



ICT COST Action IC1405

COST Action IC1405
Reversible Computation
Extending Horizons of Computing

Working Group 4 — Report on Case Studies

Grant Period 2 report

Editors:

Ulrik Pagh Schultz and Carla Ferreira

Contributors:

Carla Ferreira, Stefan Kuhn, Ivan Lanese, Markus Schordan,
Ulrik Pagh Schultz, Harun Šiljak, Irek Ulidowski,
Robert Wille, and Alwin Zulehner

June 8, 2017

Contents

1	Overview	4
1.1	Introduction	4
1.2	Selected case studies	4
1.3	Research areas	4
2	Case studies	6
2.1	Debugging of concurrent and distributed systems	6
2.1.1	Conceptual description	6
2.1.2	Concrete example	6
2.1.3	Potential impact	7
2.1.4	Progress	7
2.2	Reliability of distributed systems through reversibility	7
2.2.1	Conceptual description	7
2.2.2	Concrete example	8
2.2.3	Potential impact	8
2.2.4	Progress	9
2.3	Reversibility in chemical reactions	9
2.3.1	Conceptual description	9
2.3.2	Concrete example	10
2.3.3	Potential impact	10
2.3.4	Progress	11
2.4	Reversible Control of Robots	11
2.4.1	Conceptual description	11
2.4.2	Concrete example	12
2.4.3	Potential impact	13
2.4.4	Progress	13
2.5	Simulator for optimistic parallel discrete event simulation	13
2.5.1	Conceptual description	13
2.5.2	Concrete example	14
2.5.3	Potential impact	15
2.5.4	Progress	15
2.6	Reversibility in Massive MIMO	15
2.6.1	Conceptual description	16
2.6.2	Concrete example	16
2.6.3	Potential impact	16
2.7	Synthesis of Encoder Circuits using Reversible Logic	17
2.7.1	Conceptual description	17
2.7.2	Concrete example	17
2.7.3	Potential impact	18

Preface

This report presents the case studies selected for investigation in the context of the COST Action IC 1405 on reversible computation, collected and edited by Working Group 4. The descriptions of the specific case studies have been provided by the researchers interested in investigating the corresponding case studies. The current status and recent development are described for each of the case studies.

More on the project:

- http://www.cost.eu/COST_Actions/ict/IC1405
- <http://www.revcomp.eu/>

1 Overview

1.1 Introduction

As outlined in the Memorandum of Understanding (MoU), the theories, techniques and support software tools delivered by Working Groups 1 to 3 will be validated and their suitability for practical applications will be assessed using case studies. The case studies will be used to provide feedback on the effectiveness of the reversibility-inspired theories, techniques, and solutions, and can point to new applications and research questions.

During the first year of the action five case studies have been selected, with two additional case studies added during the second year. The case studies are developed based on an interest-driven bottom-up approach, matching member research interests with case studies carrying a potential for external collaboration, either with industry or with researchers outside the area of reversible computing.

1.2 Selected case studies

Seven case studies have been selected based on interest from members of the action:

1. Debugging of concurrent and distributed systems
2. Reliability of distributed systems through reversibility
3. Reversibility in chemical reactions
4. Reversible control of robots
5. Simulator for optimistic parallel discrete event simulation
6. Reversibility in massive MIMO
7. Synthesis of encoder circuits using reversible logic

Case studies 1-5 were added at the end of the first year, case studies 6 and 7 at the end of the second year.

1.3 Research areas

The research areas covered by the selected case studies are outlined in Table 1. The areas are divided into Reversible Computing (RC) topics and non-RC topics. There is a good coverage of RC topics from Working Groups 1 and 2, but not currently any RC topics from Working Group 3. The proposed case studies on synthesis for reversible circuits and on reversible algorithms should however involve Working Group 3.

All case studies are based on significant interaction with researchers from areas outside reversible computing. Moreover, the case studies on debugging, robots, simulation, massive MIMO, and encoder circuits all have significant potential for industrial collaboration.

	RC topics	Non-RC topics
Debugging of concurrent and distributed systems	reversible language, checkpointing, causal-consistent reversibility, causal analysis	debugging, concurrent and distributed systems
Reliability of distributed systems through reversibility	reversible language, automatic recovery, reversible transactions, causal-consistent reversibility	concurrent and distributed systems, transactions, replication
Reversibility in chemical reactions	reversible language, reversible process calculi, out-of-causal order reversibility, local reversibility	modelling of chemical reactions, life sciences modelling
Reversible control of robots	reversible language, inversion, backtracking, reversible debugging, reversible concurrency	robotics, industrial robots, domain-specific languages, modular robots, swarm robots
Simulator for optimistic parallel discrete event simulation	reversible language, out-of-causal order reversibility	simulation, high-performance computing
Reversibility in massive MIMO (Multiple Input Multiple Output)	reversible control, reversible cellular automata, reversible petri nets	MIMO, communication networks, telecommunications, 5G, antenna selection
Synthesis of Encoder Circuits using Reversible Logic	reversible function representation, reversible circuits, design of reversible circuits, reversible circuit synthesis	automatic synthesis, encoder design, network on chip interconnects, coding techniques

Table 1: Research areas covered by the selected case studies

2 Case studies

We now present the seven case studies, in each case providing an overall conceptual description, a concrete example, an estimate of the potential impact, and an update (where applicable) on progress in the area.

2.1 Debugging of concurrent and distributed systems

Involved members: Ivan Lanese, Claudio Mezzina

2.1.1 Conceptual description

Finding bugs inside software has always been a main activity in the software development life cycle [BJC12], and reversibility can play a role here [Eng12, KDC05, GLM14]. In fact, reversibility allows one to go back in the past of the program, looking for the bug that caused a given misbehaviour. Indeed, many debuggers provide features for reversible execution, including GDB [JC09] and UndoDB [UDB]. However, current reversible debuggers either do not support concurrency/distribution, or they just fix an interleaving among threads and stick to it. It is not yet clear which is the good way to debug a concurrent and/or distributed system, what is potentially very interesting. Indeed, finding bugs inside concurrent and distributed software is more difficult than inside sequential software [LPS08], since faults may appear or disappear according to the speed of the different processes and of the network communications. The bugs generating these faults, called Heisenbugs, are particularly difficult to find. Note that, by exploring back and forward the same computation one is guaranteed that Eisenbugs that occurred will not disappear (provided that the approach preserves causal dependencies between actions from different threads).

2.1.2 Concrete example

The causal-consistent approach to reversibility [LMT14] allows one to undo any action provided that its consequences, if any, are undone beforehand. This relates reversibility to causality, instead of time, making it more suitable to distributed systems where a unique notion of time may not exist. In debugging, the causal-consistent approach allows one to find bugs by following backward causal dependencies from the misbehaviour to the bug causing it. This is supported by primitives allowing one to undo a past action, together with all and only its distributed consequences [LMS11]. For instance, if one observes that, e.g., the value of a variable x is 5 instead of the expected 6, one can undo the last assignment to x , reaching a state where the assignment could have been performed. A causal-consistent reversible debugger for the toy language μOz has been presented in [GLM14], and is available at [CRD].

2.1.3 Potential impact

Debugging is a topic very relevant from the practical point of view, but mostly neglected by academic research. Concerning impact, a 2012 study by the Judge Business School of the University of Cambridge, UK, sponsored by Rogue Wave Software, goes as far as to estimate that the time spent in debugging corresponds to 49.9% of total programming time [BJC12]. According to another recent study [US14], the cost of debugging software amounts to \$312 billions annually. Hence, any improvement in debugging techniques has potential huge practical impact. Given that systems are becoming more and more complex, and concurrency and distribution is more and more relevant, the relevance of debugging will not decrease in the close future.

There is a significant potential for industrial collaboration with Undo Software (Cambridge, UK), which is a company concentrated in reversible debugging. They are mainly concerned with sequential software, and in case of concurrency their approach is currently to consider a linearization of the execution. One of their main concerns and selling points is the limited overhead of their debugger, both in time and size of the logs. Still, they want to continue improving these features. Their approach is based on taking snapshots of the state at variable rates, according to how likely the current part of the execution will need to be reversed. Most of their work is in low-level optimizations and in providing support for different compilers and operating system configurations. Undo Software currently has a limited interest in concurrency and distribution, with the more relevant open points for them being parallel logging (since sequential logging requires to linearize execution, hence an unacceptable slow down), and debugging of MPI programs (since their customers are interested in this programming library).

2.1.4 Progress

There is ongoing work to make the *rollback* command typical of causal-consistent debugging available in a core subset of the Erlang language [NPV16]. The rollback command is also being investigated for the tuple-based distributed coordination language muKlaim [GLMT17]. For details we refer to the accompanying COST IC1405 Grant Period 2 report from *Working Group 2 — Software and Systems*.

2.2 Reliability of distributed systems through reversibility

Involved members: Carla Ferreira, Ivan Lanese, Claudio Mezzina, Vasileios Koutavas

2.2.1 Conceptual description

The emergence of global services, operating over massive distribution of data and computation, has lead to the recent paradigm of cloud computing. In this

global context building reliable systems is very hard, but reversibility can be used to achieve reliability in a natural way. If something goes wrong just undo the previous actions and return to a past consistent state. However, in these highly distributed systems there is no notion of global time, so there might not be a clear previous action. Instead, actions are causally ordered. Leveraging on [DK04], that defined causal-consistent reversibility, this case study will demonstrate the use of reversibility as a recovery mechanism for highly distributed systems. The goal is that a programmer can build a reversible distributed application for free, in the sense s/he just has to program the forward actions and whenever something goes wrong all actions can be reverted automatically. This was already explored in the context of long-running transactions (see [CP13] for an overview of calculi with support for compensations / backward actions). In long-running transactions, backward actions were defined by the programmer and the relation between forward and backward actions could be quite tenuous. In this case study, we would want to explore how backward actions can be determined automatically for concurrent and distributed systems over a shared state.

Another axis to be explored by the case study is the use of replicated databases to achieve service availability and scalability, as it is common practice in global services. In this setting, clients execute operations over the closest database replica and later these operations are propagated to other remote replicas. Ensuring automatic reversibility in this case is even harder as coordination is reduced (on weaker models of consistency coordination is avoided altogether). It should be explored how to ensure reversibility in this new context.

2.2.2 Concrete example

The concrete examples should include two classes of distributed systems. Distributed and concurrent systems that operate over a shared state, and distributed systems that operate over a replicated state. Several well known distributed systems case studies used in the literature can be used: transactional web server for an online store [TPC], an auction system, and a student registration system. These specific case studies should show the applicability and relevance of reversibility for distributed systems.

2.2.3 Potential impact

Building concurrent and distributed systems is hard and reversible languages can play an important role in helping programmers developing reliable software by providing reversibility by construction. Currently, error handling has to be done explicitly by the programmer, but if reversibility could be achieved by pressing a button it would have an impact on the quality of the programs developed. In fact, reversibility could be used as a way of providing speculative execution in cloud applications.

2.2.4 Progress

Investigation of the *rollback* command typical of causal-consistent debugging [GLMT17,NPV16] is also highly relevant for reliability of distributed systems, since it is key to implementing automated rollback of each node in the system. Complementary to this, automated recovery in distributed systems is being investigated from the point of view of reversibility of multiparty protocols [NY17]. For details we refer to the accompanying COST IC1405 Grant Period 2 report from *Working Group 2 — Software and Systems*.

2.3 Reversibility in chemical reactions

Involved members: Irek Ulidowski, Stefan Kuhn, Gabriel Ciobanu, Bogdan Aman

2.3.1 Conceptual description

Biological reactions, pathways and reactions networks have been studied using various techniques, including process calculi. The initial research focused on reaction rates, by modelling and simulating networks of reactions, in order to analyse or even predict the common paths through the network. Reversibility was not considered there explicitly. Once a more structural approach was taken reversibility started to be taken into account, since it plays a crucial role in many processes, typically by going back to a previous state in the system. For example, if a gene is regulated by a protein binding to and unbinding from it, then reversibility is needed if this to be modelled directly. Such examples can be modelled by backtracking along the path taken to the current state. Beyond backtracking and causal reversibility, there are more general forms of reversibility, specifically the out-of-causal-order reversibility, which make it possible to get to states which cannot be reached by forward reactions alone. Such sequences of forward and reversible reactions are important as they lead to new chemical structures and new reactions, which would not be possible if the out-of-causal reversibility was not used.

Looking at proteins and other macromolecules means working at relatively high levels of abstraction. It also means that such biological entities are often difficult to represent as there is still no complete understanding of their behaviour. Also their interactions with other similar molecules are hard to enumerate as they are sometime not fully understood. This could also mean that some aspects of reversibility are not clearly modelled since we have abstracted too much detail. Therefore, we suggest to consider chemical reactions when studying reversibility in the context of living systems.

Chemical reactions involve various forms of reversibility and offer good case studies for the modelling of reversibility. Specifically, we can find interesting examples of the out-of-causal order reversibility in many chemical reactions, and can study the intermediate steps of such reactions where some bonds are only helping to achieve the overall aim of the reaction: are only formed to be broken before the end such reactions. The level of detail we have about such

chemical reactions should help with seeing reversibility in action in the ways not obvious at higher levels. As an additional benefit, in many cases a good understanding of the underlying reaction mechanisms exists.

2.3.2 Concrete example

We suggest to study the following three types of reactions:

- The hydration of formaldehyde in water: This is a basic textbook reaction, but there are different paths through the reaction, some quite rare, and the reaction is reversible. Clearly backtracking is not sufficient to model this, since it is experimentally possible to go forward via one path and undo reactions via a different path.
- The Monsanto process: This is a typical catalytic cycle involving a metal catalyst. We have two input molecules and want to produce another molecule by combining them. They cannot interact directly because too much energy would be required. By introducing the metal catalyst this can be overcome the catalyst breaks the input molecules and enables their recombination. The catalyst is restored to its original form after the process making it a cycle. Clearly, there must be undoing of bonds created by the catalyst. In fact, some bonds are undone via the out-of-causal order computation.
- The glucose-fructose isomerization: This is a reaction forming part of the Krebs cycle, an important pathway in the metabolism of organisms. It is an enzyme-assisted catalytic process, where an enzyme interacts with the glucose to produce fructose. Similarly to the Monsanto process the protein emerges unchanged. The mechanism, however, involving both causal and out-of-causal order reversibility, is different.

These examples display various types of reversibility, including the out-of-causal order reversibility. All examples have been studied extensively and are well understood at a very detailed level. There are also aspects of reactions which may be included at a later stage, for example the energy levels of the various steps of reactions.

The aim is to model these reactions using different existing techniques. We currently envisage the use of the kappa-calculus [DL04], P systems [CDK03, Pau00] and reversible process calculi [DK03, LMS11, PU07], including the recently developed [KU15]. The case study should give a clearer view of what the individual calculi can do for each of the examples and clarify their strengths and weaknesses.

2.3.3 Potential impact

The work can help to identify work which needs to be done to get a wider range of reversible processes represented in computing. Furthermore it can help with building knowledge about chemical reactions. Regarding the potential for

external/industrial collaboration, there is interest in reaction simulation in the chemoinformatics area. A collaboration could be a long-term goal.

2.3.4 Progress

Recent progress in the area of reversibility in chemical reactions was presented at the COST IC1405 meeting in Belgrade. Based on previous work a process calculus called Calculus for Covalent Bonding (CCB) has been developed which can model many features found in chemical reactions [KU16]. The calculus has formal semantics given by SOS rules that extended with appropriate predicates. Simulation software was developed based on CCB and used to successfully model several examples of biochemical reactions of increasing complexity.

The reversibility of biochemical reactions in natural computing has been investigated using reaction systems, in a theory of natural computing inspired by the functioning of living cells [ER07]. A controlled form reversibility was considered with additional reversing rules added to the initial set of rules with inhibitors, as well an external control by means of a special symbol informing the system that a rollback is needed. Specific results (including loop results) were proved, including an operational correspondence between reaction systems and rewriting theory. This operational correspondence was the basis for an implementation of reversible reaction systems in the rewriting engine Maude; several software experiments were performed on rather simple examples.

2.4 Reversible Control of Robots

Involved members: Ulrik Pagh Schultz, Gabriel Ciobanu

2.4.1 Conceptual description

Robots normally have one or more degrees of freedom controlled by a computational process; using reversible computing to control the robot potentially gives rise to new reverse behaviours. For example, major industrial robot manufacturers such as ABB and KUKA offer limited forms of ad-hoc reverse execution for interactive programming and debugging, but due to limitations in the underlying execution models, their programming models are incapable of reversing complex actions such as steps of an assembly process [LSE15, MAH10]. We attribute the ad-hoc limitations to the lack of an underlying reversible model. The first investigation of fully reversible robot behaviours was for self-reconfigurable robots [SBS11]. To better understand the underlying relation between reversible computation and physical reversibility, we will investigate reversible control of industrial robots and distributed robots.

Programming industrial robots is challenging due to the difficulty of precisely specifying general yet robust operations. As the complexity of these operations increases, so does the likelihood of errors. Certain classes of errors during industrial robot operations can however be addressed using reverse

execution, allowing the robot to temporarily back out of an erroneous situation, after which the operation can be automatically retried. Specifically, this approach has been shown to be useful for automatic error recovery for small-sized batch production of assembly operations [SLE15]. Moreover, reversibility can in this case be used to automatically derive a disassembly sequence from a given assembly sequence, or vice versa. These results were demonstrated using an initial design and implementation of a reversible domain-specific languages (DSL) for specifying such assembly sequences [SLE15, LSE15]. The area however remains largely unexplored, both from a theoretical and practical point of view. There is for example a large design space for different programming language approaches, both in terms of the generality of the language and the means by which reversibility is achieved. At a more fundamental level, the notion of reversible control of a reversible physical system remains largely unexplored. From a practical point of view, only the specific case of assembly operations has been investigated, and only using a specific set of industrial use cases. There has been no attempt at integration into an existing robotics platform, although we observe that many existing platforms offer limited notions of reversibility for using during programming and debugging.

Distributed robotic systems can make use of reverse execution as a means to reversing physical actions or correcting errors. As a concrete example, previous work has studied self-reconfigurable, modular robots which are distributed mechatronic devices that can autonomously change their physical shape. Self-reconfiguration from one shape to another is typically achieved through a specific sequence of actuation operations distributed across the modules of the robot. Automatically reversing the sequence of operations brings the robot back to its initial shape, which has been experimentally demonstrated using the DynaRole reversible language [SBS11]. DynaRole however only allows simple sequences of operations to be reversed, which is suitable for reversing self-reconfiguration sequences but lacks the generality needed to implement more complex behaviors. Initial ideas on generalizing the DynaRole language to support a wider range of scenarios, while retaining the possibility of reversing distributed sequences, have been proposed [Sch12]. Reversibility is used here as a practical feature, reducing the programming task of the programmer, and allowing error recovery by backing out of an error state using reverse execution. Currently there however is no underlying model allowing a distributed robotic systems to use reversibility for backing out of distributed operations that have become blocked, e.g., due to unexpected features of the environment, to under conflicting operations, or recover from partial hardware failures.

2.4.2 Concrete example

An appropriately robust and useful reversible DSL for control of industrial robots could be integrated into an existing robotics platform, and could be useful both for working interactively with the robot (programming and debugging) but also for error recovery during specific kinds of operations. Note that many robotic systems already integrate limited forms of reverse execution for interac-

tive debugging [MAH10], better understanding and improving these approaches to interactive debugging is probably interesting.

Using reversible computing to allow a subset of a distributed robotic system to optimistically proceed with an operation, but then backtrack and undo the operation if the system globally decides not to perform this operation. This is interesting not only for modular robots but also for swarm robotic system e.g. drones. A single mobile robot navigating an environment can be modeled as a 2D Turing machine [Tho14], presumably a useful model for multi-robot systems navigating an environment could be based on a similar model.

2.4.3 Potential impact

From a scientific point of view, the design of reversible languages for controlling reversible processes remains an open question. Moreover, exploring reversible control of physical processes may help to better understand the principles of reversible computing. From a society point of view, industrial robots are key to maintaining production in Europe, and reversible computation has the potential to increase robustness for specific kinds of operations such as assembly, and moreover facilitate the programming of such operations.

There is a significant potential for external/industrial collaboration for industrial robots: Universal Robots ApS is a potential industrial collaborator, as are potentially also other robotics companies or perhaps companies performing large numbers of error-prone assembly operations. For modular robots, we however note that swarm and modular robots are mostly at the level of basic research in robotics, with the obvious exception of drones (but here reversibility seems less useful).

2.4.4 Progress

Significant progress in the area of reversibility for industrial robots was presented at COST IC-1405 meeting in Belgrade. Key developments include an improved understanding of the interaction between reversible computing and real-world systems that only are partially reversible, as well a substantial experimental evaluation of the use of reversible languages to control industrial robots performing assembly and disassembly in the context of small-batch production.

2.5 Simulator for optimistic parallel discrete event simulation

Involved members: Markus Schordan

2.5.1 Conceptual description

Discrete event simulation (DES) is a simulation paradigm suitable for systems whose states are modeled as changing discontinuously and irregularly at discrete

moments of simulation time. DES is event-driven in that the times at which state changes occur are calculated dynamically rather than statically as in time-stepped simulations. Example applications include simulations of digital communication networks, vehicular traffic flow, markets and economies, epidemiological models, logistical models, ecological and population models, tactical and strategic military models, and many others. Most systems whose behavior is not describable by continuous equations and that are not suitable for simple time-stepped models are candidates for DES.

Reverse computation has become a central notion in parallel discrete event simulation (PDES) over the last decade [CPF99, Per13]. It is not just a theoretical line of research, but an immensely practical one necessary to achieve high performance for large parallel discrete event simulations. In fact, the most highly parallel and fastest discrete event simulation benchmarks ever executed have made essential use of reverse computation [BCJ13].

The fundamental issue that makes PDES so complex is the synchronization problem. Every logical process (LP) must execute all events from its own event queue in strictly non-decreasing timestamp order despite the fact that it does not generally know which LPs might be sending it events, or how many, or what timestamps they may carry. Furthermore there is no guarantee that event messages will arrive at an LP in increasing timestamp order. Optimistic synchronization is that all event executions are speculative or provisional, and are always subject to rollback if the simulation gets into local synchronization trouble. Most of the time that does not happen and the simulation proceeds forward in parallel, but occasionally a causality violation occurs that has to be corrected by rollback (achieved by reversible computation) and then the simulation continues forward again. Therefore, reversible computation is paramount for achieving high-performance in optimistic parallel discrete event simulations on supercomputers. Various approaches to reversibility have been developed over the years and have they been successfully applied in supporting optimistic PDES, with recent contributions in incremental state-saving [MOJ16, SJB15], reverse code generation [PP14, LGC14], utilization of transactional memory [SIP15], and hybrid combinations of these approaches [PVQ15, CPQ15].

2.5.2 Concrete example

We suggest to use an existing simulator as basis for optimistic PDES which allows to plug-in reverse code for implementing the rollback as required for optimistic PDES. A well known simulator is ROSS (Rensselaer’s Optimistic Simulation System), a general purpose discrete event simulator developed at RPI by C. Carothers et. al. [HC08].

ROSS helped to demonstrate that Time Warp simulators can be run efficiently both in terms of speed and memory usage relative to a high-performance sequential simulator and was used to achieve the world-record in 2013 for the most highly parallel and fastest discrete event simulation benchmarks ever executed [BCJ13]. In this strong scaling study of the ROSS simulator running Time Warp with reverse computation, had 251 million PHOLD logical processes and

was executed in several configurations with up to a peak of 7.86 million MPI tasks running on 1,966,080 cores. At the largest scale it processed 33 trillion events in 65 seconds, yielding a sustained speed of 504 billion events/second using 120 racks of Sequoia. This is by far the highest event rate reported by any parallel discrete event simulation to date, whether running PHOLD or any other benchmark.

In order to use Time Warp in a ROSS model, a reverse event function must be provided. For a given event message it is responsible for undoing the state changes that the forward event function incurred for the same message. ROSS is open-source and also comes with a number of models (implemented in C/C++) with efficient hand-written reverse functions. These can be used as a basis for comparing different automated approaches to reversibility and perform comparative evaluations with the existing reversible ROSS implementations.

2.5.3 Potential impact

It is important to investigate how different approaches to reversibility can attack the problem of making the execution of a model’s event function reversible, because any overhead introduced in the forward execution or reverse computation (dependent on the approach to reversibility) impacts the performance of the simulation. In particular, the combination of approaches is interesting for large-scale simulation models, as they are usually a combination of components with different computational properties and challenges for reversibility. If successful, the impact would be that this can bring also other approaches (from the COST project) to be used — or at least to be evaluated to be suitable — to HPC supercomputers, and may also help to improve the performance of simulations. In particular, also reversible languages such as Janus [YAG08] that can be translated to C are candidates for integration with other approaches, and can be directly evaluated and compared to existing models utilizing reverse code generation. There is potential for industrial collaboration with Lawrence Livermore National Laboratory.

2.5.4 Progress

The Backstroke compiler for generating reversible program from C++ was released to the public in March 2017 (version 2.1.0). This was the first public release of Backstroke V2 using incremental state saving. Moreover the combination of incremental state saving (i.e., using Backstroke) with Janus-generated reversible C code was investigated during an STSM by Michael Kirkedal visiting Marchus Schordan at Lawrence Livermore National Laboratory. A recent publication addresses instrumentation of binary codes [CPS17].

2.6 Reversibility in Massive MIMO

Involved members: Harun Šiljak

2.6.1 Conceptual description

Massive MIMO (Multiple Input Multiple Output) systems have been introduced less than a decade ago as a theoretical concept in telecommunications [Mar10], but since then it has entered the practical realm as well and established itself as one of primary technologies for the coming 5G framework [BHL14]. In brief, it is a concept expanding on classical MIMO by using tens and hundreds of antennas which may be distributed in the environment (distributed massive MIMO) or placed on a single panel (collocated massive MIMO).

Since massive MIMO calls for solving several control and optimisation problems in its application, we aim to introduce reversible computing in the massive MIMO context through the concept of reversible control. In addition to that, since modelling of massive MIMO and networks created by it is of significant interest, reversibility might be of use there as well. In particular, reversible cellular automata [Mor08] and reversible Petri nets [BKM16] will be investigated as a way of massive MIMO modelling, because of their flexibility and novelty. In case of massive MIMO, ability to keep track of reversible cellular automaton's history, undo actions and run it backwards can prove useful in a dynamical environment (periodic changes in environment or mobile station movements benefiting from running antenna selection backwards, algorithms defined on a higher level of abstraction benefiting from undo in case of cost function decrease).

Another natural question to pose is the relationship between reversibility and transmitter-receiver duality [VT03] and the question of its applicability in practice.

2.6.2 Concrete example

There have been some recent efforts in solving the problem of antenna selection in massive MIMO [HBD13,GET15] using cellular automata (work in progress by Harun Šiljak, proposer of this case study). Due to the dynamical nature of the environment and movement of user equipment in it, it would only be natural to allow these selection algorithms to make a step backward or a step forward from its current state to suit the needs of the user equipment in a new configuration which may have occurred earlier as well. This reversible control problem could be, for instance, solved by moving from classical cellular automaton modelling the antenna selection process to a reversible one.

2.6.3 Potential impact

There are several different potential impacts of this study. First of all, placing massive MIMO (and communication networks in general [MGM16]) and reversibility in the context of complex systems science will help us investigate relationship between emergent behaviour and reversibility, a question which has drawn some interest in the past as well [LAX14,HSC11]. Second, the interest for reversibility in telecommunications is mostly spurred by the need for energy efficiency [MA17,GYG17] offered by reversible circuits. This study will offer a

fresh perspective on reversibility in this context. Last, a third benefit is the possible direct field application of reversible control algorithms for massive MIMO prototypes enhancing their performance.

2.7 Synthesis of Encoder Circuits using Reversible Logic

Involved members: Alwin Zulehner, Robert Wille

2.7.1 Conceptual description

Encoding devices represent a vital part of numerous applications realized in today’s electronic systems such as addressing memories and caches, data demultiplexing, etc. (see [PD96,BMM98]). With the rise of *System on Chip* and *Network on Chip* architectures, they gained further importance by the fact that those architectures usually rely on a rather sophisticated interconnect solutions [BMM98,GGO11,OIM09,LLY06]: In order to address the underlying limitations with respect to performance and/or power efficiency, communications between components of those architectures is usually conducted by dedicated encodings instead of the originally provided raw data. All these applications eventually motivate the design automation of corresponding encoders, i.e. the automatic synthesis of circuits which transform an n -bit word into another n -bit word while, at the same time, guarantee a one-to-one mapping from inputs to outputs.

However, design solutions for the automatic synthesis of corresponding circuits suffer from various drawbacks, e.g. they are often not scalable, do not exploit the full degree of freedom, or are applicable to realize certain codes only. All these problems are caused by the fact that existing design solutions have to explicitly guarantee a one-to-one mapping.

This limitation can be tackled using reversible logic, since one-to-one mappings are inherently guaranteed by this paradigm. First results have been accomplished for encoders in general in [ZW17] as well as for low-power on-chip interconnects in [WDO12,WKH16].

2.7.2 Concrete example

Since specifying all input output mappings for encoders with large bit-width is intractable, *Incomplete encodings* (also known as *discretized* or *approximate* encodings; see [BMP00] have been considered where the desired input/output mapping is specified for a (non-exponential) selection of the patterns only (e.g. the most important ones). For all other input patterns, the corresponding output patterns are considered don’t care.

In many domains such as the design of low-power encoders, not an explicit input/output mapping is desired, but only one which satisfies a dedicated objective. For example, specifications derived from probability-based mapping [RSH99] only require that certain inputs are getting mapped to outputs of a dedicated Hamming weigh.

In our work we consider the cases discussed above, which both allow for a high degree of freedom in the resulting realization. This degree of freedom can be exploited using reversible logic since one-to-one mappings are inherently guaranteed.

2.7.3 Potential impact

We aim for a significant reduction in the area as well as in the power consumption of encoder circuits by exploiting the fact that we can take one-to-one mappings for granted if we use reversible logic. A reduction of hardware is always desirable, but especially in the low-power domain, where small encoders are essential to exploit the benefits low-power codings. Obviously, the larger the encoder circuit, the more power is consumed to realize the codings.

Using a first implementation of the proposed design methodology [ZW17], we have already observed a substantial reduction of the number of gates in the encoder circuits. Taking the one-to-one mapping for granted by using reversible logic allows to reduce the number of gates (the reversible circuit is mapped to conventional technology after synthesis) of up to 92% (approx. 60% on average) compared to previously proposed methods. We note that in the domain of low-power on-chip interconnect encoders a substantial reduction of power consumption has already been observed [WDO12,WKH16].

References

- [BCJ13] Barnes, Jr., P.D., Carothers, C.D., Jefferson, D.R., LaPre, J.M.: Warp speed: Executing time warp on 1,966,080 cores. In: Proceedings of the 2013 ACM SIGSIM Conference on Principles of Advanced Discrete Simulation. SIGSIM-PADS '13, New York, NY, USA, ACM (2013) 327-336.
- [BJC12] T. Britton, L. Jeng, G. Carver, P. Cheak, and T. Katzenellenbogen. Reversible debugging software - quantify the time and cost saved using reversible debuggers. <http://www.roguewave.com>, 2012.
- [BHL14] F. Boccardi, R. W. Heath, A. Lozano, T. L. Marzetta, and P. Popovski. Five disruptive technology directions for 5g. *IEEE Communications Magazine*, 52(2):74–80, February 2014.
- [BKM16] Kamila Barylska, Maciej Koutny, Łukasz Mikulski, and Marcin Piatkowski. Reversible Computation vs. Reversibility in Petri Nets. In *Reversible Computation*, pages 105–118. Springer, Cham, July 2016.
- [BMP00] L. Benini, A. Macii, M. Poncino, and R. Scarsi. Architectures and synthesis algorithms for power-efficient businterfaces. *IEEE Trans. on CAD of Integrated Circuits and Systems*, 19(9):969–980, 2000.

- [**BMM98**] L. Benini, G. D. Micheli, E. Macii, M. Poncino, and S. Quer. Power optimization of core-based systems by address bus encoding. *Trans. VLSI Syst.*, 6(4):554–562, 1998.
- [**BMS98**] L. Benini, G. D. Micheli, D. Sciuto, E. Macii, and C. Silvano. Address bus encoding techniques for system-level power optimization. In *Design, Automation & Test in Europe Conference & Exhibition (DATE)*, pages 861–866, 1998.
- [**CDK03**] G. Ciobanu, R. Desai and A. Kumar. Membrane systems and distributed computing. In *Membrane Computing. International Workshop, WMC-CdeA 2002 Curtea de Arges 2002*, pages 187-202. Springer, 2003.
- [**CP13**] Christian Colombo and Gordon J. Pace. 2013. Recovery within long-running transactions. *ACM Comput. Surv.* 45, 3, pages 28:1 28:35
- [**CPF99**] Carothers, C.D., Perumalla, K.S., Fujimoto, R.M.: Efficient optimistic parallel simulations using reverse computation. *ACM Trans. Model. Comput. Simul.* 9(3) (July 1999) 224-253.
- [**CPQ15**] D. Cingolani, A. Pellegrini, and F. Quaglia. Transparently mixing undo logs and software reversibility for state recovery in optimistic pdes. In *Proceedings of the 3rd ACM SIGSIM Conference on Principles of Advanced Discrete Simulation, SIGSIM PADS '15*, pages 211-222, New York, NY, USA, ACM, 2015.
- [**CPS17**] Dealing with Reversibility of Shared Libraries in PDES Davide Cingolani, Alessandro Pellegrini, Markus Schordan, Francesco Quaglia, David R. Jefferson. *Proceedings of the 2017 annual ACM Conference on SIGSIM Principles of Advanced Discrete Simulation, SIGSIM-PADS 2017*, NIU, Singapore, May 24-26, 2017, ACM. To appear.
- [**CRD**] CaReDeb <http://www.cs.unibo.it/caredeb>
- [**DK03**] V. Danos and J. Krivine. Formal molecular biology done in CCS-R. In *Proceedings of the 1st Workshop on Concurrent Models in Molecular Biology BioConcur 2003*, volume 180 of ENTCS, pages 3149, 2007.
- [**DK04**] V. Danos and J. Krivine. Reversible communicating systems. In *CONCUR04*, volume 3170 of LNCS. Springer, 2004.
- [**DL04**] V. Danos and C. Laneve. Formal molecular biology. *Theoretical Computer Science*, 325(1):69110, 2004.
- [**Eng12**] J. Engblom, A review of reverse debugging, in: *System, Software, SoC and Silicon Debug*, IEEE, 2012, pp. 1-6.
- [**ER07**] A. Ehrenfeucht and G. Rozenberg. Reaction Systems. *Fundamenta Informaticae* 75(1), 263280 (2007).

- [**GGO11**] A. Garca-Ortiz, D. Gregorek, and C. Osewold. Optimization of interconnect architectures through coding: A review. In *Electronics, Communications and Photonics Conference*, pages 1–6, April 2011.
- [**GET15**] X. Gao, O. Edfors, F. Tufvesson, and E. G. Larsson. Massive MIMO in Real Propagation Environments: Do All Antennas Contribute Equally? *IEEE Transactions on Communications*, 63(11):3917–3928, November 2015.
- [**GLM14**] E. Giachino, I. Lanese, C. A. Mezzina: Causal-Consistent Reversible Debugging. FASE 2014: 370-384
- [**GLMT17**] E. Giachino, I. Lanese, C.A. Mezzina, and F. Tiezzi. Causal-consistent rollback in a tuple-based language. *J. Log. Algebr. Meth. Program.* Volume 88, pages 99–120. 2017.
- [**GYG17**] Xiaohu Ge, Jing Yang, Hamid Gharavi, and Yang Sun. Energy Efficiency Challenges of 5g Small Cell Networks. *arXiv:1702.03503 [cs]*, February 2017. arXiv: 1702.03503.
- [**HBD13**] J. Hoydis, S. ten Brink, and M. Debbah. Massive MIMO in the UL/DL of Cellular Networks: How Many Antennas Do We Need? *IEEE Journal on Selected Areas in Communications*, 31(2):160–171, February 2013.
- [**HC08**] A. O. Holder and C. D. Carothers. Analysis of time warp on a 32,768 processor IBM Blue Gene/L supercomputer. In Proceedings of the European Modeling and Simulation Symposium (EMSS), 2008.
- [**HSC11**] H. Hamann, T. Schmickl, and K. Crailsheim. Thermodynamics of emergence: Langton’s ant meets Boltzmann. In *2011 IEEE Symposium on Artificial Life (ALIFE)*, pages 62–69, April 2011.
- [**JC09**] J. Conrod. Tutorial: Reverse debugging with GDB 7, 2009 <http://jayconrod.com/posts/28/tutorial-reverse-debugging-with-gdb-7>
- [**KDC05**] S. T. King, G. W. Dunlap, and P. M. Chen. Debugging operating systems with time-traveling virtual machines. In USENIX Annual Technical Conference, General Track, pages 1-15, 2005.
- [**KU15**] S. Kuhn, and I. Ulidowski. Towards modelling of local reversibility. In Proceedings of Reversible Computation 2015, pages 279–284. LNCS 9138. Springer, 2015.
- [**KU16**] S. Kuhn, and I. Ulidowski. A Calculus for Local Reversibility. In Proceedings of Reversible Computation 2016, pages 20–35. LNCS 9720. Springer, 2016.
- [**LAX14**] Jia Lee, Susumu Adachi, Yun-Ni Xia, and Qing-Sheng Zhu. Emergence of universal global behavior from reversible local transitions in asynchronous systems. *Information Sciences*, 282:38–56, October 2014.

- [**LGC14**] J. M. LaPre, E. J. Gonsiorowski, and C. D. Carothers. Lorain: A step closer to the pdes 'holy grail'. In Proceedings of the 2Nd ACM SIGSIM/PADS Conference on Principles of Advanced Discrete Simulation, SIGSIM-PADS '14, pages 3–14, New York, NY, USA, 2014. ACM.
- [**LLY06**] K. Lee, S. Lee, and H. Yoo. Low-power network-on-chip for high-performance soc design. *IEEE Trans. VLSI Syst.*, 14(2):148–160, 2006.
- [**LMS11**] I. Lanese, C.A. Mezzina, A. Schmitt, and J-B. Stefani. Controlling reversibility in Higher-Order Pi. In Proceedings of CONCUR 2011, pages 297-311, volume 6901 of Lecture Notes
- [**LMT14**] I. Lanese, C. A. Mezzina, F. Tiezzi: Causal-Consistent Reversibility. Bulletin of the EATCS 114 (2014)
- [**LPS08**] Shan Lu, Soyeon Park, Eunsoo Seo, Yuanyuan Zhou: Learning from mistakes: a comprehensive study on real world concurrency bug characteristics. ASPLOS 2008: 329-339
- [**LSE15**] J.S. Laursen, U.P. Schultz, and L.-P. Ellekilde, "Automatic error recovery in robot assembly operations using reverse execution," in International Conference on Intelligent Robots and Systems (IROS 2015), IEEE/RSJ, 2015.
- [**MA17**] A. Mammela and A. Anttonen. Why Will Computing Power Need Particular Attention in Future Wireless Devices? *IEEE Circuits and Systems Magazine*, 17(1):12–26, 2017.
- [**MAH10**] Henrik Mhe and Andreas Angerer and Alwin Hoffmann and Wolfgang Reif. On reverse-engineering the KUKA Robot Language, 1st International Workshop on Domain-Specific Languages and models for ROBotic systems. 2010.
- [**Mar10**] T. L. Marzetta. Noncooperative Cellular Wireless with Unlimited Numbers of Base Station Antennas. *IEEE Transactions on Wireless Communications*, 9(11):3590–3600, November 2010.
- [**MGM16**] Irene Macaluso, Carlo Galiotto, Nicola Marchetti, and Linda Doyle. A complex systems science perspective on wireless networks. *Journal of Systems Science and Complexity*, 29(4):1034–1056, August 2016.
- [**MOJ16**] Markus Schordan, Tomas Opielstrup, David Jefferson, Peter D. Barnes Jr., Daniel J Quinlan. Automatic Generation of Reversible C++ Code and Its Performance in a Scalable Kinetic Monte-Carlo Application. SIGSIM-PADS 2016: 111-122, 2016.
- [**Mor08**] Kenichi Morita. Reversible computing and cellular automata—A survey. *Theoretical Computer Science*, 395(1):101–131, April 2008.

- [**NPV16**] Naoki Nishida, Adrián Palacios, Germán Vidal. A Reversible Semantics for Erlang. Proc. of the 26th International Symposium on Logic-Based Program Synthesis and Transformation, LOPSTR 2016. LNCS 10184, Springer. To appear.
- [**NY17**] Rumyana Neykova and Nobuko Yoshida. Let It Recover: Multiparty Protocol-Induced Recovery. 26th International Conference on Compiler Construction (CC), pages 98–108. ACM. 2017.
- [**OIM09**] A. G. Ortiz, L. S. Indrusiak, T. Murgan, and M. Glesner. Low-power coding for networks-on-chip with virtual channels. *J. Low Power Electronics*, 5(1):77–84, 2009.
- [**PD96**] P. R. Panda and N. D. Dutt. Reducing address bus transitions for low power memory mapping. In *European Design and Test Conference, ED&TC*, pages 63–71, 1996.
- [**Per13**] Perumalla, K.S.: Introduction to Reversible Computing, Chapman and Hall/CRC, 325 Pages, ISBN 9781439873403, 2013.
- [**PP14**] K. S. Perumalla and A. J. Park. Reverse computation for rollback-based fault tolerance in large parallel systems. *Cluster Computing*, 17(2):303–313, June 2014.
- [**PU07**] I.C.C. Phillips and I. Ulidowski. Reversing algebraic process calculi. *Journal of Logic and Algebraic Programming*, 73:7096, 2007.
- [**Pau00**] Gh. Paun. Computing with Membranes, *Journal of Computer and System Sciences*, 61(1): 108-143, 2000.
- [**PVQ15**] A. Pellegrini, R. Vitali, and F. Quaglia. Autonomic state management for optimistic simulation platforms. *IEEE Transactions on Parallel and Distributed Systems*, 26(6):1560–1569, June 2015.
- [**RSH99**] S. Ramprasad, N. R. Shanbhag, and I. N. Hajj. A coding framework for low-power address and data busses. *IEEE Trans. VLSI Syst.*, 7(2):212–221, 1999.
- [**SBS11**] Schultz, Ulrik, Mirko Bordignon, and Kasper Stoy. "Robust and reversible execution of self-reconfiguration sequences." *Robotica* 29.01 (2011): 35-57.
- [**SIP15**] E. Santini, M. Ianni, A. Pellegrini, and F. Quaglia. Hardware-transactional-memory based speculative parallel discrete event simulation of very ne grain models. In 2015 IEEE 22nd International Conference on High Performance Computing (HiPC), pages 145-154, Dec 2v015.
- [**SJB15**] Reverse Code Generation for Parallel Discrete Event Simulation. Markus Schordan, David Jefferson, Peter Barnes, Tomas Ooppelstrup, Daniel Quinlan. In Proceedings of the 7th Conference on Reversible Computation.

- Jean Krivine and Jean-Bernard Stefani (Eds.): Reversible Computation, LNCS 9138, pp. 95-110, ISBN 978-3-319-20859-6, Springer, 2015.
- [SLE15] U.P. Schultz, J.S. Laursen, L.-P. Ellekilde, and H.B. Axelsen, "Towards a domain-specific language for reversible assembly sequence," in Reversible Computation, 2015.
- [Sch12] Schultz, Ulrik. "Towards a general-purpose, reversible language for controlling self-reconfigurable robots." Reversible Computation. Springer Berlin Heidelberg, 2012. 97-111.
- [TPC] Overview of the TPC-C Benchmark. <http://www.tpc.org/tpcc/detail.asp>
- [Tho14] Mobile Robot API for Janus. Master thesis defence by Martin Nyborg Thomsen. University of Copenhagen, Denmark, 2014.
- [UDB] UndoDB. Commercial debugger <http://undo-software.com/products/undodb/>
- [US14] Undo Software. Increasing software development productivity with reversible debugging, 2014. in Computer Science, pages 297311. Springer, 2011.
- [VT03] P. Viswanath and D. N. C. Tse. Sum capacity of the vector Gaussian broadcast channel and uplink-downlink duality. *IEEE Transactions on Information Theory*, 49(8):1912–1921, August 2003.
- [WDO12] R. Wille, R. Drechsler, C. Osewold, and A. G. Ortiz. Automatic design of low-power encoders using reversible circuit synthesis. In *Design, Automation & Test in Europe Conference & Exhibition (DATE)*, pages 1036–1041, 2012.
- [WKH16] R. Wille, O. Keszocze, S. Hillmich, M. Walter, and A. G. Ortiz. Synthesis of approximate coders for on-chip interconnects using reversible logic. In *Design, Automation & Test in Europe Conference & Exhibition (DATE)*, pages 1140–1143, 2016.
- [ZW17] A. Zulehner and R. Wille. Taking one-to-one mappings for granted: Advanced logic design of encoder circuits. In *Design, Automation and Test in Europe*, 2017.
- [YAG08] Tetsuo Yokoyama, Holger Bock Axelsen, and Robert Glck. 2008. Principles of a reversible programming language. In Proceedings of the 5th conference on Computing frontiers (CF '08). ACM, New York, NY, USA, 43-54.