



ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

ARCHIVIO ISTITUZIONALE DELLA RICERCA

Alma Mater Studiorum Università di Bologna Archivio istituzionale della ricerca

Towards a Taxonomy for Reversible Computation Approaches

This is the final peer-reviewed author's accepted manuscript (postprint) of the following publication:

Published Version:

Glück R., Lanese I., Mezzina C.A., Miszczak J.A., Phillips I., Ulidowski I., et al. (2023). Towards a Taxonomy for Reversible Computation Approaches [10.1007/978-3-031-38100-3_3].

Availability:

This version is available at: <https://hdl.handle.net/11585/953336> since: 2024-01-17

Published:

DOI: http://doi.org/10.1007/978-3-031-38100-3_3

Terms of use:

Some rights reserved. The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (<https://cris.unibo.it/>).
When citing, please refer to the published version.

(Article begins on next page)

This is the final peer-reviewed accepted manuscript of:

Glück, R. *et al.* (2023). Towards a Taxonomy for Reversible Computation Approaches. In: Kutrib, M., Meyer, U. (eds) Reversible Computation. RC 2023. Lecture Notes in Computer Science, vol 13960. Springer, Cham.

The final published version is available online at: https://doi.org/10.1007/978-3-031-38100-3_3

Terms of use:

Some rights reserved. The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (<https://cris.unibo.it/>)

When citing, please refer to the published version.

Towards a Taxonomy for Reversible Computation Approaches^{*}

Robert Glück¹[0000-0001-6990-3935], Ivan Lanese²[0000-0003-2527-9995], Claudio Antares Mezzina³[0000-0003-1556-2623], Jarosław Adam Miszczak⁴[0000-0001-8790-101X], Iain Phillips⁵[0000-0001-5013-5876], Irek Ulidowski^{6,7}[0000-0002-3834-2036], and Germán Vidal⁸[0000-0002-1857-6951]

¹ DIKU, Department of Computer Science, University of Copenhagen, Denmark

² Focus Team, University of Bologna/INRIA, Italy

³ Dipartimento di Scienze Pure e Applicate, Università di Urbino, Urbino, Italy

⁴ Institute of Theoretical and Applied Informatics, Polish Academy of Sciences, Gliwice, Poland

⁵ Imperial College London, England

⁶ Department of Applied Informatics, AGH, Kraków, Poland,

⁷ SCMS, University of Leicester, England

⁸ MIST, VRAIN, Universitat Politècnica de València, Spain

Abstract. Reversible computation is a paradigm allowing computation to proceed not only in the usual, forward direction, but also backwards. Reversible computation has been studied in a variety of models, including sequential and concurrent programming languages, automata, process calculi, Turing machines, circuits, Petri nets, event structures, term rewriting, quantum computing, and others. Also, it has found applications in areas as different as low-power computing, debugging, simulation, robotics, database design, and biochemical modeling. Thus, while the broad idea of reversible computation is the same in all the areas, it has been interpreted and adapted to fit the various settings. The existing notions of reversible computation however have never been compared and categorized in detail. This work aims at being a first stepping stone towards a taxonomy of the approaches that co-exist under the term reversible computation. We hope that such a work will shed light on the relation among the various approaches.

Keywords: Reversible computing · Taxonomy · Models and languages.

^{*} I. Lanese has been partially supported by French ANR project DCore ANR-18-CE25-0007 and INdAM-GNCS Project CUP_E55F22000270001 “Proprietà qualitative e quantitative di sistemi reversibili”. I. Ulidowski has been partially supported by JSPS Fellowship grant S21050. G. Vidal has been partially supported by grant PID2019-104735RB-C41 funded by MCIN/AEI/ 10.13039/501100011033. J.A. Miszczak has been partially supported by NCN grant 2019/33/B/ST6/02011.

1 Introduction

Reversible computation is a paradigm considering computation to proceed not only in the usual, forward direction, but also backwards [22, 76, 68]. Reversible computation has been studied in a variety of models, including sequential programming languages (both imperative [97, 96, 30, 31], functional [95, 88, 38, 43], and object-oriented [85, 35, 36]), concurrent programming languages [53, 54, 37], process calculi [19, 83, 18, 51], universal logic elements [90, 28, 67], Turing machines [8, 67, 3, 70], automata [68, 46], cellular automata [91, 66, 40, 69], modal logics [80], Petri nets [6, 62, 60], event structures [79], term rewriting [1, 72], Markov chains [41], circuits [27, 23], and others. Also, it has found applications in areas as different as low-power computing [47, 27], debugging [59, 93], bidirectional transformations [58, 73], database design [11], simulation [84], robotics [86, 56], quantum computing [74, 65], and biochemical modeling [16, 82, 44, 45]. In some of those applications, including quantum computing, the reversibility of the computational process is enforced by the very nature of the physical process of computation. In some other areas, the reversibility is treated as a crucial feature, implemented, for example, by database transactions. Thus, while the broad idea of reversible computation is the same in all the areas, it has been interpreted and adapted to fit the various settings. The existing notions of reversible computation however have never been compared and categorized in detail. This work aims at being a first stepping stone towards a taxonomy of the approaches that co-exist under the term reversible computation.

We remark that defining a taxonomy for a field as heterogeneous as reversible computation is a very difficult task, and as far as we are aware this is the first effort in this direction. As such, we provide a *possible* classification, with the aim to start the discussion. We do not claim that our taxonomy is the final word on the subject, and indeed other dimensions may be worth considering, in addition to or instead of some of the ones that we propose (cf. Section 2), and of course we do not claim to be complete on our coverage of models⁹ and languages (cf. Section 3). Indeed, our examples are concentrated in the area of formal methods and programming languages, which are the main expertise of most of the authors, and we may have missed very significant examples from other areas.

2 Taxonomy

In this section we present our taxonomy for approaches in the area of reversible computation. The taxonomy includes six *dimensions*, and for each dimension we describe different *positions* that a given approach may fit. In many cases, positions on a dimension are ordered, in the sense that one generalizes the other.

⁹ In the following we mainly use the term model to refer to the instances of reversible computation that we consider. Indeed, many of our examples are (formal) models. However, we think that our taxonomy can be applied also to more concrete entities, such as languages, applications or systems.

Hence, of course, if an approach fits a position it fits also all the more general ones. We write (dim, pos) to refer to position pos in dimension dim . We write $(d1, p1) \implies (d2, p2)$ to say that pair $(d1, p1)$ implies $(d2, p2)$, and dually the latter is a generalization of the former.

We note that the same approach may fit different positions, depending on the level of abstraction. Vice versa, very different models may fit the same position in the classification. This is the case for instance of the reversible imperative language Janus [97] and of reversible logic elements [28]. This is partially by construction, in the sense that we tried to focus in the taxonomy on the features of the reversibility mechanisms, abstracting away as far as possible from the features of the underlying model.

We describe below each dimension, by explaining the different positions, with examples of approaches that fit each of them. We refer to Section 3 for a more comprehensive description about where models from the literature fit in our taxonomy.

Reversibility focus (FOC): This is the main dimension in the proposed taxonomy. It refers to which aspects of a model are looked at to check whether it is reversible. It features three positions, listed below.

Functional behavior (FUN): In this case a system is said to be reversible if it computes an injective function. Indeed, injectivity ensures that there is a single input which can result in a given output, hence from the output one can recompute the input. As examples, reversible Turing machines, Janus programs, reversible circuits, quantum circuits, and the biorthogonal automata of [1] all define injective functions, hence they fit this position. The functional behavior can be computationally as powerful as reversible Turing machines (r-Turing-complete) [3], or subuniversal [70] and total (always terminating) such as in reversible primitive recursive programs [75] and reversible circuits [28]. Reversible circuits compute exactly the bijective Boolean functions, which are a proper subset of the partial injective functions that are computable by reversible Turing machines and r-Turing complete reversible languages, like Janus.

Reachable states (STA): In this case a system is said to be reversible if it can go back to past states. Checkpointing and SVN are real world techniques fitting this position. Some notions of reversibility in Petri nets [25], requiring that the initial state is reachable from any state, fit this dimension too. Notably, this class of approaches does not consider how past states are reachable, allowing one to reach them via transitions unrelated to the ones used in the past of the computation. Actually, approaches such as rollback directly restore past states, without taking a step-by-step approach to reach them. Notably, not all past states may be reachable, or they may be reachable only with some approximation.

Undoing steps (UND): In this case a system is said to be reversible if it can undo previous steps. This may require or not using special memory or history information. Reversible process calculi [19, 83, 51, 10, 12], cellular automata [91] and Janus [97] fit in this position. Note that Janus fits the

FUN position too: the position depends on the level of abstraction. If we consider a small step semantics, then Janus fits position **UND**; if we abstract away execution details and just look at the functional behavior, then Janus fits position **FUN**.

Note that, if one is able to undo steps, then by undoing steps one can reach past states. Hence, we have the relation $(\text{FOC}, \text{UND}) \implies (\text{FOC}, \text{STA})$.

Also, if a functional behavior can be defined, by undoing steps one can compute the unique inverse function. Hence, the computed function is injective (keeping into account additional memory if present), and we have the implication $(\text{FOC}, \text{UND}) \implies (\text{FOC}, \text{FUN})$.

Resources for reversibility (RES): This dimension refers to whether a model is directly reversible, or whether additional resources (e.g., memory) are needed to enable backward execution.

None (NON): The model is directly reversible, without needing additional memory. Janus, reversible Turing machines and reversible cellular automata fit here. Janus [97] is the standard representative of the class of *clean* (without garbage) reversible programming languages, which all fit this position [32]. This class includes imperative [30, 31], functional [95, 88, 38, 43], and object-oriented [85, 35, 36] languages; reversible flowchart languages [96] to model the high-level structured languages, as well as low-level machine code for reversible von Neumann machines [4, 89].

We remark that models designed without reversibility in mind (e.g., mainstream programming languages) in most of the cases do not fit this position (quantum circuits are an exception to this observation though). In order for models to fit this position, one normally restricts a general class of models. For example, the reversible Turing machines [8] are a forward- and backward-deterministic subset of the Turing machines.

Inside the model (INS): In order to enable reversibility some history information is needed. This information is represented in the same formalism as the original system. This happens, e.g., for some Petri nets [60], where additional places and tokens can be used to keep such history information. This is also the case for reversible rewrite systems [72], where some additional terms are added to make a function injective, an approach which is similar to the addition of a *complement* in the bidirectionalization of functional programs [58, 73]. Another example is the Reverse C Compiler [76] that instruments C programs with statements that trace the computation history. Earlier work that trace at the source level are for Pascal programs [13] and for irreversible Turing machines [8].

Outside the model (OUT): In order to enable reversibility some history information is needed, but to represent this information the model needs to be extended. This happens normally in process calculi [19, 83, 51] and when mainstream programming languages are made reversible: RCCS processes [19] are not CCS processes, and reversible Erlang [53] is not

plain Erlang (since the interpreter is instrumented to additionally store history information).

In reversible event structures, additional relations on events such as precedence or direct causation are used to work out how to reverse events [92].

It is easy to note that no history information is a particular case of history information, and history information outside the model can mimic history information inside the model. Thus, $(\text{RES}, \text{NON}) \implies (\text{RES}, \text{INS})$ and $(\text{RES}, \text{INS}) \implies (\text{RES}, \text{OUT})$.

Moreover, the classification in this dimension depends on the definition of the model. Notably, by considering a model together with the additional memory needed to make it reversible, one moves from position OUT to INS, or even to NON if one considers history information as part of the normal runtime information of the system. Notably, a model of category NON is able to run backwards without having first run forwards, while for models in category IN or OUT one first needs to run forward to generate and store history information. However, if one looks at history information as part of the state, then one can imagine running backwards directly, just by providing history information as part of the starting state. In practice, the history is often difficult to construct without running a program because it depends on the operational internals of the program.

To summarize this discussion, the categorization of a model inside this dimension critically relies on a clear definition of which is the basic model and which is the history information. This distinction comes out naturally when a reversible model is obtained by extending a non-reversible one: in this case what is added to the non-reversible model can be considered as history information kept to enable reversibility. This is the case of, e.g., RCCS [19], which extends CCS with reversibility by equipping each process with a dedicated memory, and in general of Landauer embedding [47].

When reversibility is enabled (WHE): This dimension considers whether reversibility is always enabled or not.

Always (ALW): Reversibility is always enabled, one can take any state and compute backwards. This happens, e.g., in Janus. Process calculi require history information to compute backwards, but we fit them here if they can always go backwards provided that history information is available. The distinction between the Janus case and the process calculi case can be made by looking at dimension RES.

Sometimes (SOM): Reversibility is not always enabled, i.e. there are irreversible steps or other conditions that need to be satisfied for enabling reversibility. Some of the examples include RCCS with irreversible actions [20] (and in general models or languages featuring control mechanisms for reversibility [50]), robotics [56], where some actions (e.g., gluing objects together and drilling holes) cannot be physically reversed, or hybrid quantum-classical algorithms, where only part of the calculation is executed using a reversible quantum circuit.

In this dimension we have $(\text{WHE}, \text{ALW}) \implies (\text{WHE}, \text{SOM})$. Notably, we stated above that Janus fits position ALW, since one can execute backwards from any state, however Janus also has mechanisms to change the direction of execution, in particular the *uncall* of a function computes its inverse function, which can be seen as a control mechanism to decide when reversibility is enabled. Clean reversible programming languages, including reversible machine code, typically include mechanisms that allow to change the computation direction at run time. However, no such mechanism is available in reversible Turing machines [2].

Order of undoing (ORD): This dimension is a sub-dimension of the location (FOC, UND), and refers to which transitions can be reversed at a given point in the execution.

Reverse order (REV): This requires actions to be undone in reverse order of completion. This is the typical notion of reversibility in sequential systems (e.g., reversible Turing machines, Janus), and backtracking in concurrent systems [9] is also an example of REV. Notably, REV ensures that at any point in time a single backward action is enabled, hence the model is backward deterministic.

Causal order (CAU): This requires actions to be undone only if their consequences, if any, are undone beforehand. Equivalently, causal dependent actions need to be undone in reverse order, while independent actions can be undone in any order. This approach, born in the area of process calculi, is known as causal-consistent reversibility [19, 52]. This is the typical notion of reversibility in concurrent process calculi and languages (e.g., RCCS [19], reversible Erlang [53, 54], ...). It has also been used in reversible event structures [92], and reversible Occurrence Nets [62]. In this position the notion of backward determinism from position REV is weakened into backward confluence.

Out of causal order (OCO): This position does not prescribe any constraint on when actions can be undone. This has been used, e.g., in biological systems and models for them (in some process calculi [82, 45], some Petri nets [77, 61]) and in modeling distributed antenna selection for massive MIMO [87] systems.

We have $(\text{ORD}, \text{REV}) \implies (\text{ORD}, \text{CAU})$ and $(\text{ORD}, \text{CAU}) \implies (\text{ORD}, \text{OCO})$. Some Petri net models [77] can be tuned so as to cover all three positions in this dimension.

State reachability (STR): This dimension is a sub-dimension of (FOC, STA), and roughly corresponds to the dimension ORD above. This describes which states can be reached by backward execution.

Only past states (PAS): In this position only past states can be reached. This is typical of sequential models (e.g., Janus) or concurrent models when backtracking is used.

Only past states up-to concurrency (CON): Only states that could have been reached in the past by swapping the order of concurrent actions can be reached. This is the typical behavior of concurrent systems

based on the causal-consistent approach, such as concurrent process calculi and languages (e.g., RCCS [19], reversible Erlang [53]).

States reachable by going forward (FOR): In this case backward execution does not introduce new states, but may allow to reach states in different ways. This happens for instance in Petri nets [6], where one would like to avoid introducing new states, but it does not matter whether the states were in the past of the computation or not.

Also states not reachable by going forward (NOT): In this case backward execution allows computation to reach new states. This behavior may happen in the presence of out of causal order reversibility (ORD, OCO), hence typically in biological systems. In Petri nets there is a line of work [6] trying to understand whether the specific net falls under location NOT or under location FOR.

For state reachability, we have $(\text{STR}, \text{PAS}) \implies (\text{STR}, \text{CON})$, $(\text{STR}, \text{CON}) \implies (\text{STR}, \text{FOR})$ and $(\text{STR}, \text{FOR}) \implies (\text{STR}, \text{NOT})$. This dimension is clearly related to dimension ORD: if a system can be looked at both from the point of view of undoing actions and from the point of view of reachable states, (ORD, REV) corresponds to (STR, PAS), (ORD, CAU) to (STR, CON), and (ORD, OCO) to either (STR, FOR) or (STR, NOT). It would be interesting to find a position in classification ORD corresponding to (STR, FOR), but it is not clear whether any such position exists.

Preciseness of reversibility (PRE): This dimension refers to whether by going backwards one perfectly undoes the effect of forward moves or not.

Precise (PRC): Going forwards and then backwards exactly restores the original state. This happens in most of the models (e.g., Janus, process calculi). This has been captured in causal-consistency theory by the Loop Lemma [19].

With additional information (ADD): When going backwards one keeps some information on the undone computation, e.g., that an unsuccessful try has been performed (to avoid doing the same try again), or that a possible solution of the problem has been found (but one would like to find all the solutions). This approach has been partially explored in the area of reversible process calculi using alternatives [48] (which allow one to select a different computation upon rollback) or predictions [94] (which are not involved in backward computation, hence keep trace of what happened). It has also been studied in the field of session types [17], where branches of a choice are discarded upon rollback, and of reversible contracts [5], where different alternatives are explored looking for a compatible behavior with another process.

Approximate (APP): By going forwards and backwards one can reach a state which is close in some sense to the starting one, but not exactly the same. This happens typically in long-running transactions with compensations [14, 15], where the compensation does an approximate undo, and in robotics [56], where perfect reversibility is not possible due

to small imprecisions in physical actions. Similarly, in reversible neural networks when inputs are recalculated from outputs (not using precise arithmetic), one only gets inputs equal to the original ones up to some small error [33, 7].

We have $(\text{PRE}, \text{PRC}) \implies (\text{PRE}, \text{ADD})$ and $(\text{PRE}, \text{ADD}) \implies (\text{PRE}, \text{APP})$.

Another possible dimension concerns control of reversibility, namely whether there is any policy to decide which action to take when more than one (forward or backward) action is enabled. Possible positions include uncontrolled (no such policy), semantic control (policy hardwired in the language definition), internal control (there are specific constructs in the model to specify the policy) and external control (the policy comes from outside the program, e.g., from the user or from another program). This dimension has been discussed in [50]. We note that frequently uncontrolled reversibility corresponds to (WHE, ALW) while forms of control correspond to (WHE, SOM) , since the policy may disallow backward actions under some conditions.

3 Application of the Taxonomy

While in the previous section we discussed the different dimensions of the taxonomy, here we focus on which approach fits which position in the taxonomy. While there is a partial overlap with the examples given in the previous section, this dual view provides interesting insights as well. The results of this section are captured in Table 1.

Research on reversible computing first tackled sequential models of computation, such as finite state automata and Turing machines. The basic idea was to take the original models and restrict to those instances which were reversible. This naturally led to approaches focused on undoing actions at the small step level, computing injective functions at the global level. Actions were undone in reverse order, as natural for sequential systems, leading back to past states in a precise way. This is the case, e.g., of the language Janus and the biorthogonal automata of [1]. In turn, some sequential models were extended in order to become reversible by introducing a so-called Landauer embedding [47]. Here, we find, e.g., reversible rewrite systems [72] and the *bidirectionalization* of functional programs in [58].

Such an approach was less suitable for concurrent systems, where reverse order of undoing was too strict in many cases, and one would like to be able to undo independent actions in any order, while undoing dependent actions in reverse order. This was first argued in [19], which introduced the notion of causal-consistent reversibility. Instead of restricting calculi to their injective part, memories were added to keep track of past execution (thus fitting position (RES, OUT)), and enable backward computation. Given that in concurrency functional behavior is of limited interest, since interaction with the environment is important, the focus is mainly on undoing actions. A similar approach has been applied to programming languages for concurrency, in particular Erlang [53, 54], where causal

Formalism	Approach	FOCUS	RESOURCE	WHEN	ORDER	STATE R.	PRECIS.
Reversible Turing machines	[8, 67]	FUN UND	NON	ALW	REV	PAS	PRC
Janus	[97]	FUN UND	NON	ALW	REV	PAS	PRC
Biorthogonal automata	[1]	FUN UND	NON	ALW	REV	PAS	PRC
Reversible cellular automata	[91, 66]	FUN UND	NON	ALW	REV	PAS	PRC
Reversible logic elements	[28, 67]	FUN UND	NON	ALW	REV	PAS	PRC
Reversible rewriting	[72]	UND	INS OUT	ALW	REV	PAS	PRC
Causal-consistent calculi	[19, 83] [18, 51] [10, 12]	UND	OUT	ALW	CAU	CON	PRC
Calculi + control	[20, 49]	UND	OUT	SOM	CAU	CON	PRC
Calculi + predictions	[94]	UND	OUT	ALW	CAU	NOT	ADD
Reversible Erlang	[53, 54]	UND	OUT	ALW SOM	CAU	CON	PRC
Petri nets	[6]	STA	NON	ALW	OCO	FOR NOT	PRC
Reversing Petri nets	[77, 78]	UND	INS	ALW	CAU OCO	CON FOR NOT	PRC
Occurrence nets	[62]	UND	NON	ALW	CAU	CON	PRC
Petri nets	[60]	UND	INS	ALW	CAU	CON	PRC
Biological models	[82, 45]	UND	OUT	SOM	OCO	NOT	PRC
Event structures	[81, 92]	UND	OUT	SOM	CAU OCO	CON FOR NOT	PRC
Quantum circuits	[24, 26]	FUN UND	NON	ALW	REV	PAS	PRC
Quantum programming languages	[42, 74]	UND	INS	SOM	REV	PAS	PRC
Reversible neural networks	[33]	UND	NON OUT	SOM	REV	NOT	APP
Reversible Markov chains	[41]	STA	NON	ALW	REV	PAS	PRC
Sagas	[14]	UND	INT	SOM	CAU	X	APP
SVN	[63]	STA	INT	ALW	X	PAS	PRC

Table 1. Application of the taxonomy to sample approaches from the literature

consistency is ensured for both forward (replay) and backward computations during debugging [55, 29].

While the first approaches considered precise reversibility which was always enabled, further studies introduced control mechanisms [20, 49] as well as forms of reversibility which were not precise [94]. Some applications, most notably in the biochemical setting, triggered the need for weakening causal order, thus introducing out of causal order reversibility [82, 45]. We note that CCSK [83], with the addition of a control mechanism in the form of a rollback operator inspired by [49], has been modeled using reversible event structures exploiting out of causal order reversibility [34].

Petri nets, while being a model for concurrency like process calculi, resulted in a number of different approaches. The fact that Petri nets have a clear representation of state (in terms of tokens inside places), triggered approaches [25] focusing on state reachability more than on action undoing. Approaches based on action undoing were also considered and contrasted with the ones based on state reachability [6]. Other works [62] considered the causal-consistent approach, thus matching the one of process calculi. Further work tailored Petri nets for biological applications [77], allowing one to explore different forms of reversibility, most notably the out of causal order one.

In the quantum circuit model [24, 26] used for developing most quantum mechanical algorithms, the set of allowed operations is represented by unitary matrices or unitary gates. Such matrices act on an isolated physical system and, in this scenario, one is always able to undo the last action. Hence, the term reversible is, in quantum computer science, synonymous with the term ‘unitary’. Compared with classical reversible gates, unitary matrices provide us with a larger set of operations. However, to read out the result of the computation, one needs to translate the final state into the classical result. Such a process requires a measurement which is achieved through interaction with the system executing the computation. The main feature of such a process is its irreversibility. Thus, reversibility is lost at the moment of ‘interfacing’ with a classical machine or with the readout procedure. Architecture-specific limitations of current quantum hardware lead to the problem of optimizing quantum circuits [71], most importantly taking into account the hardware topology [21]. Such optimization is part of the process of transpilation – translation of quantum circuits to the form suitable for the target quantum computer.

This need of interfacing between the reversible and the irreversible elements motivated the development of quantum programming languages [42, 74, 64]. Also, many quantum algorithms (NISQ algorithms in particular) use classical subroutines. Quantum programming languages include a specialized type system for handling quantum structures used in purely quantum, reversible computation. Additionally, they also include an irreversible subsystem, suitable for dealing with classical – which in this case means irreversible – computation.

Reversibility is used in *Convolutional Neural Networks* [57] (CNNs) to undo computation of the networks’ layers. This removes the need to store, retrieve and delete layers’ inputs and outputs, which can be recomputed instead. Some layers perform transformations (of inputs to outputs) which have inverses, such as multiplication by a matrix, so are directly reversible. Other transformations,

such as applying a convolution or max pooling, lose data so can only be reversed by enriching the network with additional components. A *Reversible Residual Network* [33] (ResNet) is a form of CNN that adds *shortcuts* between layers. This makes it possible to undo computation of most layers. Calculation is not in precise arithmetic, so only approximate values of inputs can be uncomputed from outputs (up to an agreed precision), and thus new states can be reached.

In the field of performance evaluation, a Markov chain is (time) reversible [41] if it has the same behavior as its inverse, in terms of probabilistic distribution. Hence the focus is on states, and, since the approach restricts attention to Markov chains which naturally satisfy the reversibility property, no additional resources are required. Reversible transitions are always enabled, though they are subject to a probabilistic distribution, and the order of reversing is the inverse of the forward one. An initial work relating causal-consistent reversibility with reversible Markov chains in the setting of a stochastic process algebra is described in [10].

We conclude the table with a few approaches which are at the boundary of reversible computation, namely Sagas [14], used to model long-running transactions with compensation, and the well-known tool SVN for version control [63, 39]. Given the distance from classical reversible computation, it is not clear whether some of the dimensions make sense in these cases. We put ‘X’ in the cells which we believe are not interesting.

4 Conclusion, Related and Future Work

We have presented a first proposal of taxonomy for reversible computation approaches, and discussed how various models fit in it. We focused on approaches from programming languages and concurrency theory, hence in future work it would be good to put our taxonomy at work also on other kinds of models.

We are not aware of other works putting forward proposals of taxonomies for reversible computing. A partial analysis in this direction is the classification of control mechanisms in [50] and an account of reversible computing from a programming language perspective [32]. Also, a few works in the context of Petri nets contrast different approaches [6, 77], taking advantage of the existence of many such approaches.

Table 1, while not covering all the literature, highlights some holes which are interesting targets for future work. For instance, a large part of the approaches concern precise reversibility, and indeed this is the main focus of the reversible computing community. Approaches where reversibility is not perfect are however of interest as well, motivated, e.g., by applications in robotics and neural networks, and are an interesting research direction for the reversible computation community. Another interesting point is that most of the approaches focus on undoing actions, while a focus on functional behavior and on states has been adopted only in a few cases. From a theoretical perspective, it would also be interesting to investigate the computational power and inherent complexity of reversible computing models.

Acknowledgements This work refines and extends the results of discussions that occurred during the meetings of the COST Action IC1405 on Reversible Computation – Extending Horizons of Computing. We thank all the participants to such discussions. The authors were partially supported by the COST Action IC1405. We thank the anonymous referees for their helpful comments.

References

1. Samson Abramsky. A structural approach to reversible computation. *Theor. Comput. Sci.*, 347(3):441–464, 2005.
2. Holger Bock Axelsen and Robert Glück. A simple and efficient universal reversible Turing machine. In *LATA 2011*, volume 6638 of *LNCS*, pages 117–128. Springer, 2011.
3. Holger Bock Axelsen and Robert Glück. On reversible Turing machines and their function universality. *Acta Inform.*, 53(5):509–543, 2016.
4. Holger Bock Axelsen, Robert Glück, and Tetsuo Yokoyama. Reversible machine code and its abstract processor architecture. In *CSR 2007*, volume 4649 of *LNCS*, pages 56–69. Springer, 2007.
5. Franco Barbanera, Ivan Lanese, and Ugo de'Liguoro. A theory of retractable and speculative contracts. *Sci. Comput. Program.*, 167:25–50, 2018.
6. Kamila Barylska, Maciej Koutny, Łukasz Mikulski, and Marcin Piątkowski. Reversible computation vs. reversibility in Petri nets. *Sci. Comput. Program.*, 151:48–60, 2018.
7. Jens Behrmann, Paul Vicol, Kuan-Chieh Wang, Roger B. Grosse, and Jörn-Henrik Jacobsen. Understanding and mitigating exploding inverses in invertible neural networks. In *AISTATS 2021*, volume 130 of *Proceedings of Machine Learning Research*, pages 1792–1800. PMLR, 2021.
8. Charles H. Bennett. Logical reversibility of computation. *IBM J. Res. Dev.*, 17(6):525–532, 1973.
9. Jan A. Bergstra, Alban Ponse, and Jos van Wamel. Process algebra with backtracking. In *A Decade of Concurrency, Reflections and Perspectives, REX School/Symposium*, volume 803 of *LNCS*, pages 46–91. Springer, 1993.
10. Marco Bernardo and Claudio Antares Mezzina. Towards bridging time and causal reversibility. In *FORTE 2020*, volume 12136 of *LNCS*, pages 22–38. Springer, 2020.
11. Phillip A. Bernstein and Eric Newcomer. *Principles of Transaction Processing*. Morgan Kaufmann Publishers Inc, US, 2nd edition, 2009.
12. Laura Bocchi, Ivan Lanese, Claudio Antares Mezzina, and Shoji Yuen. The reversible temporal process language. In *FORTE 2022*, volume 13273 of *LNCS*, pages 31–49. Springer, 2022.
13. James S. Briggs. Generating reversible programs. *Software Pract. Exper.*, 17(7):439–453, 1987.
14. Roberto Bruni, Hernán C. Melgratti, and Ugo Montanari. Theoretical foundations for compensations in flow composition languages. In *POPL 2005*, pages 209–220. ACM, 2005.
15. Luís Caires, Carla Ferreira, and Hugo Torres Vieira. A process calculus analysis of compensations. In *TGC 2008*, volume 5474 of *LNCS*, pages 87–103. Springer, 2008.
16. Luca Cardelli and Cosimo Laneve. Reversibility in massive concurrent systems. *Sci. Ann. Comput. Sci.*, 21(2):175–198, 2011.

17. Ilaria Castellani, Mariangiola Dezani-Ciancaglini, and Paola Giannini. Reversible sessions with flexible choices. *Acta Inform.*, 56(7-8):553–583, 2019.
18. Ioana Cristescu, Jean Krivine, and Daniele Varacca. A compositional semantics for the reversible π -calculus. In *LICS 2013*, pages 388–397. IEEE Computer Society, 2013.
19. Vincent Danos and Jean Krivine. Reversible communicating systems. In *CONCUR 2004*, volume 3170 of *LNCS*, pages 292–307. Springer, 2004.
20. Vincent Danos and Jean Krivine. Transactions in RCCS. In *CONCUR 2005*, volume 3653 of *LNCS*, pages 398–412. Springer, 2005.
21. Marc G Davis, Ethan Smith, Ana Tudor, Koushik Sen, Irfan Siddiqi, and Costin Iancu. Towards optimal topology aware quantum circuit synthesis. In *QCE 2020*, pages 223–234. IEEE, 2020.
22. Alexis De Vos. *Reversible Computing: Fundamentals, Quantum Computing, and Applications*. Wiley-VCH, 2010.
23. Alexis De Vos, Stijn De Baerdemacker, and Yvan Van Rentergem. *Synthesis of Quantum Circuits vs. Synthesis of Classical Reversible Circuits*. Synthesis Lectures on Digital Circuits and Systems. Morgan & Claypool Publishers, 2018.
24. David Deutsch. Quantum theory, the Church–Turing principle and the universal quantum computer. *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences*, 400(1818):97–117, 1985.
25. Javier Esparza and Mogens Nielsen. Decidability issues for Petri nets. *BRICS Report Series*, 1(8), 1994.
26. Richard P. Feynman. Quantum mechanical computers. *Found. Phys.*, 16(6):507–531, 1986.
27. Michael P. Frank. *Reversibility for Efficient Computing*. PhD thesis, Massachusetts Institute of Technology, Cambridge, MA, USA, 1999.
28. Edward Fredkin and Tommaso Toffoli. Conservative logic. *Intl. J. Theor. Phy.*, 21(3-4):219–253, 1982.
29. Elena Giachino, Ivan Lanese, and Claudio Antares Mezzina. Causal-consistent reversible debugging. In *FASE 2014*, volume 8411 of *LNCS*, pages 370–384. Springer, 2014.
30. Robert Glück and Tetsuo Yokoyama. A linear-time self-interpreter of a reversible imperative language. *Computer Software*, 33(3):108–128, 2016.
31. Robert Glück and Tetsuo Yokoyama. A minimalist’s reversible while language. *IEICE Trans. Inf. Syst.*, E100-D(5):1026–1034, 2017.
32. Robert Glück and Tetsuo Yokoyama. Reversible computing from a programming language perspective. *Theor. Comput. Sci.*, 953:Article 113429, 2023.
33. Aidan N. Gomez, Mengye Ren, Raquel Urtasun, and Roger B. Grosse. The reversible residual network: Backpropagation without storing activations. In *Advances in Neural Information Processing Systems 30: NIPS 2017*, volume 30, pages 2214–2224. Curran Associates, Inc., 2017.
34. Eva Graversen, Iain C. C. Phillips, and Nobuko Yoshida. Event structure semantics of (controlled) reversible CCS. *J. Log. Algebraic Methods Program.*, 121:100686, 2021.
35. Tue Haulund, Torben Æ. Mogensen, and Robert Glück. Implementing reversible object-oriented language features on reversible machines. In *Reversible Computation 2017*, volume 10301 of *LNCS*, pages 66–73. Springer, 2017.
36. Lasse Hay-Schmidt, Robert Glück, Martin Holm Cservenka, and Tue Haulund. Towards a unified language architecture for reversible object-oriented programming. In *Reversible Computation 2021*, volume 12805 of *LNCS*, pages 96–106. Springer, 2021.

37. James Hoey and Irek Ulidowski. Reversing an imperative concurrent programming language. *Sci. Comput. Program.*, 223:102873, 2022.
38. Petur Andrias Højgaard Jacobsen, Robin Kaarsgaard, and Michael Kirkedal Thomsen. CoreFun: A typed functional reversible core language. In *Reversible Computation 2018*, volume 11106 of *LNCS*, pages 304–321. Springer, 2018.
39. Judah Jacobson. A formalization of darcs patch theory using inverse semigroups. Technical report, UCLA, 2009.
40. Jarkko Kari. Reversible cellular automata: From fundamental classical results to recent developments. *New Gener. Comput.*, 36(3):145–172, 2018.
41. Frank P. Kelly. *Reversibility and Stochastic Networks*. Wiley, 1979.
42. Emanuel Knill. Conventions for quantum pseudocode. Technical Report LAUR-96-2724, Los Alamos National Lab, 1996.
43. Joachim Tilsted Kristensen, Robin Kaarsgaard, and Michael Kirkedal Thomsen. Jeopardy: An invertible functional programming language. *CoRR*, arXiv:2209.02422, 2022.
44. Stefan Kuhn and Irek Ulidowski. Local reversibility in a calculus of covalent bonding. *Sci. Comput. Program.*, 151:18–47, 2018.
45. Stefan Kuhn and Irek Ulidowski. Modelling of DNA mismatch repair with a reversible process calculus. *Theor. Comput. Sci.*, 925:68–86, 2022.
46. Martin Kutrib and Andreas Malcher. Reversible pushdown automata. *J. Comput. Syst. Sci.*, 78(6):1814–1827, 2012.
47. Rolf Landauer. Irreversibility and heat generation in the computing process. *IBM J. Res. Dev.*, 5(3):183–191, 1961.
48. Ivan Lanese, Michael Lienhardt, Claudio Antares Mezzina, Alan Schmitt, and Jean-Bernard Stefani. Concurrent flexible reversibility. In *ESOP 2013*, volume 7792 of *LNCS*, pages 370–390. Springer, 2013.
49. Ivan Lanese, Claudio Antares Mezzina, Alan Schmitt, and Jean-Bernard Stefani. Controlling reversibility in higher-order pi. In *CONCUR 2011*, volume 6901 of *LNCS*, pages 297–311. Springer, 2011.
50. Ivan Lanese, Claudio Antares Mezzina, and Jean-Bernard Stefani. Controlled reversibility and compensations. In *Reversible Computation 2012*, volume 7581 of *LNCS*, pages 233–240. Springer, 2012.
51. Ivan Lanese, Claudio Antares Mezzina, and Jean-Bernard Stefani. Reversibility in the higher-order π -calculus. *Theor. Comput. Sci.*, 625:25–84, 2016.
52. Ivan Lanese, Claudio Antares Mezzina, and Francesco Tiezzi. Causal-consistent reversibility. *Bulletin of the EATCS*, 114, 2014.
53. Ivan Lanese, Naoki Nishida, Adrián Palacios, and Germán Vidal. A theory of reversibility for Erlang. *J. Log. Algebr. Meth. Program.*, 100:71–97, 2018.
54. Ivan Lanese, Adrián Palacios, and Germán Vidal. Causal-consistent replay reversible semantics for message passing concurrent programs. *Fundam. Informaticae*, 178(3):229–266, 2021.
55. Ivan Lanese, Ulrik Pagh Schultz, and Irek Ulidowski. Reversible computing in debugging of Erlang programs. *IT Prof.*, 24(1):74–80, 2022.
56. Johan Sund Laursen, Lars-Peter Ellekilde, and Ulrik Pagh Schultz. Modelling reversible execution of robotic assembly. *Robotica*, 36(5):625–654, 2018.
57. Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proc. IEEE*, 86(11):2278–2324, 1998.
58. Kazutaka Matsuda, Zhenjiang Hu, Keisuke Nakano, Makoto Hamana, and Masato Takeichi. Bidirectionalization transformation based on automatic derivation of view complement functions. In *ICFP 2007*, pages 47–58. ACM, 2007.

59. James McNellis, Jordi Mola, and Ken Sykes. Time travel debugging: Root causing bugs in commercial scale software. CppCon talk, https://www.youtube.com/watch?v=11YJTg_A914, 2017.
60. Hernán C. Melgratti, Claudio Antares Mezzina, and G. Michele Pinna. Towards a truly concurrent semantics for reversible CCS. In *Reversible Computation 2021*, volume 12805 of *LNCS*, pages 109–125. Springer, 2021.
61. Hernán C. Melgratti, Claudio Antares Mezzina, and G. Michele Pinna. A Petri net view of covalent bonds. *Theor. Comput. Sci.*, 908:89–119, 2022.
62. Hernán C. Melgratti, Claudio Antares Mezzina, and Irek Ulidowski. Reversing place transition nets. *Log. Methods Comput. Sci.*, 16(4), 2020.
63. Samuel Mimram and Cinzia Di Giusto. A categorical theory of patches. In *MFPS XXIX*, volume 298 of *Electr. Notes Theor. Comput. Sci.*, pages 283–307. Elsevier, 2013.
64. Jarosław Miszczak. Models of quantum computation and quantum programming languages. *Bulletin of the Polish Academy of Sciences: Technical Sciences*, 59(3):305–324, Sep 2011.
65. Jarosław Miszczak. *High Level Structures for Quantum Computing*. Springer International Publishing, 2012.
66. Kenichi Morita. Computation-universality of one-dimensional one-way reversible cellular automata. *Inf. Process. Lett.*, 42(6):325–329, 1992.
67. Kenichi Morita. Reversible computing and cellular automata—a survey. *Theor. Comput. Sci.*, 395(1):101–131, 2008.
68. Kenichi Morita. *Theory of Reversible Computing*. Monographs in Theoretical Computer Science. An EATCS Series. Springer, 2017.
69. Daniel Morrison and Irek Ulidowski. Direction-reversible self-timed cellular automata for delay-insensitive circuits. *J. Cell. Autom.*, 12(1-2):101–120, 2016.
70. Keisuke Nakano. Time-symmetric Turing machines for computable involutions. *Sci. Comput. Program.*, 215:102748, 2022.
71. Beatrice Nash, Vlad Gheorghiu, and Michele Mosca. Quantum circuit optimizations for NISQ architectures. *Quantum Sci. Technol.*, 5(2):025010, 2020.
72. Naoki Nishida, Adrián Palacios, and Germán Vidal. Reversible computation in term rewriting. *J. Log. Algebraic Methods Program.*, 94:128–149, 2018.
73. Naoki Nishida and Germán Vidal. Characterizing compatible view updates in syntactic bidirectionalization. In *Reversible Computation 2019*, volume 11497 of *LNCS*, pages 67–83. Springer, 2019.
74. Bernard Ömer. *Structured Quantum Programming*. PhD thesis, Vienna University of Technology, 2003.
75. Luca Paolini, Mauro Piccolo, and Luca Roversi. On a class of reversible primitive recursive functions and its Turing-complete extensions. *New Gener. Comput.*, 36(3):233–256, 2018.
76. Kalyan S. Perumalla. *Introduction to Reversible Computing*. CRC Press/Taylor & Francis Group, 2014.
77. Anna Philippou and Kyriaki Psara. Reversible computation in Petri nets. In *Reversible Computation 2018*, volume 11106 of *LNCS*, pages 84–101. Springer, 2018.
78. Anna Philippou and Kyriaki Psara. A collective interpretation semantics for reversing Petri nets. *Theor. Comput. Sci.*, 924:148–170, 2022.
79. Iain Phillips and Irek Ulidowski. Reversibility and asymmetric conflict in event structures. In *CONCUR 2013*, volume 8052 of *LNCS*, pages 303–318. Springer, 2013.

80. Iain Phillips and Irek Ulidowski. Event identifier logic. *Math. Struct. Comput. Sci.*, 24(2), 2014.
81. Iain Phillips and Irek Ulidowski. Reversibility and asymmetric conflict in event structures. *J. Log. Algebraic Methods Program.*, 84(6):781–805, 2015.
82. Iain Phillips, Irek Ulidowski, and Shoji Yuen. A reversible process calculus and the modelling of the ERK signalling pathway. In *Reversible Computation 2012*, volume 7581 of *LNCS*, pages 218–232. Springer, 2012.
83. Iain C. C. Phillips and Irek Ulidowski. Reversing algebraic process calculi. *J. Log. Algebr. Program.*, 73(1-2):70–96, 2007.
84. Markus Schordan, Tomas Opielstrup, David R. Jefferson, and Peter D. Barnes Jr. Generation of reversible C++ code for optimistic parallel discrete event simulation. *New Gener. Comput.*, 36(3):257–280, 2018.
85. Ulrik Pagh Schultz and Holger Bock Axelsen. Elements of a reversible object-oriented language. In *Reversible Computation 2016*, volume 9720 of *LNCS*, pages 153–159. Springer, 2016.
86. Ulrik Pagh Schultz, Mirko Bordignon, and Kasper Støy. Robust and reversible execution of self-reconfiguration sequences. *Robotica*, 29(1):35–57, 2011.
87. Harun Siljak, Kyriaki Psara, and Anna Philippou. Distributed antenna selection for massive MIMO using reversing Petri nets. *IEEE Wirel. Commun. Lett.*, 8(5):1427–1430, 2019.
88. Michael Kirkedal Thomsen and Holger Bock Axelsen. Interpretation and programming of the reversible functional language RFUN. In *IFL 2015*, pages 8:1–8:13. ACM, 2015.
89. Michael Kirkedal Thomsen, Holger Bock Axelsen, and Robert Glück. A reversible processor architecture and its reversible logic design. In *Reversible Computation 2011*, volume 7165 of *LNCS*, pages 30–42. Springer, 2012.
90. Tommaso Toffoli. Reversible computing. In *Automata, Languages and Programming 1980*, volume 85 of *LNCS*, pages 632–644. Springer, 1980.
91. Tommaso Toffoli and Norman Margolus. *Cellular Automata Machines. A New Environment for Modeling*. MIT Press, 1987.
92. Irek Ulidowski, Iain Phillips, and Shoji Yuen. Reversing event structures. *New Gener. Comput.*, 36(3):281–306, 2018.
93. Undo. *UDB - reverse debugger for C/C++*. <https://undo.io>, 2020.
94. Martin Vassor. Reversibility and predictions. In *Reversible Computation 2021*, volume 12805 of *LNCS*, pages 163–181. Springer, 2021.
95. Tetsuo Yokoyama, Holger Bock Axelsen, and Robert Glück. Towards a reversible functional language. In *Reversible Computation 2011*, volume 7165 of *LNCS*, pages 14–29. Springer, 2011.
96. Tetsuo Yokoyama, Holger Bock Axelsen, and Robert Glück. Fundamentals of reversible flowchart languages. *Theor. Comput. Sci.*, 611:87–115, 2016.
97. Tetsuo Yokoyama and Robert Glück. A reversible programming language and its invertible self-interpreter. In *PEPM 2007*, pages 144–153. ACM, 2007.