

Separating Algorithmic Thinking and Programming

Maurice Chandoo

chandoo@thi.uni-hannover.de

Institut für Theoretische Informatik

Hannover, Germany

ABSTRACT

We show how the first steps of algorithmic thinking and programming can be trained separately. The learner is not assumed to have any prior experience. A general framework and a sequence of training tasks is described and tested in practice with two 10th graders. Both were able to write relatively complex programs using only pen and paper within two days.

To train algorithmic thinking, computational problems are presented as games to the learner. Roughly speaking, a winning strategy corresponds to an algorithm which solves the problem. Thus, if the learner consistently wins the game for various instances then this indicates that they have found an algorithm. We describe a general mechanism to translate a computational problem into such a game. For the programming part, the learner is shown how a program can be constructed from traces. Programs are specified in a simple language which depends on the underlying model of computation (think of Turing machines, pushdown automata or instruction set architectures); such a model can be seen as a notional machine.

KEYWORDS

trace-based programming, model of computation, notional machine, programming education

1 INTRODUCTION

Complex tasks are usually taught by breaking them into smaller ones which can be trained separately. Writing a computer program requires to devise an algorithm for a given task and then implement it in a programming language. We shall call the ability to devise an algorithm *algorithmic thinking* and the task of implementing it *programming* (some prefer to call the latter coding).

To effectively train something the learner requires feedback to assess their success or performance. This enables them to appropriately adjust their course of action. Ideally, the feedback is immediate, informative and easy to obtain. To train algorithmic thinking separately, the learner needs to be able to verify a potential solution without having to implement it as a program. One way to do this is to give an informal description of the algorithm to someone else who verifies it. However, coming up with a description can be difficult by itself, especially for learners with limited means to express themselves verbally, which makes feedback neither immediate nor easy to obtain. We describe a game which enables the learner to quickly and easily verify their solution. Playing this game essentially amounts to executing the algorithm and a recording of the game corresponds to a trace (therefore playing it is also useful because the implicitly produced traces can be utilized to construct a program which implements the algorithm).

A notional machine is described as “the general properties of the machine that one is learning to control” by du Boulay who introduced the concept [6]. In the context of a programming language a notional machine can be seen as an execution model for the language which describes the semantics of its various constructs. Forming a mental representation of such a machine is one of the difficulties that novice programmers face. Sorva has given an extensive survey on notional machines and their relation to different research threads in programming education [15]. One of the conclusions of the survey is that it is recommended to have notional machines as an explicit learning goal, which we follow in our approach.

Models of computation are a central part of our framework and form a subclass of notional machines. A model of computation can be seen as an abstract description of a machine: it consists of buttons (operations) and indicator lamps (predicates), and it resides in a particular state (machine state). The indicator lamps (partially) describe its current state. When a button is pressed the machine changes its state depending only on its current state and which button has been pressed. For example, a Turing machine with tape alphabet $\{0, 1, \square\}$ has 5 buttons and 3 indicator lamps. For each character x from the alphabet it has an indicator lamp which is on iff the current cell contains x and a button which writes x to the current cell. It also has buttons to move the head one cell to the left and right. Its state consists of the tape contents and the position of the head. This concept has also been described in [14] to uniformly describe various kinds of automata.

Besides automata, it can also be used to describe computing agents such as a drawing turtle [11] or a lumberjack [5]. For example, in a dimly lit maze the agent’s actions (operations) could be to turn left by 90 degrees or move a step forward and the agent can see whether there is a wall in front and whether the exit is reached (predicates). The machine state describes the maze’s layout and the agent’s position. A limitation of models of computation is that the machine state cannot change unless an operation is executed. For example, an agent in a dynamic real-time environment cannot be described as such a model. Our framework (game, programming language & method) is applicable to anything that can be framed as a model of computation.

Our motivation for finding a way to teach algorithmic thinking and programming separately comes from the experience that students who struggle the most with programming tasks tend to think about algorithms in terms of a programming language ([12] explains why this is problematic). A reason for this might be that they equate algorithms with programs. Since it is usually more difficult to liberate someone from a misconception than to prevent it from happening, we tried to find a way of letting the learner experience that an algorithm can be constructed without even being aware of a programming language. Futschek defines algorithmic thinking as a set of abilities related to constructing and understanding algorithms

and also argues that it can be taught independently of programming [8]. He and Moschitz also propose playing algorithms as means to verify them [9]. Moreover, we also wanted to emphasize that programming requires to take a myriad of details into account and even if an algorithm is easy to conceive, implementing it can be surprisingly difficult.

In general, the complexity of translating a mental representation of an algorithm into a program seems to be underestimated. In [3] we show how to deal with this translation process systematically. The programming method presented here is a simpler variant thereof (the difference is that operations and predicates are already given). Hilton et al. describe a similar approach in [10]. Both have in common that one starts with executing the algorithm by hand for a concrete input and then generalizes the steps to build a program incrementally. A related approach called direct-manipulation programming is tested in [1]. The idea is that the learner can manipulate a program in terms of its traces using special software which provides a visual link between the concrete values in the traces and their abstract representations in the code. These works emphasize the value of traces as means to bridge the gap between the mental representation of an algorithm and its implementation.

In our approach a trace for some input is written first (or generated by playing the game) and used as specification of the program's expected behavior on that input. Then the learner performs two tasks on the trace to add additional information. One task is to assign a program state to each row of the trace. From this enriched trace a (partial) program is derived that is consistent with the trace. This process is repeated for additional traces until the program is complete. When adding additional traces the program is always modified such that it remains consistent with the previous traces, otherwise a contradiction occurs. Compared to ad hoc programming where the learner is left alone with figuring out how to construct programs, this approach imposes a structure on this process with clearly defined steps, which facilitates detecting errors.

In practice, the two 10th graders quickly understood the method and were able to intuitively apply it to construct complex programs. The method can be carried out using only pen & paper. An advantage of working away from computers can be that it eliminates the temptation of blindly guessing parts of a program and verifying its correctness by running it on some arbitrary input [4].

2 MACHINE-COMPUTER GAME

A model of computation consists of a set of machine states S , a set of operations (total functions from S to S) and a set of predicates (total functions from S to $\{0, 1\}$). For example, a counter machine with k registers can hold a non-negative integer in each register; a machine state is a sequence of k such numbers. Each register can be incremented or decremented by one (decrementing a register that contains 0 has no effect). This means the machine has $2k$ operations. For each register it has a predicate which holds iff that register contains 0.

A simple task on a counter machine with 3 registers is to copy the contents of register 1 to register 2. More specifically, the machine is initialized with an unknown state and a (human) computer has to operate the machine such that register 1 and 2 contain the value that register 1 had in the beginning. The difficulty is that the

indicator lamps (predicates) only reveal partial information about the machine state, i.e. whether a register contains 0 or not.

This can be played as a 2-player game where one person assumes the role of the machine and the other is the computer. The machine starts by choosing the initial machine state and writes it down on a piece of paper which the computer cannot see. Moreover, the machine has a card for each predicate with 'on' written on one side and 'off' on the other. These cards are used to simulate the indicator lamps for the computer. The computer can tell the machine to execute an operation. The machine writes down the new machine state after applying this operation and updates the indicator lamps by possibly flipping the cards. When the computer says 'end' the machine shows the hidden paper to the computer. The computer wins if the final contents of register 1 and 2 contain the value which was in register 1 in the beginning.

A trivial winning strategy for the computer is to determine the contents of register 1 by counting (decrement until 0) and then set register 1 and 2 to this value. This strategy can be generalized to solve arbitrary tasks: determine the initial machine state, solve the problem without the machine, enter the result into the machine. The issue is that no algorithm is executed on the machine itself. To prevent the learner from resorting to this strategy an additional requirement must be made that it should be possible to teach the winning strategy to a person who cannot count (e.g. a preschool child). More generally, the computer is only allowed to remember a constant amount of information independent of the input.

An actual computer can be used to simulate the machine which makes the game much faster and less error-prone. However, initially playing the game with another person can facilitate relating computation to 'physical reality'. For instance, in the case of the counter machine the machine-player can use a stack of beer mats to represent each register's content (number of beer mats = register value) instead of writing it down on a paper. This allows the learner to associate the abstract quantities in the registers with something tangible. By writing characters from a fixed alphabet on the beer mats this also extends to stack machines (defined below) and makes it intuitively clear why counter machines are isomorphic to stack machines over a unary alphabet. In [7] it is found that programming a physical object versus programming a simulation thereof does not significantly affect learning gains. Assuming this generalizes, using a computer to simulate the game with the beer mats might just be as effective; since the computer cannot see the beer mats until the end, the learner should also be able to play the machine in this simulation.

The learner can train the first steps of algorithmic thinking by trying to find algorithms for simple models of computation and tasks. If the learner consistently wins the game as computer (without resorting to the trivial winning strategy described above) then this indicates that they found a correct algorithm. The feedback for the learner is relatively quick (provided the inputs are not too large) and easy to obtain since there is no need to convert the mental representation of the algorithm into a program or an informal description.

A stack machine with k registers over an alphabet Σ can hold a string over Σ (including the empty string ϵ) in each register; a machine state is a sequence of k such strings. For each register it has an operation to remove the last character or append a given

character from Σ ; removing the last character from an empty string has no effect. For each register and each character x from Σ it has a predicate which holds iff x is the last character of the string in that register.

In the following, sequences of training tasks for the counter and stack machine are given. The number in the parenthesis indicates the number of registers. We write R_i and R_i' for the initial and final value of the i -th register, respectively. For a string x let $|x|$ denote its length. The first sequence contains tasks for counter machines and the second one for stack machines over the alphabet $\{A, B, C\}$.

- MOVE (2): $R2' = R1$
- COPY (3): $R1' = R2' = R1$
- ADD (3): $R3' = R1 + R2$
- MULT (4): $R3' = R1 \cdot R2$
- PARITY (1): $R1' = R1 \bmod 2$
- PRIME (6): $R2' = 1$ if $R1$ is prime, otherwise 0
- REVERSE (2): $R2' =$ the reverse of $R1$
- MOVE (3): $R2' = R1$
- COPY (3): $R1' = R2' = R1$
- CONCAT (3): $R3' = R1$ concatenated with $R2$
- REPEAT (4): $R3' = R1$ concatenated with itself $|R2|$ times
- SUBSTR (4): $R4' =$ substring of $R1$ starting from the $(|R2| + 1)$ -th character with length $|R3|$
- BINARY (4): $R2' = x$ times A where x is the value of the binary number in $R1$ ($A=1, B=0$, last character is LSB)

For the task SUBSTR the initial machine state must be chosen such that it describes a valid substring, i.e. $|R2| + |R3| \leq |R1|$. The game for these tasks can be played at [2].

This game can be used for any computational problem which can be framed as the task of reaching a particular target machine state from an unknown initial machine state in some model of computation. Two examples are given in the following.

There are four sacks $A1, B1$ and $A2, B2$. The sacks $A1$ and $B1$ belong to person 1 and the sacks $A2$ and $B2$ to person 2. $A1$ contains some apples and $B2$ some bananas. Person 1 and 2 agree to trade one apple for one banana and as many as possible. A small child which cannot count is given the task of completing this transaction. It can see whether a sack is empty (4 predicates) and it can put one apple from $A1$ into $A2$ (transfer apple from person 1 to 2) and it can put one banana from $B2$ into $B1$ (2 operations). The solution is to transfer apples and bananas in alternation until person 1 runs out of apples or person 2 runs out of bananas.

In a gummy bear factory there are two containers: one with sugar and one with food dye. Beneath these two containers is a mixer. A gummy bear consists of three pieces of sugar and two pieces of food dye. There is a button which drops a piece of sugar in the mixer and another one which drops a piece of food dye in the mixer, provided the respective container is not empty. Moreover, there are two indicator lamps which show whether the sugar or food dye container is empty. The task is to put as many piece of sugar and food dye as possible into the mixer in the correct ratio (3:2) to produce gummy bears. For example, if there were 60 pieces of sugar and 45 pieces of food dye in the beginning then there should be 60 pieces of sugar and 40 pieces of food dye in the mixer in the end. The challenge is that one does not know how many

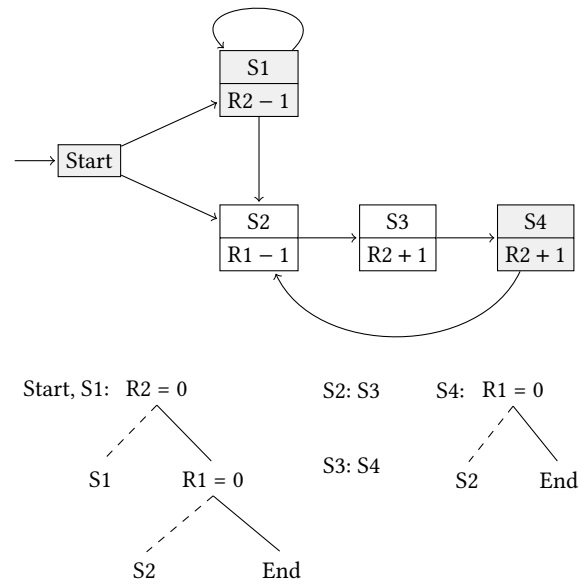


Figure 1: Example program for a counter machine

pieces of sugar and food dye are in the containers in the beginning. In certain cases it is not possible to achieve a correct ratio (e.g. 5 pc. sugar, 4 pc. food dye) Bonus question: how can the indicator lamps be modified to solve this issue?

3 PROGRAMMING WITH TRACES

We define (machine) programs w.r.t. a model of computation. A program is represented as control flow graph (CFG) with a designated start node. Nodes of the CFG are called program states. Each program state is associated with an operation (except the start node) and with a binary decision tree (BDT) which specifies what program state to visit next or whether the program should terminate. The inner nodes of such a tree are labeled with predicates and the leaves with program states (except 'Start') or 'End'.

Figure 1 shows a program for a counter machine with two registers. If a predicate does not hold then the left child (dashed line) is taken. A program state is gray if the program can terminate at that state. The program doubles the value of register 1 and writes it to register 2, clearing register 1 in the process. Table 1 shows the trace produced when executing this program on input (2,1). Blank cells in the register columns indicate that the value has not changed.

Observe that the CFG is redundant since it can be reconstructed from the BDTs. The outgoing edges of a program state are given by the leaves of its BDT. Therefore a list of program states along with their operation and BDT is a complete description of a machine program. Nonetheless, we deem the CFG important since it provides a different perspective on the program for the learner.

Since this definition works for any model of computation it can be used to describe programs for various kinds of automata, computing agents and even instruction set architectures with certain restrictions. In contrast to assembler programs the control flow is explicit and jump/branching operations are replaced by BDTs. The

Table 1: Execution trace of program from Fig. 1

Program State	Operation	R1	R2	Predicate Sequence
Start		2	1	$R2 = 0 \dashrightarrow$
S1	$R2 - 1$		0	$R2 = 0 \rightarrow R1 = 0 \dashrightarrow$
S2	$R1 - 1$	1		
S3	$R2 + 1$		1	
S4	$R2 + 1$		2	$R1 = 0 \dashrightarrow$
S2	$R1 - 1$	0		
S3	$R2 + 1$		3	
S4	$R2 + 1$		4	$R1 = 0 \rightarrow$
End				

lack of abstraction mechanisms leads to a compact definition which can be quickly understood. Subprogram calls and more advanced mechanisms can be introduced later on.

Programming Method. Take a look at Table 1 again. The last column ‘Predicate Sequence’ describes the path that was taken in the BDT of the program state in that row. For example, in the second row the machine state is (2, 0) after executing the operation $R2 - 1$. Thus, in the BDT of S1 the predicate $R2 = 0$ is checked first (true) and then $R1 = 0$ (false), which leads to S2. If we append the program state in the next row to the predicate sequence we get ‘ $R2 = 0 \rightarrow R1 = 0 \dashrightarrow S2$ ’. For the second last row this would be ‘ $R1 = 0 \rightarrow \text{End}$ ’ and for the third row simply ‘S3’.

This trace table can be used to derive a partial version of the program in Figure 1. The program states and their operations can be taken from the table directly. The BDT of a program state can be constructed by combining the paths resulting from the predicate sequences. Whenever an inner node in a BDT has only one child we add a second child which leads to a special program state \emptyset representing undefined behavior. Applying this procedure leads to the partial program shown in Figure 2. The program state \emptyset has BDT ‘End’ and an arbitrary operation can be assigned to it (represented by * in the CFG). Further traces can be added and incorporated until the program is complete (no more undefined behavior).

In this case we already had the trace table. How can the learner obtain such a table? Playing the machine-computer game yields a trace containing the operations and machine states. The learner then has to fill in the column for the program states manually. Whenever two rows have different operations they cannot be assigned the same program state. The converse is not necessarily true, i.e. in Table 1 the fourth and fifth row have the same operation but different program states. The predicate sequences can also be filled in manually but this is more tricky. Alternatively, the machine-computer game can be slightly modified to yield these sequences as well. The predicates are hidden from the computer. In order to see one the computer has to ask the machine. The order in which the computer asks to see the predicates is the predicate sequence. After an operation is executed all predicates are hidden from the computer again.

The key challenge for the learner is to understand under what circumstances two rows correspond to the same program state. After obtaining such a table the rest of this procedure can be automated.

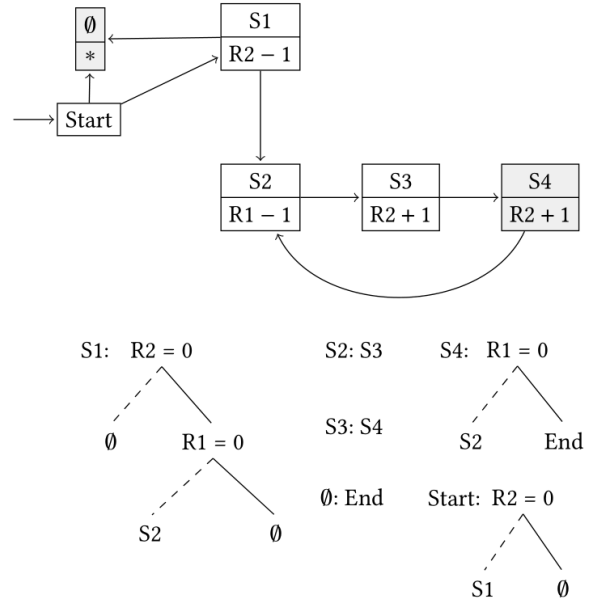


Figure 2: Partial program constructed from Table 1

A less rigid way of applying this method is to leave out the predicate sequences. In this case the traces are only used to construct the CFG until one is convinced that no edges and program states are missing anymore. Then the BDTs are developed separately, knowing which states must occur as leaves. Manually constructing BDTs makes it easier to exploit symmetries in the program behavior. A hybrid approach would be to partially determine a BDT using a few predicate sequences and then complete it manually to exploit symmetries.

This method has a consistency guarantee. Suppose that we want to construct a program P from some traces T containing program states and pred. sequences. We derive the partial program P' from T . If P is consistent with T , i.e. its execution produces the traces T , then P' will either behave identically to P or run into undefined behavior. This means the partial program in Fig. 2 behaves identically to the complete one for all inputs (x, y) such that $x \geq 1$ and $y = 1$.

4 EXPERIENCE REPORT

We tested our framework and the previous tasks with two 10th graders A and B (both 16 years old) who had no prior experience with programming. They attended a 5-day workshop for ca. 3 hours per day. They were given work sheets with multiple problems. We engaged only if we noticed a conceptual misunderstanding, the interpretation of a problem was unclear or to verify solutions. They were allowed to interact, but only did so to compare their solutions for the first worksheet.

Worksheets. Worksheet 1 contains the two problems described at the end of section 2 (apples & bananas and gummy bears). Solutions for these problems should be exemplified by filling out tables which are provided as well. These tables contain a column for the operation and 4 resp. 2 for the machine states. The first rows are already filled

in, i.e. the initial machine state is fixed. Algorithms or programs are not mentioned.

Worksheet 2 explains the machine-computer game for counter machines using beer mats. Instead of predicate cards the machine-player has to notify the computer as soon as a stack becomes empty. The problem is to find winning strategies for MOVE and COPY without counting. Worksheet 3 contains a description of counter machines and the 6 tasks for them. The problem is to find algorithms for each of them.

Worksheet 4 contains the program from Fig. 1 and two traces when executing it (similar to Table 1). The first problem is to write a trace when executing this program for a given input. The aim of this problem is to make sure that the semantics of machine programs are clear. The second problem is to find out how the program can be partially reconstructed from the given traces. The purpose of this problem is to raise awareness of the connection between traces and programs. The last problem is to write programs for the 6 previous tasks.

Worksheet 5 describes the machine-computer game for stack machines using beer mats. The problem is to find winning strategies for the first three tasks for stack machines without memorizing the stack contents. Worksheet 6 describes stack machines and the seven tasks for it. The first problem is to find algorithms for each task and the second problem is to write programs implementing them. For the task BINARY it is explained how a binary number can be converted by summing the appropriate powers of two.

Similar to worksheet 1, preprinted trace tables are provided for the other worksheets as well in order to demonstrate algorithms in terms of their traces. They contain only columns for the operations and machine states. The first rows in these tables are given, i.e. the initial machine states (inputs) are fixed. This makes it easier to verify solutions.

Observation. In the following we give a chronological account of how the two dealt with these worksheets.

Both of them found correct algorithms for all tasks up to and including PARTY after just one hour (worksheets 1 through 3). They did not play the game at all (either on the computer or with beer mats) to come up with the algorithms. Both failed to find an algorithm for PRIME. When we asked them to play the game for the more difficult tasks, they usually needed two or three tries until they were able to reliably execute their algorithms.

After less than 20 minutes both had written the execution trace for the program on worksheet 4 and figured out that the sequence of program states in a trace describes a path through the CFG. Next, A decided to write a program for MULT first and B decided to write a program for ADD. We gave both a laptop with an interpreter where they could type in their programs to run them and explained the syntax of the interpreter. After 1 hour B claimed to have a correct program. We found an error during testing, which was resolved by B after another half an hour. A had a correct program for MULT after 2 hours. This concluded the first day.

On day 2 we demonstrated the programming method. During the demonstration we asked them how to proceed at certain points to check whether they could follow. This took 15 minutes. Then we gave them trace tables with an additional column for the predicate sequences to facilitate the use of the method. We asked them to

write a program for MULT using only pen & paper. Initially, A was not sure how to fill out the predicate sequence column: instead of writing down the sequence of predicates considered for that particular input, A seemed to write down the subset of predicates that has to be considered at the corresponding program state over all inputs.

After less than 45 minutes A was done. We typed it into the computer and it worked correctly immediately without further modifications. After 90 minutes B was done. We typed it in but it returned incorrect results in certain cases. It took B only a few minutes of looking at the trace produced by the program to identify the cause of the problem (in two places the wrong register was referenced). After this correction the program worked correctly. B's program for multiplication had 15 program states and implemented a more efficient algorithm than A's whose program had only 9; for certain inputs B's program was nearly twice as fast. Their programs can be found in the appendix

It took A and B another 30 minutes to write programs for the other tasks on the counter machine using only pen & paper. All of their programs worked immediately when typed in with the exception of B's program for ADD (a predicate was missing). A said that the method was useful for MULT but not so much for the other tasks since the programs were quite simple.

Both solved worksheet 5 before the end of day 2. It took both 20 minutes to understand the game and find correct winning strategies for all problems with the following exception. B's strategy for COPY was incorrect because it involved implicitly counting the length of a string. This concluded day 2.

The remaining 3 days they were working on worksheet 6. A found algorithms for all of the tasks. B found algorithms for all tasks but BINARY. We gave B a trace produced by an algorithm which solves it. After 12 minutes B was able to understand the algorithm from the trace alone and execute it by playing the game. B wrote correct programs for all tasks but CONCAT and BINARY before the workshop ended. A wrote correct programs for all tasks but REPEAT. While B used the method for all problems, A tried to write programs for REPEAT and SUBSTR on the computer. In both cases A failed to write correct programs without the method. After two hours A gave up on the program for SUBSTR because it had multiple errors. A tried again on paper using the method which lead to a correct program for SUBSTR after 90 minutes. A expressed surprise that programming on the computer can be more difficult than on paper.

In general, when A and B used the method their programs were either immediately correct or had a single error which they resolved after one correction. If they had an incorrect program we gave them the option to continue with the next task and return to the current one later on.

We noticed that after MULT both only used the method to construct the CFG and did not write predicate sequences. Instead, they recognized certain patterns in the BDTs and that some subtrees occurred multiple times in a program. They would draw a colored circle around the first occurrence of such a subtree and then used the color to reference it.

Interpretation. Both had a relatively high level of abstraction as exhibited by the fact that they did not need to play the game.

Also, the connection between the beer mat game and the counter and stack machines was obvious to them. It took them only a few minutes to come up with algorithms for most tasks. We suppose that their performance during the workshop was above average with respect to their age group.

We claimed that the machine-computer game can help learners to train the first steps of algorithmic thinking. However, the two did not play it for themselves. Why is that? Clearly, playing the game itself does not help in finding an algorithm. The learner needs to come up with an algorithmic idea beforehand. The game serves two other purposes. It gives the learner a concrete idea of what constitutes an algorithm (i.e. what they are looking for) and it can be used by the learner to confirm that an algorithm works correctly without having to implement it. If the learner is already confident in their algorithm's correctness then playing the game provides no additional value in that regard. In particular, we believe that the game can help a learner to build confidence in their ability to reason (about the correctness of algorithms) until the point where it is not needed anymore. A and B were already beyond that point and thus found no use in playing the game to check their algorithms. Nonetheless, we found it useful to let them play the game to confirm that they had found a correct algorithm.

The programs that they had to write were quite large. For example, both their programs for SUBSTR had 27 program states and over 60 nodes in the BDTs. Writing programs of such complexity is not trivial. We let them try the ad hoc approach on the computer to see how it compares to the method. For instance, A's ad hoc program for MULT had 10 program states and 44 nodes in the BDTs whereas the one constructed with the method had only 9 program states and 31 nodes. A required less than half of the time to write the program with the method. An even stronger example of the method's usefulness is that A failed to write a program for SUBSTR without it. We would expect that a weaker learner can profit even more from the method because they would already fail to write simpler programs ad hoc.

However, programming with traces alone is not a silver bullet. In order to write programs effectively it has to be combined with other techniques. A good example is the task CONCAT for which B failed to write a correct program using the method. It can be solved using only 4 calls to a subprogram for REVERSE. Knowing how to compile a program with subprogram calls makes it trivial to write a program for CONCAT if a program for REVERSE is already known. In contrast, trying to write a program for it using only the method is tedious and error-prone because it requires a lot of needless repetition.

Similarly, neither A nor B used predicate sequences for the tasks on the stack machine. Instead they recognized recurring subtrees and used them to efficiently construct the BDTs after having derived the control flow and program states with the method. Deriving BDTs strictly from traces in these cases would have been much more laborious.

The learner should be taught how the method can be used flexibly and how techniques such as subprogram calls can be integrated into its workflow; something that we did not do. For example, subprogram calls can be incorporated by adding them as operation to the model of computation. In this case the learners figured out themselves how to adapt the method to their needs.

We conjecture that writing the more complex programs would have been infeasible for the two within the given time without knowing the method. The fact that A tried to implement simple algorithms for REPEAT and SUBSTR ad hoc and failed despite strong analytical thinking skills and then managed to do it with the method shows that our programming method can be useful even for such learners and that weaker learners would probably face even greater difficulties if left without guidance on how to implement algorithms as machine programs.

5 CONCLUSION & FUTURE WORK

We described a theoretical framework consisting of models of computation and machine programs, a game which allows the learner to test and execute arbitrary algorithms within this framework, a method where a program is constructed from execution traces and a sequence of training tasks.

The presented framework is simple yet general. The syntax and semantics of machine programs can be inferred from a single example even without prior knowledge of programming languages. Despite this simplicity it is also possible to use such programs to describe algorithms on the same level of abstraction as modern programming languages by choosing an adequate model of computation [3]. During the workshop the two 10th graders quickly understood the framework just from the description and problems on the worksheets without further instructions.

The game was not played by the 10th graders to verify their algorithms. From this one might infer that it provided no value. However, it also serves the purpose of implicitly describing what constitutes an algorithm (a certain kind of winning strategy). Before a learner can be asked to find algorithms they need to know what algorithms are. An abstract definition is likely too vague in order for the learner to understand what they are looking for. The alternative is to directly introduce programs. But this can cause the conflation of programming and algorithmic thinking, which we tried to circumvent. A learner which struggles to find a correct implementation might not understand whether the algorithm or the implementation is wrong. Instead of writing a program, the learner can represent their algorithm in terms of traces. During the workshop we found traces to be quite effective as means to communicate algorithms, especially when paired with a description of the idea behind the algorithm. Writing such traces implies the ability to mentally simulate the game. Thus, while the 10th graders did not play the game they must have simulated it mentally.

A disadvantage of the game compared to programs is that not every winning strategy corresponds to an algorithm. More specifically, any strategy which requires the computer to remember a non-constant amount of information does not correspond to a program. In general, it might not be easy for the learner to understand what constitutes a valid strategy. In the context of counting machines the requirement that a winning strategy should not depend on counting seemed to give a decent idea of what not to do.

The method to derive a program from traces worked well in practice. The two 10th graders were able to understand and use it surprisingly quickly. But the workshop also showed that in certain instances a rigid application of the method can be exhausting and error-prone. Therefore the learner should be taught basic forms of

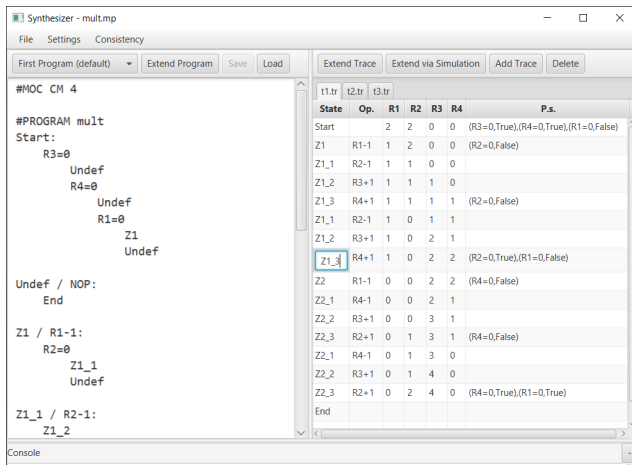


Figure 3: Synthesizer

abstraction early on to complement the method. Also, applying the method by hand requires a lot of concentration.

We plan to conduct the algorithmic part of this workshop for a larger group of younger students. The notion of a program will not be presented to the students in order to evaluate whether teaching algorithmic thinking using only the machine-computer game is feasible. They should learn the following abilities: recognizing whether a winning strategy constitutes an algorithm, describing an algorithm using traces and verbal descriptions, translating a given idea into an algorithm and designing algorithms from scratch.

A prototype which facilitates programming with traces has been developed by Yannik Mahlau [13]. The user can play a modified version of the machine-computer game which generates predicate sequences as well. Then state labels can be added to a trace and the software automatically synthesizes a (partial) program from it. When another trace is added, it also checks whether it is consistent with the current program. This makes the use of the method more comfortable by automating trivial tasks. We intend to test this software with students and evaluate how it compares to pen-&-paper application of the method.

REFERENCES

- [1] Michel Adam, Moncef Daoud, and Patrice Frison. 2019. Direct Manipulation versus Text-based Programming: An experiment report. In *Proceedings of the 2019 ACM Conference on Innovation and Technology in Computer Science Education, Aberdeen, Scotland, UK, July 15-17, 2019*. 353–359. <https://doi.org/10.1145/3304221.3319738>
- [2] Maurice Chandoo. 2019. Games and Interpreter for Counter and Stack Machines. <https://upsl.uber.space/aws19/info.txt>
- [3] Maurice Chandoo. 2019. A Systematic Approach to Programming. (2019). <https://arxiv.org/abs/1808.08989>
- [4] Quintin I. Cutts, Matthew Barr, Mireilla Bikanga Ada, Peter Donaldson, Steve Draper, Jack Parkinson, Jeremy Singer, and Lovisa Sundin. 2019. Experience Report: Thinkathon - Countering an “I Got It Working” Mentality with Pencil-and-Paper Exercises. In *Proceedings of the 2019 ACM Conference on Innovation and Technology in Computer Science Education, Aberdeen, Scotland, UK, July 15-17, 2019*. 203–209. <https://doi.org/10.1145/3304221.3319785>
- [5] Sylvia da Rosa Zipitria and Andres Aguirre Dorelo. 2018. Students Teach a Computer How to Play a Game. In *Informatics in Schools. Fundamentals of Computer Science and Software Engineering - 11th International Conference on Informatics in Schools: Situation, Evolution, and Perspectives, ISSEP 2018, St. Petersburg, Russia, October 10-12, 2018, Proceedings*. 55–67. https://doi.org/10.1007/978-3-030-02750-6_5
- [6] Benedict du Boulay. 1986. Some Difficulties of Learning to Program. *J. Educational Computing Research* 2, 1 (1986).
- [7] Grégoire Fessard, Patrick Wang, and Ilaria Renna. 2019. Are There Differences in Learning Gains When Programming a Tangible Object or a Simulation?. In *Proceedings of the 2019 ACM Conference on Innovation and Technology in Computer Science Education, Aberdeen, Scotland, UK, July 15-17, 2019*. 78–84. <https://doi.org/10.1145/3304221.3319747>
- [8] Gerald Futschek. 2006. Algorithmic Thinking: The Key for Understanding Computer Science. In *Informatics Education - The Bridge between Using and Understanding Computers, International Conference in Informatics in Secondary Schools - Evolution and Perspectives, ISSEP 2006, Vilnius, Lithuania, November 7-11, 2006, Proceedings*. 159–168. https://doi.org/10.1007/11915355_15
- [9] Gerald Futschek and Julia Moschitz. 2010. Developing Algorithmic Thinking by Inventing and Playing Algorithms. In *Constructionism, Paris*.
- [10] Andrew D. Hilton, Genevieve M. Lipp, and Susan H. Rodger. 2019. Translation from Problem to Code in Seven Steps. In *Proceedings of the ACM Conference on Global Computing Education (CompEd '19)*. ACM, New York, NY, USA, 78–84. <https://doi.org/10.1145/3300115.3309508>
- [11] Tobias Kohn and Dennis Komm. 2018. Teaching Programming and Algorithmic Complexity with Tangible Machines. In *Informatics in Schools. Fundamentals of Computer Science and Software Engineering - 11th International Conference on Informatics in Schools: Situation, Evolution, and Perspectives, ISSEP 2018, St. Petersburg, Russia, October 10-12, 2018, Proceedings*. 68–83. https://doi.org/10.1007/978-3-030-02750-6_6
- [12] Leslie Lamport. 2018. If You’re Not Writing a Program, Don’t Use a Programming Language. *Bulletin of the EATCS* 125 (2018). <http://eatcs.org/beatcs/index.php/beatcs/article/view/539>
- [13] Yannik Mahlau. 2020. *Synthesis of Machine Programs from Execution Traces*. Bachelor Thesis. Leibniz Universität Hannover. <https://www.thi.uni-hannover.de/fileadmin/thi/abschlussarbeiten/mahlau-ba.pdf>
- [14] Dana S. Scott. 1967. Some Definitional Suggestions for Automata Theory. *J. Comput. Syst. Sci.* 1, 2 (1967), 187–212. [https://doi.org/10.1016/S0022-0000\(67\)80014-X](https://doi.org/10.1016/S0022-0000(67)80014-X)
- [15] Juha Sorva. 2013. Notional machines and introductory programming education. *TOCE* 13, 2 (2013), 8:1–8:31. <https://doi.org/10.1145/2483710.2483713>

A PROGRAMS

This appendix includes the programs for MULT written by A and B, the program for SUBSTR written by B and the program for BINARY written by A. The BDTs are written in prefix notation. The first child is the false branch, e.g. (R2=0 Z2 Z5) means if register 2 contains 0 then go to Z5 else to Z2. On the stack machine the operation RxFy appends *y* to the string in register *x*, the operation NOP does not modify the machine state (no operation), the operation RxL removes the last character from the string in register *x* and the predicate RxL holds iff register *x* contains the empty string.

An online interpreter to execute the programs can be found on <https://upsl.uber.space/aws19/info.txt>

A.1 Multiplication (A)

Start:
(R3=0 Z7 (R4=0 Z8 (R1=0 (R2=0 Z1 End) End)))

Z1 / R1-1:
Z2

Z2 / R2-1:
Z3

Z3 / R3+1:
Z4

Z4 / R4+1:
(R2=0 Z2 Z5)

Z5 / R4-1:
Z6

Z6 / R2+1:
(R4=0 Z5 (R1=0 Z1 End))

Z7 / R3-1:
(R3=0 Z7 Z8)

Z8 / R4-1:
(R4=0 Z8 (R1=0 (R2=0 Z1 End) End))

A.2 Multiplication (B)

Start:
(R3=0 Z1 (R4=0 Z2 (R1=0 (R2=0 Z3 End) End)))

Z1 / R3-1:
(R3=0 Z1 (R4=0 Z2 (R1=0 (R2=0 Z3 End) End)))

Z2 / R4-1:
(R4=0 Z2 (R1=0 (R2=0 Z3 End) End))

Z3 / R1-1:
(R1=0 Z4 Z15)

Z4 / R2-1:
Z5

Z5 / R3+1:
Z6

Z6 / R4+1:
(R2=0 Z4 (R1=0 Z7 Z3))

Z7 / R1-1:
(R1=0 Z8 Z12)

Z8 / R4-1:
Z9

Z9 / R3+1:
Z10

Z10 / R2+1:
(R4=0 Z8 (R1=0 Z3 Z7))

Z12 / R4-1:
Z13

Z13 / R3+1:
(R4=0 Z12 End)

Z15 / R2-1:
Z16

Z16 / R3+1:

(R2=0 Z15 End)

A.3 Substring (B)

Start:
Lila

Lila / NOP:
(R1L (R3L (R4L Z1 Gruen) (R4L Z1 End)) (R4L Z1 End))

Gruen / NOP:
(R1A (R1B Z2C Z2B) Z2A)

Orange / NOP:
(R4A (R4B Z6C Z6B) Z6A)

Pink / NOP:
(R1A (R1B Z10C Z10B) Z10A)

Braun / NOP:
(R2A (R2B Z12C Z12B) Z12A)

Z1 / R4E:
Lila

Z2A / R4FA:
Z3

Z2B / R4FB:
Z3

Z2C / R4FC:
Z3

Z3 / R1E:
(R1L Gruen (R2L Z4 Z5))

Z4 / R4E:
Z5

Z5 / R2E:
(R2L Z4 Orange)

Z6A / R1FA:
Z7

Z6B / R1FB:
Z7

Z6C / R1FC:
Z7

Z7 / R4E:
Z8

Z8 / R3E:
(R3L Orange Z9)

Z9 / R4E:
(R4L Z9 Pink)

Z10A / R2FA:
Z11

Z10B / R2FB:
Z11

Z10C / R2FC:
Z11

Z11 / R1E:
(R1L Pink Braun)

Z12A / R4FA:
Z13

Z12B / R4FB:
Z13

Z12C / R4FC:
Z13

Z13 / R2E:
(R2L Braun End)

A.4 Binary (A)

Start:
(R1L (R2L 2LEER (R3L 3LEER (R4L 4LEER Z1))) Z30)

Z1 / R3FA:
Z2

Z2 / NOP:
(R1L (R1A Z101 Z3) End)

Z3 / R1E:
Z4

Z4 / R2FA:
Z5

Z5 / R4FA:
Z6

Z6 / R3E:
(R3L Z4 Z7)

Z7 / R3FA:
Z8

Z8 / R3FA:
Z9

Z9 / R4E:
(R4L Z7 Z2)

Z101 / R1E:
Z10

Z10 / NOP:
(R3L Z11 Z7)

Z11 / R4FA:
Z12

Z12 / R3E:
Z10

2LEER / R2E:
(R2L 2LEER (R3L 3LEER (R4L 4LEER Z1)))

3LEER / R3E:
(R3L 3LEER (R4L 4LEER Z1))

4LEER / R4E:
(R4L 4LEER Z1)

Z30 / NOP:
(R2L 2LEER1 End)

2LEER1 / R2E:
Z30