

Reversible Execution for Robustness in Embodied AI and Industrial Robots

Ivan Lanese , University of Bologna/INRIA, 40126, Italy

Ulrik P. Schultz , University of Southern Denmark, 5230, Denmark

Irek Ulidowski , University of Leicester, LE1 7RH, U.K.

Reversible computation is a computing paradigm where execution can progress backward as well as in the usual, forward direction. It has found applications in many areas of computer science, such as circuit design, programming languages, simulation, modeling of chemical reactions, debugging, and robotics. In this article, we give an overview of reversible computation focusing on its use in robotics. We present an example of programming industrial robots for assembly operations where we combine classical AI planning with reversibility and embodied AI to increase the robustness and versatility of industrial robots.

Reversibility can be defined as the ability of a program or a system to execute in reverse in order to undo the effects of its (forward) computation. Reversibility has interested scientists for many years. Landauer has discovered over 60 years ago that erasing information in computers requires energy and that loss of information, such as erasing a value stored in a variable, during computation is manifested by the release of heat.¹ The scientists thought at the time that if we could build logic circuits and, ultimately, hardware that reduces or even avoids completely the need to remove information, then computers would be more energy efficient. Subsequently, Fredkin and Toffoli developed reversible universal logic gates as an alternative to the traditional CMOS technology gates.² This meant that, at least in theory, it was possible to design and manufacture reversible computers. There has been a significant amount of research on reversible computers since the discovery of reversible logic gates, culminating in many projects to develop reversible circuits and hardware, but these have not changed the way modern hardware is built yet. Apart from this original

motivation for physical reversibility, there are many other reasons for, and benefits of, logical reversibility.³ The latter form of reversibility concerns enhancing systems and software (that run on a physically irreversible hardware) with the ability to undo (or simulate undoing of) computation. There are reversible programming languages such as Janus⁴ and there are techniques for reversing traditional imperative programming languages such as C.⁵

We have also discovered the basics of how to reverse the computation of concurrent programs and systems.^{6–9}

The purpose of this article is to introduce the topic of reversible computation by presenting a robotics case study where logical reversibility has made a difference. The case study, and more generally, reversible computation research in Europe were partially supported by COST Action IC1405 on Reversible Computation—Extending Horizons of Computing.¹⁰ We shall touch gently on the theories we have developed and explain how they assisted us in solving practical problems of the case study. We will also indicate how we have adjusted our formal techniques to strengthen a traditional AI planning approach to produce a full working solution.

Our case study is about programming industrial robots performing assembly operations (i.e., building a physical product) in a way that, based on a fixed assembly sequence generated by an AI-based planner,




FIGURE 1. Experimental setup: industrial robot (left) and the two product parts being automatically assembled and disassembled (right).

achieves automatic error recovery and even automatic disassembly. Error recovery is achieved by temporarily reversing the direction of execution, effectively undoing recent steps, and then trying again. This approach works well in the physical world of robots because slight imprecisions can cause the robot to get stuck, but partially disassembling the object and trying again can often solve the problem. Taken to an extreme, the entire assembly sequence can be reversed, effectively providing an automatic way to disassemble an object.

We thus demonstrate how a traditional AI-based planning approach is enriched by an underlying reversible execution model that relies on the embodiment of the robot system to provide a robust, probabilistic way of executing the plan. The approach is based on the principles of the Janus reversible programming language,⁴ where every step of the computation must in itself be reversible, thus ensuring that the program as a whole is reversible. In Janus, this means that certain irreversible operations, such as multiplication by zero, are not allowed. Similarly, for the robots, reversible execution can not be applied to intrinsically irreversible steps such as cutting or welding.

REVERSIBILITY IN ROBOTICS

Robots act upon the physical world, and depending on the type of robot, may be capable of performing actions that can be considered reversible.

Consider the specific case of an industrial robot, i.e., a general-purpose robot arm as depicted in Figure 1 (left), normally consisting of six or more joints

connected in series and programmed using a special-purpose robot programming language. Moving an object from one location to another, or screwing two pieces of metal together using a bolt, could be considered reversible actions. Conversely, breaking an object in two or welding two pieces of metal together would not be considered reversible. If a robot is performing a sequence of operations that can be considered reversible, such as the steps required to assemble a kitchen appliance or a photocopier, then could the entire sequence of operations be perhaps considered a reversible program?

This thought experiment motivated the study and development of reversible domain-specific robot programming languages. (Domain-specific languages are special-purpose languages designed to solve specific problems such as robot programming¹¹). The key insight is that if the robot is constrained to only perform operations that are physically reversible, then an entire sequence of operations (i.e., a robot program) can be considered physically reversible. Reversible robot programming languages have been studied for industrial robots and modular self-reconfigurable robots. Industrial robots are the topic of this article.

Self-reconfigurable robots are robots that can physically rebuild themselves to take on different shapes.¹² As a concrete example, the ATRON modular robot¹³ is shown in Figure 2: each module is an individual robot, and a snake robot composed of multiple modules can rebuild itself into a car robot. A modular robot shaped as a car can for example rebuild itself into a snake to traverse an obstacle, and then afterward return to the car shape to continue normal operations.




FIGURE 2. ATRON modular robot: a single module (left), car/snake configurations (middle), and the snake changing shape (right).

This process, referred to as self-reconfiguration, can be considered reversible if each of the steps performed by the individual modules is reversible and can be repeated in reverse order.¹⁴

Industrial Robots, AI, and Reversibility

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. The classical AI-based approach is to derive a plan for the sequence of operations required to complete a given task.^{15, 16}

Such plans however break down in case of errors, resulting in the need to replan the sequence of operations. Replanning can be costly, in particular in complex operations where errors could occur quite often.¹⁷

We propose to generate plans as reversible operation sequences, such that if random failure causes a step to fail, the system can automatically backtrack and retry without replanning. Reverse execution here allows the robot to back out of an erroneous situation, after which the operation can be automatically retried. In a perfectly predictable system retrying would usually result in the same errors, but the embodiment of the robot here works to our advantage: retrying the same physical operation multiple times can produce different results. In effect, the AI-generated plan is made robust toward specific kinds of errors and can be executed robustly without any need for costly dynamic replanning. The combination of automatic retries based on reversibility and probabilistic operations can be considered a form of embodied AI, where reversibly retrying the same operation multiple times results in increased robustness.

As we will see, the combination of AI-based planning and reversibility is amongst others useful for automatic error recovery for small-sized batch production of assembly operations, where precisely specifying error-free operations would be time-consuming and expensive, and dynamic replanning would add a significant overhead to the operation. Moreover, reversibility can in this case be used to automatically derive a disassembly sequence from a given assembly sequence or vice versa. These capabilities can be achieved using a reversible domain-specific language for specifying assembly sequences.

Robotic assembly and disassembly are done in terms of sequences of operations, such as placement of objects, insertions, screwing operations, and so forth. All are challenged by uncertainties from sensors, robot kinematics, and part tolerances. Not all operations are reversible, some are not even repeatable! Many physical phenomena and actions can nevertheless be considered reversible, depending on the abstraction level at which they are observed. For example, an industrial robot that pushes an object to a new position could easily move this object back to its original position, but cannot simply do this by reversing its pushing movements, as pulling requires gripping the object first. Moreover, some operations, such as cutting and welding, are in practice nonreversible. A study of 13 real-world industrial cases showed roughly 76% of the operations to be reversible,¹⁸ but many of the operations require the robot to perform a different physical action to reverse a given action. Taking inspiration from this study, we divide reversible operations into two categories: directly reversible operations that simply can be reversed by performing the forward action in reverse; and indirectly reversible operations, which can be reversed, but require a

```
// SCREWING OPERATION
sequence("insert_screw_operation").
  action(insert_screw).
  reverseWith("remove_screw").nonreversible().
  call("insert_screw_suboperation").
  move(qScrewInside).
  move(qScrewOutside);
```

FIGURE 3. Sample reversible assembly program: a sequence is defined to consist of an “insert screw” action, a call to another sequence, and two move actions. The “insert screw” action is not itself reversible and is marked as indirectly reversible using the “remove screw” action.

different sequence of instructions, which must be manually specified by the programmer.

In the design of our robot programming language, we take inspiration from the Janus reversible programming language, where programs are said to be time-invertible.⁴ Each computational step in Janus has a specific inverse, and a given program that when executed forwards computes a function will compute the inverse of this function when executed backward. Subtracting a constant is for example the inverse of adding a constant. In our robot programming language, each physically reversible operation similarly has an inverse. In the case of directly reversible operations, the inverse is automatically derived by the system. In the case of indirectly reversible operations, the programmer must manually specify the sequence of operations that constitutes the reverse of a given operation. Executing such a manually specified inverse can temporarily bring the system into a state not normally encountered during forward execution. Moreover, switching execution direction in the middle of such a manually specified inverse might again take the system into a new state. This contrasts a main property of reversibility in programming languages like Janus and of causal-consistent reversibility,^{6, 7} a notion used in concurrent systems, which says that any reachable state is forwards reachable. In causal-consistent reversibility, any step of computation of a concurrent system can be undone provided that all its effects, if any, are undone first.⁹ Such a property also fails in other contexts, e.g., in some biological systems.¹⁹ Since an operation and its (indirect) reverse are paired, the program has unique starting and ending states, and execution will only terminate in one of these two states. Infinite loops of error correction can manifest and are handled using a monitoring heuristic that detects if the assembly operation might be stuck.

The programming model we have developed is based on this abstract semantics-based model extended with various features required for reversible control of industrial robots in real-world scenarios.¹⁸

The actual implementation is in the form of an internal DSL in C++, meaning that a sequence of C++ method calls is used to build a model of the reversible assembly sequence, as shown in Figure 3. A robot assembly task is programmed as a sequential flow of operations. It is sequential since in practice assembly tasks tend to be a simple sequence of operations (except for error handling, but we aim to automatically handle errors using reverse execution). Reversibility is nevertheless still relevant due to the presence of random behavior of the physical operations: reversing and re-executing an operation may produce a different result. Each operation (denoted by the keyword “sequence”) represents a high-level assembly case logic and is a sequence of primitive instructions. Each instruction is either directly reversible (default), indirectly reversible (indicated by the keyword “reverseWith”), or nonreversible (indicated by “nonreversible”).

Our approach was evaluated experimentally using two industrial assembly use-cases,¹⁸ Figure 1 shows the physical robot platform (left) and the two assembled use-cases (right). Both use-cases were used to test the principle of reversible assembly and the use of reverse execution for error correction. For reversible assembly, the program was executed forward to assemble each use-case. Afterward the finished object was manually placed back into the system, and the program was executed backward to disassemble the object. This was done multiple times for each use-case with no errors.

The use of reverse execution as an effective error correction tool was experimentally demonstrated by assembling a large number of objects, as follows. The workcell was set to assemble 100 objects of each type consecutively and without pause. During these 200 assemblies a total of 22 errors occurred, of which 18, corresponding to 82%, were automatically resolved and corrected using reverse execution. Errors that were automatically corrected include failed peg-in-hole operations (fixed by backtracking and trying

again), dropping a tube (fixed by reversing until a new tube was picked from the feeder), failed to grasp a screw, and screwing failing due to misalignment.

Errors that could not be automatically corrected include air-tubing from the gripper getting stuck on the platform, causing the gripper to misalign, and a screw being inserted at a skewed angle causing a bracket to misalign, which could not be corrected as the system had no means of detecting the bracket misalignment.

CONCLUSION

Reversing of computation is conceptually and technically a challenging task even if we only consider logical reversibility. We have illustrated significant potential benefits of reversibility to improve AI-planning in the robotics case study.

We have presented briefly some of the recently developed theoretical underpinnings for the case study, concentrating mainly on explaining how reversibility helps. Exploring this application area helped us to exemplify the richness of different forms of reversibility.

While the case study we have discussed is based on sequential reversibility, the notion of causal-consistent reversibility,^{6, 7} is key to scalable reversible programming of modular robotic systems such as the ATRON robot shown in Figure 2. This is because the modules perform operations in parallel and hence reversing the system must respect dependencies between the actions of individual modules. Provided the ATRON robot does not perform any irreversible steps, we have this strong property: any reachable (by an arbitrary combination of reverse and forward steps) state is forwards reachable. Applying causal consistency to achieve scalable and robust reversible programming for swarming robot systems like the ATRON and unmanned aerial vehicles (UAVs, drones) is considered future work. In the specific case of UAVs, a programming model based on reversibility would hypothetically allow a drone swarm as a whole to “reverse” distributed control decisions, thus easing requirements on reaching consensus (before beginning new operations) and increasing the robustness of the system.

Robots are controlled through programming, but we cannot be certain of their actions since they interact with an unpredictable physical world. Contrary to causal-consistent reversibility, we have seen that some inverses of indirectly reversible operations may lead to new “get-out-of-trouble” states, albeit temporarily, which are not forwards reachable. Such states are needed due to the irreversibility of the physical world with which the robot interacts.

There are also other forms of reversibility suitable for different applications. Probably the best known is backtracking, where steps of computation are undone in the inverse order of execution.

Apart from many traditional applications of backtracking such as, for example, in search algorithms or logic programming, it has been used to undo concurrent C-like programs for debugging.^{10, 20}

ACKNOWLEDGMENTS

The authors acknowledge partial support from COST Action IC1405 on Reversible Computation—Extending Horizons of Computing. The work of I. Lanese was supported in part by French ANR project DCore ANR-18-CE25-0007.

REFERENCES

1. R. Landauer, “Irreversibility and heat generated in the computing process,” *IBM J. Res. Develop.*, vol. 5, pp. 183–191, 1961, doi: [10.1147/rd.53.0183](https://doi.org/10.1147/rd.53.0183).
2. E. Fredkin and T. Toffoli, “Conservative logic,” *Int. J. Theor. Phys.*, vol. 21, pp. 219–253, 1982, doi: [10.1007/BF01857727](https://doi.org/10.1007/BF01857727).
3. C. Bennett, “Logical reversibility of computation,” *IBM J. Res. Develop.*, vol. 17, pp. 525–532, 1973, doi: [10.1147/rd.176.0525](https://doi.org/10.1147/rd.176.0525).
4. T. Yokoyama, H. B. Axelsen, and R. Glück, “Principles of a reversible programming language,” in *Proc. Comput. Frontiers*, 2008, pp. 43–54, doi: [10.1145/1366230.1366239](https://doi.org/10.1145/1366230.1366239).
5. K. Perumalla, *Introduction to Reversible Computing*. Boca Raton, FL, USA: CRC Press, 2014, doi: [ISBN 9781439873403](https://doi.org/10.1007/9781439873403).
6. V. Danos and J. Krivine, “Reversible communicating systems,” in *Proc. CONCUR*, 2004, vol. 3170, pp. 292–307, doi: [10.1007/978-3-540-28644-8_19](https://doi.org/10.1007/978-3-540-28644-8_19).
7. I. Phillips and I. Ulidowski, “Reversing algebraic process calculi,” *J. Logic Algebr. Program.*, vol. 73, pp. 70–96, 2007, doi: [10.1016/j.jlap.2006.11.002](https://doi.org/10.1016/j.jlap.2006.11.002).
8. I. Lanese, C. Mezzina, and J.-B. Stefani, “Reversing higher-order pi,” in *Proc. CONCUR*, 2010, vol. 6269, pp. 478–493, doi: [10.1007/978-3-642-15375-4_33](https://doi.org/10.1007/978-3-642-15375-4_33).
9. I. Lanese, I. Phillips, and I. Ulidowski, “An axiomatics approach to reversible computation,” in *Proc. FOSSACS*, 2020, vol. 12077, pp. 442–461, doi: [10.1007/978-3-030-45231-5_23](https://doi.org/10.1007/978-3-030-45231-5_23).
10. I. Ulidowski, I. Lanese, U. P. Schultz, and C. Ferreira, eds., *Reversible Computation: Extending Horizons of Computing - Selected Results of the COST Action IC1405, ser. LNCS*. Berlin, Germany: Springer, 2020, vol. 12070, doi: [10.1007/978-3-030-47361-7](https://doi.org/10.1007/978-3-030-47361-7).

11. M. Fowler, *Domain-Specific Languages*. Boston, MA, USA: Addison-Wesley, 2010, doi: [10.5555/1809745](https://doi.org/10.5555/1809745).
12. M. Yim et al., "Modular self-reconfigurable robot systems [Grand challenges of robotics]," *IEEE Robot. Automat. Mag.*, vol. 14, no. 1, pp. 43–52, Mar. 2007, doi: [10.1109/MRA.2007.339623](https://doi.org/10.1109/MRA.2007.339623).
13. M. W. Jorgensen, E. H. Ostergaard, and H. H. Lund, "Modular ATRON: Modules for a self-reconfigurable robot," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2004, pp. 2068–2073, doi: [10.1109/IROS.2004.1389702](https://doi.org/10.1109/IROS.2004.1389702).
14. U. Schultz, M. Bordignon, and K. Stoy, "Robust and reversible execution of self-reconfiguration sequences," *Robotica*, vol. 29, pp. 35–57, 2011, doi: [10.1017/S0263574710000664](https://doi.org/10.1017/S0263574710000664).
15. R. E. Fikes and N. J. Nilsson, "Strips: A new approach to the application of theorem proving to problem solving," in *Artif. Intell.*, vol. 2, no. 3/4, 1971, pp. 189–208, doi: [10.1016/0004-3702\(71\)90010-5](https://doi.org/10.1016/0004-3702(71)90010-5).
16. C. R. Garrett, T. Lozano-Pérez, and L. P. Kaelbling, "Backward-forward search for manipulation planning," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2015, pp. 6366–6373, doi: [10.1109/IROS.2015.7354287](https://doi.org/10.1109/IROS.2015.7354287).
17. P. Loborg, "Error recovery supporting manufacturing control systems," Ph.D. dissertation, School Eng., Linköping Univ., Linköping, Sweden, <http://urn.kb.se/resolve?urn=urn%3Anbn%3Ase%3Aliu%3Adiva-163284>, 1994, doi: [ISBN 9178713765](https://doi.org/10.1016/S0263574717000613).
18. J. Laursen, L. Ellekilde, and U. Schultz, "Modelling reversible execution of robotic assembly," *Robotica*, vol. 36, no. 5, pp. 625–654, 2018, doi: [10.1017/S0263574717000613](https://doi.org/10.1017/S0263574717000613).
19. I. Phillips, I. Ulidowski, and S. Yuen, "A reversible process calculus and the modelling of the ERK signalling pathway," in *Proc. Reversible Comput.*, 2013, vol. 7581, pp. 218–232, doi: [10.1007/978-3-642-36315-3_18](https://doi.org/10.1007/978-3-642-36315-3_18).
20. J. Hoey and I. Ulidowski, "Reversing imperative parallel programs and debugging," in *Proc. Reversible Comput.*, 2019, vol. 11497, pp. 108–127, doi: [10.1007/978-3-030-21500-2_7](https://doi.org/10.1007/978-3-030-21500-2_7).

IVAN LANESE is an Associate Professor with the University of Bologna, Bologna, Italy. He acted as the Vice Chair of the COST Action IC1405 on Reversible Computation. He is interested in formal methods and concurrency theory, with special focus on reversibility and reversible debugging. Further information is available at <https://www.unibo.it/sitoweb/ivan.lanese/en>. He can be contacted at ivan.lanese@gmail.com.

ULRIK P. SCHULTZ is the Professor in aerial robotics with the University of Southern Denmark, Odense, Denmark. He recently participated in the COST Action IC1405 on reversible computing where he chaired the Working Group on Applications. He is interested in programming languages for robotics, his full biography is available at <http://www.sdu.dk/staff/ups> and he can be contacted at ups@sdu.dk.

IREK ULIDOWSKI is an Associate Professor with the University of Leicester, Leicester, U.K. He chaired the COST Action IC1405 on reversible computation. His interests include reversible computation and its application, and concurrent systems. Further information is available at <https://www.cs.le.ac.uk/people/iu3>. He can be contacted at irekulidowski@gmail.com.