal.inria.fr/hal-03860442>, VCS: <https://gitlab.inria.fr/ribbit/ribbit>, SWHID: ([swh:1:dir:b3be924ea99bdee1c1df8d9b4310e2f767536766;origin=https://hal.archives-ouvertes.fr/hal-03860442;visit=swh:1:snp:6c1082239a4d9372b0546b82945f453576931dd2;anchor=swh:1:rel:eeff41f174f0ec95c25dba2cea8acb78b859d98a;path=/](https://sw.hal.archives-ouvertes.fr/hal-03860442)).

12.4 Cited publications

- [0] C. Alias, A. Darte, P. Feautrier and L. Gonnord. ‘Multi-dimensional Rankings, Program Termination, and Complexity Bounds of Flowchart Programs’. In: *International Static Analysis Symposium (SAS’10)*. 2010.
- [0] C. Alias and A. Plesco. ‘Data-Aware Process Networks’. In: *Proceedings of the 30th ACM SIGPLAN International Conference on Compiler Construction*. CC 2021. Virtual, Republic of Korea: Association for Computing Machinery, 2021, pp. 1–11. DOI: [10.1145/3446804.3446847](https://doi.org/10.1145/3446804.3446847). URL: <https://doi.org/10.1145/3446804.3446847>.
- [0] C. Alias and A. Plesco. *Data-aware Process Networks*. Research Report RR-8735. Inria - Research Centre Grenoble – Rhône-Alpes, June 2015, p. 32. URL: <https://hal.inria.fr/hal-01158726>.
- [0] C. Alias and A. Plesco. *Method of Automatic Synthesis of Circuits, Device and Computer Program associated therewith*. Patent FR1453308. Apr. 2014.
- [0] C. Allain, G. Radanne and L. Gonnord. ‘Isomorphisms are back!’ In: *ML 2021 - ML Workshop*. Virtual, France, Aug. 2021, pp. 1–3. URL: <https://hal.archives-ouvertes.fr/hal-03355381>.
- [0] I. Amer, C. Lucarz, G. Roquier, M. Mattavelli, M. Raulet, J.-F. Nezan and O. Deforges. ‘Reconfigurable video coding on multicore’. In: *Signal Processing Magazine, IEEE* 26.6 (2009), pp. 113–123. URL: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=5230810.
- [0] S. Ananian. ‘The Static Single Information Form’. MA thesis. MIT, Sept. 1999.

- [0] C. B. Aoun, L. Andrade, T. Maehne, F. Pêcheux, M.-M. Louërat and A. Vachoux. 'Pre-simulation elaboration of heterogeneous systems: The SystemC multi-disciplinary virtual prototyping approach'. In: *Embedded Computer Systems: Architectures, Modeling, and Simulation (SAMOS), 2015 International Conference on*. IEEE. 2015, pp. 278–285.
- [0] P. Aubry, P.-E. Beaucamps, F. Blanc, B. Bodin, S. Carpov, L. Cudennec, V. David, P. Doré, P. Dubrulle, B. Dupont De Dinechin, F. Galea, T. Goubier, M. Harrant, S. Jones, J.-D. Lesage, S. Louise, N. Morey Chaisemartin, T. H. Nguyen, X. Raynaud and R. Sirdey. 'Extended Cyclostatic Dataflow Program Compilation and Execution for an Integrated Manycore Processor'. In: *Alchemy 2013 - Architecture, Languages, Compilation and Hardware support for Emerging Manycore systems*. Vol. 18. Proceedings of the International Conference on Computational Science, ICCS 2013. Barcelona, Spain, June 2013, pp. 1624–1633. DOI: [10.1016/j.procs.2013.05.330](https://doi.org/10.1016/j.procs.2013.05.330). URL: <https://hal.inria.fr/hal-00832504>.
- [0] T. Baudon, C. Fuhs and L. Gonnord. 'Parallel Complexity of Term Rewriting Systems'. In: *WST 2021 - 17th International Workshop on Termination*. Virtual, France, July 2021, pp. 1–6. URL: <https://hal.science/hal-03418400>.
- [0] T. Baudon, G. Radanne and L. Gonnord. *Compositional Flexible Memory Representations for Algebraic Data Types*. Tech. rep. 9495. Inria, Aug. 2022. URL: <https://hal.inria.fr/hal-03940742>.
- [0] D. Becker, M. Moy and J. Cornet. 'Parallel Simulation of Loosely Timed SystemC/TLM Programs: Challenges Raised by an Industrial Case Study'. In: *MDPI Electronics* 5.2 (2016). Ed. by F. Rousseau, G. Nicolescu, A. Baghdadi and M. Bassiouni, p. 22. DOI: [10.3390/electronics5020022](https://doi.org/10.3390/electronics5020022). URL: <https://hal.archives-ouvertes.fr/hal-01321055>.
- [0] D. Caromel and L. Henrio. *A Theory of Distributed Objects*. Springer-Verlag, 2004.
- [0] N. Courant and X. Leroy. 'Verified Code Generation for the Polyhedral Model'. In: *Proceedings of the ACM on Programming Languages* 5.POPL (Jan. 2021), 40:1–40:24. DOI: [10.1145/3434321](https://doi.org/10.1145/3434321). URL: <https://hal.archives-ouvertes.fr/hal-03000244>.
- [0] P. Cousot and R. Cousot. 'Abstract interpretation: A unified lattice model for static analysis of programs by construction or approximation of fixpoints'. In: *4th ACM Symposium on Principles of Programming Languages (POPL'77)*. Los Angeles, Jan. 1977, pp. 238–252.
- [0] F. De Boer, V. Serbanescu, R. Hähnle, L. Henrio, J. Rochas, C. C. Din, E. Broch Johnsen, M. Sirjani, E. Khamespanah, K. Fernandez-Reyes and A. M. Yang. 'A Survey of Active Object Languages'. In: *ACM Comput. Surv.* 50.5 (Oct. 2017), 76:1–76:39. DOI: [10.1145/3122848](https://doi.org/10.1145/3122848). URL: <http://doi.acm.org/10.1145/3122848>.
- [0] J. Emmanuel, M. Moy, L. Henrio and G. Pichon. 'Simulation of the Portals 4 protocol, and case study on the BXI interconnect'. In: *HPCS 2020 - International Conference on High Performance Computing & Simulation*. Barcelona, Spain, Dec. 2020, pp. 1–8. URL: <https://hal.archives-ouvertes.fr/hal-02972297>.
- [0] P. Feautrier. 'Dataflow analysis of array and scalar references'. In: *International Journal of Parallel Programming* 20.1 (1991), pp. 23–53.
- [0] P. Feautrier. 'Scalable and Structured Scheduling'. In: *International Journal of Parallel Programming* 34.5 (Oct. 2006), pp. 459–487.
- [0] P. Feautrier, A. Gamatié and L. Gonnord. 'Enhancing the Compilation of Synchronous Dataflow Programs with a Combined Numerical-Boolean Abstraction'. In: *CSI Journal of Computing* 1.4 (2012), 8:86–8:99. URL: <http://hal.inria.fr/hal-00860785>.
- [0] P. Feautrier and C. Lengauer. 'Polyhedron Model'. In: *Encyclopedia of Parallel Computing*. Springer, 2011, pp. 1581–1592.
- [0] K. Fernandez-Reyes, D. Clarke, E. Castegren and H.-P. Vo. 'Forward to a Promising Future'. In: *Conference proceedings COORDINATION 2018*. 2018.
- [0] K. Fernandez-Reyes, D. Clarke, L. Henrio, E. Broch Johnsen and T. Wrigstad. 'Godot: All the Benefits of Implicit and Explicit Futures'. In: *ECOOP 2019 - 33rd European Conference on Object-Oriented Programming*. Leibniz International Proceedings in Informatics (LIPIcs). London, United Kingdom, July 2019, pp. 1–28. URL: <https://hal.archives-ouvertes.fr/hal-02302214>.

- [0] M. I. Gordon. ‘Compiler techniques for scalable performance of stream programs on multicore architectures’. PhD thesis. Massachusetts Institute of Technology. Dept. of Electrical Engineering and Computer Science, 2010.
- [0] O. Hakjoo, L. Wonchan, H. Kihong, Y. Hongseok and Y. Kwangkeun. ‘Selective context-sensitivity guided by impact pre-analysis’. In: *ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI ’14, Edinburgh, United Kingdom - June 09 - 11, 2014*. ACM, 2014, p. 49.
- [0] N. Halbwachs, P. Caspi, P. Raymond and D. Pilaud. ‘The synchronous data flow programming language LUSTRE’. In: *Proceedings of the IEEE* 79.9 (Sept. 1991), pp. 1305–1320.
- [0] L. Henrio. *Data-flow Explicit Futures*. Research Report. I3S, Université Côte d’Azur, Apr. 2018. URL: <https://hal.archives-ouvertes.fr/hal-01758734>.
- [0] P. Iannetta, L. Gonnord and G. Radanne. ‘Compiling pattern matching to in-place modifications’. In: *GPCE 2021 - 20th International Conference on Generative Programming: Concepts & Experiences*. Chicago & Virtual, United States, Oct. 2021. DOI: [10.1145/3486609.3487204](https://doi.org/10.1145/3486609.3487204). URL: <https://hal.archives-ouvertes.fr/hal-03355377>.
- [0] P. Iannetta, L. Gonnord and G. Radanne. *Parallelizing Structural Transformations on Tarbres*. Research Report RR-9405. ENS Lyon, CNRS & INRIA, Apr. 2021, p. 21. URL: <https://hal.inria.fr/hal-03208466>.
- [0] *IEEE 1666 Standard: SystemC Language Reference Manual*. Open SystemC Initiative. 2011. URL: <http://www.accellera.org/>.
- [0] G. Kahn. ‘The semantics of a simple language for parallel programming’. In: *Information processing*. North-Holland, 1974.
- [0] A. Krizhevsky, I. Sutskever and G. E. Hinton. ‘Imagenet classification with deep convolutional neural networks’. In: *Advances in neural information processing systems*. 2012, pp. 1097–1105.
- [0] M. Maalej, V. Paisante, P. Ramos, L. Gonnord and F. Pereira. ‘Pointer Disambiguation via Strict Inequalities’. In: *Code Generation and Optimisation*. Austin, United States, Feb. 2017. URL: <https://hal.archives-ouvertes.fr/hal-01387031>.
- [0] M. Maalej Kammoun. ‘Low-cost memory analyses for efficient compilers’. Thèse de doctorat, Université Lyon1. PhD thesis. Université Lyon 1, 2017. URL: <http://www.theses.fr/2017LYSE1167>.
- [0] M. Moy. ‘Parallel Programming with SystemC for Loosely Timed Models: A Non-Intrusive Approach’. In: *DATE*. Grenoble, France, Mar. 2013, p. 9. URL: <https://hal.archives-ouvertes.fr/hal-00761047>.
- [0] *OSCI TLM-2.0 Language Reference Manual*. Open SystemC Initiative (OSCI). June 2008. URL: <http://www.accellera.org/downloads/standards>.
- [0] V. Paisante, M. Maalej, L. Barbosa, L. Gonnord and F. M. Q. Pereira. ‘Symbolic Range Analysis of Pointers’. In: *International Symposium of Code Generation and Optimization*. Barcelon, Spain, Mar. 2016, pp. 791–809. URL: <https://hal.inria.fr/hal-01228928>.
- [0] L.-N. Pouchet. *Polybench: The polyhedral benchmark suite*. 2012. URL: <http://www.cs.ucla.edu/~pouchet/software/polybench/>.
- [0] W. H. Press, S. A. Teukolsky, W. T. Vetterling and B. P. Flannery. ‘Numerical recipes in C++’. In: *The art of scientific computing* (2015).
- [0] P. Quinton. ‘Automatic synthesis of systolic arrays from uniform recurrent equations’. In: *ACM SIGARCH Computer Architecture News* 12.3 (1984), pp. 208–214.
- [0] H. Rihani, M. Moy, C. Maïza, R. I. Davis and S. Altmeyer. ‘Response Time Analysis of Synchronous Data Flow Programs on a Many-Core Processor’. In: *Proceedings of the 24th International Conference on Real-Time Networks and Systems*. RTNS ’16. Brest, France: ACM, 2016, pp. 67–76. DOI: [10.1145/2997465.2997472](https://doi.org/10.1145/2997465.2997472). URL: <http://doi.acm.org/10.1145/2997465.2997472>.
- [0] H. N. W. Santos, I. Maffra, L. Oliveira, F. Pereira and L. Gonnord. ‘Validation of Memory Accesses Through Symbolic Analyses’. In: *Proceedings of the 2014 ACM International Conference on Object Oriented Programming Systems Languages And Applications (OOPSLA’14)*. Portland, Oregon, United States, Oct. 2014. URL: <http://hal.inria.fr/hal-01006209>.

- [0] W. Thies. ‘Language and compiler support for stream programs’. PhD thesis. Massachusetts Institute of Technology, 2009.
- [0] J. Travis and J. Kring. *LabVIEW for everyone: graphical programming made easy and fun*. Prentice-Hall, 2007.
- [0] A. Turjan. ‘Compiling Nested Loop Programs to Process Networks’. PhD thesis. Universiteit Leiden, 2007.
- [0] N. Ventroux and T. Sassolas. ‘A new parallel SystemC kernel leveraging manycore architectures’. In: *Design, Automation & Test in Europe Conference & Exhibition (DATE), 2016*. IEEE. 2016, pp. 487–492.
- [0] S. Verdoolaege. ‘Polyhedral Process Networks’. In: *Handbook of Signal Processing Systems*. Springer, 2010, pp. 931–965.
- [0] S. Williams, A. Waterman and D. Patterson. ‘Roofline: an insightful visual performance model for multicore architectures’. In: *Communications of the ACM* 52.4 (2009), pp. 65–76.
- [0] P. Wilmott. *Quantitative Finance*. Wiley, 2006.
- [0] L. Xia, Y. Zakowski, P. He, C. Hur, G. Malecha, B. C. Pierce and S. Zdancewic. ‘Interaction Trees’. In: *Proceedings of the ACM on Programming Languages* 4.POPL (2020). DOI: [10.1145/3371119](https://doi.org/10.1145/3371119).