

This is a repository copy of *Time and Space Measures for a Complete Graph Computation Model*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/194716/>

Version: Published Version

Proceedings Paper:

Courtehouse, Brian and Plump, Detlef orcid.org/0000-0002-1148-822X (2022) Time and Space Measures for a Complete Graph Computation Model. In: Proceedings of the Thirteenth International Workshop on Graph Computation Models. International Workshop on Graph Computation Models, 06 Jul 2022, France. Electronic Proceedings in Theoretical Computer Science. Open Publishing Association, Nantes, pp. 23-44.

<https://doi.org/10.4204/EPTCS.374.4>

Reuse

This article is distributed under the terms of the Creative Commons Attribution (CC BY) licence. This licence allows you to distribute, remix, tweak, and build upon the work, even commercially, as long as you credit the authors for the original work. More information and the full terms of the licence here:

<https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

Time and Space Measures for a Complete Graph Computation Model

Brian Courtehouse and Detlef Plump

Department of Computer Science, University of York, York, UK

{bc956,detlef.plump}@york.ac.uk

We present a computation model based on a subclass of GP 2 graph programs which can simulate any off-line Turing machine of space complexity $O(s(n) \log s(n))$ in space $O(s(n))$. The simulation only requires a quadratic time overhead. Our model shares this property with Schönhage’s storage modification machines and Kolmogorov-Uspenskii machines [17]. These machines use low-level pointer instructions whereas our GP 2-based model uses pattern-based transformation rules and high-level control constructs.

1 Introduction

Schönhage’s storage modification machines (SMMs) and Kolmogorov-Uspenskii machines (KUMs) are graph-based computation models that do not use graph transformation rules. They have the remarkable feature of being able to simulate Turing machines using less space with only quadratic time overhead, using a uniform space measure [17]. Although these are graph-based models, the literature on rule-based graph transformation so far seems to have ignored the work of Schönhage [15] and Kolmogorov [11], and in particular the results of van Emde Boas [17] and Luginbuhl [12].

One fundamental difference between SMMs and KUMs on the one hand, and GP 2 on the other hand is that the inputs and outputs of SMMs and KUMs are strings, whereas GP 2 programs compute relations on graphs. Moreover, SMMs and KUMs differ from modern graph transformation languages in that they use low-level pointer instructions instead of pattern-based rules where the programmer can specify a subgraph and how it is transformed, allowing for a more natural implementation of graph algorithms.

An advantage of structured programming languages such as GP 2 is the absence of “go to” statements which are considered to be harmful [8]. Furthermore, GP 2 is based on the double-pushout approach for graph transformation, which comes with a large amount of theoretical results. We are not aware of any comparable theory for SMMs and KUMs.

The fact that the above mentioned space compression property applies to a high-level language is not obvious. Even so, we show in this paper that GP 2 exhibits the same space compression feature. Specifically, we show that a Turing machine using $O(s(n) \log s(n))$ space can be simulated in $O(s(n))$ space with quadratic time overhead, where $s(n)$ is an arbitrary function. The complexity class of the simulation is a strict subset of the original class (for non-constant $s(n)$). Rather than using full GP 2, we show that a subset of the language suffices to establish the compression result.

Note that another way to show space compression is to efficiently simulate SMMs in general. This would also mean other properties are applicable to GP 2, such as the real-time simulation of Turing machines shown by Luginbuhl [12]. This paper’s approach however comes with a concrete class of programs that show in much detail that the simulation is possible. In particular, we show *how* to establish a result akin to van Emde Boas’ in a rule-based graph transformation language such as GP 2. Moreover, the van Emde Boas paper lacks a lot of technical details, which we provide in our setting.

In Section 2, we define the subset of GP 2 sufficient to get our results. We also show how this subset finds matches for rules efficiently.

We show how we simulate Turing machines in Section 3. We describe the encoding of Turing machine configurations, give the simulation, provide an example, and prove correctness.

In Section 4 we identify a class of efficient computation models based on GP 2, which come with time and space complexity measures. We then show that our simulation gives rise to such a model.

Finally, in Section 5, we show the aforementioned time and space complexities of the simulation, and discuss the use of logarithmic versus uniform space measures.

2 Rule-Based Graph Transformation

Throughout this paper, we use a subset of GP 2 called SGP 2 (Small GP 2, defined in Subsection 2.1), which suffices to establish the simulation and compression result. This subset uses simpler host graphs. Specifically labels are limited to being either an integer or one of three characters, and lists of these are not allowed. Rules are simpler as well. They use neither label expressions with variables, nor application conditions.

In Subsection 2.2, we analyse the complexity of rule matching by showing a new variant of a Theorem that allows for constant-time rule matching.

2.1 Small GP 2

SGP 2 programs transform input graphs into output graphs, where graphs are directed and may contain parallel edges and loops. Both nodes and edges are labelled.

Definition 1 (Graph). Let \mathcal{L} be the set of labels defined by the grammar in Figure 1. A *graph* over \mathcal{L} is a system $\langle V, E, P, s, t, l, m \rangle$, where V and E are finite sets of nodes (or vertices) and edges, $P \subseteq V$ is the set of roots, $s: E \rightarrow V$ and $t: E \rightarrow V$ are source and target functions for edges, $l: V \rightarrow \mathcal{L}$ is the partial node labelling function and $m: E \rightarrow \mathcal{L}$ is the (total) edge labelling function.

A graph is *totally labelled* if l is a total function. Unlabelled nodes will only occur in the interfaces of rules and are used to relabel nodes. There is no need to relabel edges as they can always be deleted and reinserted with different labels.

Label	::=	Atom [Mark]
Atom	::=	Integer "L" "R" "I" empty
Mark	::=	red green blue grey dashed

Figure 1: Abstract syntax of host and rule graph labels

Both nodes and edges are labelled with either an integer, with the special characters "L", "R", "I", or with the constant empty. By convention, we draw items labelled with empty in pictures without labels. Figure 1 defines the labels that may occur in both rules and host graphs. Integers within a label are represented in base 10. In this paper however, we only use digits up to 2 in labels.

We call the graphs to which rules are applied *host graphs*, they consist of all totally labelled graphs.

Nodes and edges can be *marked* red, green or blue. In addition, nodes can be marked grey and edges can be dashed. For example, rules in CacheNext in Figure 10 contains red and unmarked nodes and blue edges. Marks are convenient, among other things, to record visited items during a graph traversal and to differentiate between similar structures.

The principal SGP2 programming construct used in this paper are graph transformation rules according to the double-pushout approach with injective matching. For example, the rules in CacheNext in Figure 10 consist of a left-hand graph and a right-hand graph which are specified graphically. The small numbers attached to nodes are identifiers, all other black text in the graphs are labels.

Definition 2 (Rule). A rule $\langle L \leftarrow K \rightarrow R \rangle$ consists of totally labelled graphs L and R , a graph K consisting of unlabelled nodes only, where $V_K \subseteq V_L$ and $V_K \subseteq V_R$.

When a rule is declared, such as in CacheNext in Figure 10, graph K is implicitly represented by the node identifiers in L and R (which much coincide). Nodes without identifiers in L are to be deleted and nodes without identifiers in R are to be created. Roots are represented by double circles.

Definition 3 (Graph Morphism). Given graphs L and G , a *graph morphism* $g: L \rightarrow G$ consists of two functions $g_V: V_L \rightarrow V_G$ and $g_E: E_L \rightarrow E_G$ that preserve sources, targets and labels, that is, $s_G \circ g_E = g_V \circ s_L$, $t_G \circ g_E = g_V \circ t_L$, $l_G \circ g_E = l_L$ and $m_G \circ g_E = m_L$. Moreover, we require that roots are preserved and reflected, that is, $g_V(P_L) \subseteq P_G$ and $g_V^{-1}(P_G) \subseteq P_L$.

Definition 4 (Rule Application). A rule $r = \langle L \leftarrow K \rightarrow R \rangle$ is applied to a host graph G as follows:

1. Find an injective graph morphism $g: L \rightarrow G$ that satisfies the *dangling condition*: no node in $g_V(V_L - V_K)$ is incident to an edge in $E_G - g_E(E_L)$.
2. Construct a host graph H from G as follows:
 - (a) Remove all edges in $g_E(E_L)$ and all nodes in $g_V(V_L) - g_V(V_K)$, obtaining a subgraph D of G .
 - (b) $V_H = V_D + (V_R - V_K)$ and $E_H = E_D + E_R$.
 - (c) For each edge $e \in E_D$, $s_H(e) = s_D(e)$. For each edge $e \in E_R$, $s_H(e) = s_R(e)$ if $s_R(e) \in V_R - V_K$, otherwise $s_H(e) = g_V(s_R(e))$. Targets are defined analogously.
 - (d) For each node $v \in V_D - g_V(V_K)$, $l_H(v) = l_D(v)$. For each node $v \in g_V(V_K)$ with $g_V(\bar{v}) = v$, $l_H(v) = l_R(\bar{v})$. For each node $v \in V_R - V_K$, $l_H(v) = l_R(v)$.
 - (e) For each edge $e \in E_D$, $m_H(e) = m_D(e)$. For each edge $e \in E_R$, $m_H(e) = m_R(e)$.
 - (f) $P_H = (P_G - g_V(V_L)) \cup P_R$.

We write $G \Rightarrow_{r,g} H$ to express that H results from G by applying r with match g .

Prog	::=	Decl {Decl}
Decl	::=	RuleDecl ProcDecl MainDecl
ProcDecl	::=	ProcId '=' ComSeq
MainDecl	::=	Main '=' ComSeq
ComSeq	::=	Com {';' Com}
Com	::=	RuleSetCall ProcCall
		if ComSeq then ComSeq [else ComSeq]
		try ComSeq [then ComSeq] [else ComSeq]
		ComSeq '!' '(' ComSeq ')' break
RuleSetCall	::=	RuleId {' [RuleId {';' RuleId}] '}'
ProcCall	::=	ProcId

Figure 2: Abstract syntax of Small GP2 programs

SGP2 programs follow the syntax described in Figure 2. We omit the the definition of RuleDecl because we do not use textual rule declaration in this paper. We declare rules graphically so they are

easier to read. An example of a graphical rule declaration can be seen in Figure 7. The structural operational semantics of programs is defined in Figure 3, where inference rules inductively define a small-step transition relation \rightarrow on *configurations*. In the setting of SGP 2, a configuration is either a command sequence together with a host graph, just a host graph or the special element fail:

$$\rightarrow \subseteq (\text{ComSeq} \times \mathcal{G}) \times ((\text{ComSeq} \times \mathcal{G}) \cup \mathcal{G} \cup \{\text{fail}\})$$

where \mathcal{G} is the set of SGP 2 host graphs. Configurations in $\text{ComSeq} \times \mathcal{G}$, given by a rest program and a host graph, represent states of unfinished computations while graphs in \mathcal{G} are final states or *results* of computations. The element fail represents a failure state. A configuration γ is said to be *terminal* if there is no configuration δ such that $\gamma \rightarrow \delta$.

Figure 3 shows the inference rules for the core commands of SGP 2. We use an older version of the semantics because it is simpler, and equivalent to the current one for terminating programs (which is all we care about for time complexity analysis) [7]. The rules contain meta-variables for command sequences and graphs, where R stands for a call of a rule set or of a rule, C, P, P', Q stand for command sequences, and G, H stand for host graphs. Branching statements without then/else clauses are omitted, because they can be defined with the other commands.

The transitive and reflexive-transitive closures of \rightarrow are written \rightarrow^+ and \rightarrow^* , respectively. We write $G \Rightarrow_R H$ if H results from host graph G by applying the rule set R , while $G \not\Rightarrow_R H$ means that there is no graph H such that $G \Rightarrow_R H$ (application of R fails).

$$\begin{array}{l}
[\text{call}_1] \frac{G \Rightarrow_R H}{\langle R, G \rangle \rightarrow H} \quad [\text{call}_2] \frac{G \not\Rightarrow_R}{\langle R, G \rangle \rightarrow \text{fail}} \\
[\text{seq}_1] \frac{\langle P, G \rangle \rightarrow \langle P', H \rangle}{\langle P; Q, G \rangle \rightarrow \langle P'; Q, H \rangle} \quad [\text{seq}_2] \frac{\langle P, G \rangle \rightarrow H}{\langle P; Q, G \rangle \rightarrow \langle Q, H \rangle} \quad [\text{seq}_3] \frac{\langle P, G \rangle \rightarrow \text{fail}}{\langle P; Q, G \rangle \rightarrow \text{fail}} \\
[\text{if}_1] \frac{\langle C, G \rangle \rightarrow^+ H}{\langle \text{if } C \text{ then } P \text{ else } Q, G \rangle \rightarrow \langle P, G \rangle} \quad [\text{if}_2] \frac{\langle C, G \rangle \rightarrow^+ \text{fail}}{\langle \text{if } C \text{ then } P \text{ else } Q, G \rangle \rightarrow \langle Q, G \rangle} \\
[\text{try}_1] \frac{\langle C, G \rangle \rightarrow^+ H}{\langle \text{try } C \text{ then } P \text{ else } Q, G \rangle \rightarrow \langle P, H \rangle} \quad [\text{try}_2] \frac{\langle C, G \rangle \rightarrow^+ \text{fail}}{\langle \text{try } C \text{ then } P \text{ else } Q, G \rangle \rightarrow \langle Q, G \rangle} \\
[\text{alap}_1] \frac{\langle P, G \rangle \rightarrow^+ H}{\langle P!, G \rangle \rightarrow \langle P!, H \rangle} \quad [\text{alap}_2] \frac{\langle P, G \rangle \rightarrow^+ \text{fail}}{\langle P!, G \rangle \rightarrow G} \quad [\text{alap}_3] \frac{\langle P, G \rangle \rightarrow^* \langle \text{break}, H \rangle}{\langle P!, G \rangle \rightarrow H} \\
[\text{break}] \langle \text{break}; P, G \rangle \rightarrow \langle \text{break}, G \rangle
\end{array}$$

Figure 3: Inference rules for Small GP 2 core commands

2.2 Time Complexity of Rule Matching

When analysing the time and space complexity of programs, we assume that these are fixed. This is customary in algorithm analysis where programs are fixed and running time or space is measured in terms of input size [2, 16]. In our setting, the input size is the *size* of a host graph, which we define to be the total number of nodes and edges. When we match a rule graph L into a host graph G , L is constant since it is given by the program, but G is not since it depends on the input.

Since unrestricted matching is very time-consuming we use rooted rules. They contain special nodes called *roots* that can match with other roots in the host graph in constant time, allowing rule matching to happen locally.

In order to achieve this efficient rule matching, we use a variant of the fast matching theorem first introduced in [4]. The trade-off with the original is that we can have a more general host graph, i.e. we allow the indegree to be unbounded, but there are more restrictions on rules, namely there needs to be a directed path from some root node to every node in the left-hand side.

To prove that matching can be implemented in constant time, we give Algorithm 1, which matches a graph L into a graph G . It is identical to the one given in [3], except for the following two details. First, we consider morphisms instead of pre-morphisms (which ignore labels) since label expressions with variables or lists are not part of SGP 2. Second, we adjust the definition of edge enumerations to take into account edge direction. Because these are the only changes from the algorithm in [3], a correctness proof would be very similar and is hence omitted.

The definition of edge enumeration is very similar to the one given in [3], except that every time an edge is required to be incident to a node, we require that node to be the source of the edge. This is to make sure edges are matched via their source node. Since we use host graphs of bounded outdegree only, this ensures the number of matches is bounded.

Definition 5 (Edge Enumeration). *Given a graph L and a node $p \in P_L$ (the set of roots of L), an edge enumeration e_1, \dots, e_n for p is a list of edges such that $s(e_1) = p$, and for each $i \geq 2$, $s(e_i)$ is the source or target of some edge in e_1, \dots, e_{i-1} .*

Note that the source of each edge in an enumeration of p is reachable from p via a directed path. Furthermore, the union of edge enumerations over all roots of L only cover edges reachable from a root via a directed path. If we restrict the left-hand side of rules to only have nodes reachable from a root, the source of each edge must be reachable from a root, and so must the edge itself. Matching algorithms for other rules are out of the scope of this paper and can be found in [3].

Definition 6 (Fast Rule). *A rule $L \Rightarrow R$ is fast if each node in L is reachable from some root via a directed path.*

The other requirements for a rule to be fast from [4] are omitted since they concern variable labels and application conditions, which do not exist in SGP 2.

Definition 7 (Partial Morphism). *A graph morphism $g: L' \rightarrow G$ is a partial morphism $g: L \xrightarrow{\text{par}} G$ if $L' \subseteq L$. We say that g is undefined for the items in $L - L'$ and call L' the domain of g , denoted by $\text{Dom}(g)$.*

Consider partial morphisms $f, g: L \xrightarrow{\text{par}} G$ such that $f(x) = g(x)$ for all items x in $\text{Dom}(g)$. We say that f extends g by a node v if $\text{Dom}(f)_V = \text{Dom}(g)_V \cup \{v\}$ and $\text{Dom}(f)_E = \text{Dom}(g)_E$. We say that f extends g by an edge e if $\text{Dom}(f)_E = \text{Dom}(g)_E \cup \{e\}$ and $\text{Dom}(f)_V = \text{Dom}(g)_V \cup \{s_G(e), t_G(e)\}$.

In order to analyse the complexity of Algorithm 1 below, we need to assume that the extending a partial morphism by a single node or edge takes constant time. Constant-time integer comparison is a common assumption in algorithm analysis. Only in applications that require computing with large integers is a more sensitive model needed [2].

Theorem 1 (Fast Rule Matching). *Matching can be implemented to run in constant time for fast rules, using host graphs with a bounded outdegree that contain a bounded number of roots.*

Proof. We argue that Algorithm 1 implements matching as described by the theorem.

For this proof, we show that the number of times a morphism is updated with a single node or edge is bounded. Note that L has a bounded number of nodes and edges. So subsets of nodes or edges of L such as P_L and edge enumerations in L are also bounded. The while loop has constantly many iterations because it terminates once every node in the bounded set P_L is tagged, which happens at least once in every iteration. The for loop has constantly many iterations because it iterates over an edge enumeration.

Algorithm 1: A graph matching algorithm**Input:** Graphs L and G ; and for each $p \in P_L$, an edge enumeration e_{p_1}, \dots, e_{p_n} **Output:** The set A of all injective graph morphisms $L \rightarrow G$ that are root-preserving and root-reflecting

```

1  $A \leftarrow \{h : L \xrightarrow{par} G \mid \text{Dom}(h) = \emptyset\}$ 
2 while there is an untagged root  $p \in P_L$  do
3    $A_0 \leftarrow \{h : L \xrightarrow{par} G \mid h \text{ is injective, root-preserving and root-reflecting, and there exists}$ 
    $h' \in A \text{ such that } h \text{ extends } h' \text{ by } p\}$ 
4   tag  $p$ 
5   for  $i = 1$  to  $n$  do
6      $A_{p_i} \leftarrow \{h : L \xrightarrow{par} G \mid h \text{ is injective, root-preserving and root-reflecting, and there exists}$ 
      $h' \in A_{p_{i-1}} \text{ such that } h \text{ extends } h' \text{ by } e_{p_i}\}$ 
7     if  $s(e_{p_i}) \in P_L$  then tag  $s(e_{p_i})$ 
8     if  $t(e_{p_i}) \in P_L$  then tag  $t(e_{p_i})$ 
9    $A \leftarrow A_{p_n}$ 
10 return  $A$ 

```

In line 3, a partial morphism is added to A_0 for each match of the root p in G . Since we use host graphs with a bounded number of roots, the number of these matches is constant as well. Each of these morphisms extends the empty morphism by one node, so there are constantly many single-node morphism updates.

In line 8, we extend morphisms by an edge e_{p_i} , whose source is already in the extended morphism by definition of edge enumerations. Since G has bounded outdegree, the number of matches for e_{p_i} is bounded, and each morphism can only be extended in constantly many ways. The fact that there can only be constantly many morphisms can be shown by induction on A_0, A_1, A_2, \dots . By the previous paragraph, A_0 has bounded size. For $k \geq 0$, if A_k has bounded size, then so does A_{k+1} since each of the morphisms in A_k can only be extended in constantly many ways. So there are constantly many single-edge morphism updates. Additionally, for each such update, the target of the edge may need to be added to the definition of the morphism. Hence there are also constantly many single-node morphism updates. □

3 Simulating Turing Machines

First, we define what Turing machines we are simulating in Subsection 3.1. We then outline the basic idea of how we achieve space compression in Subsection 3.2. In Subsection 3.3, we describe exactly how Turing machine configurations are represented as graphs. Next, we give and describe the class of programs that simulate Turing machines in Subsection 3.4. We then provide an example of a Turing machine simulation in Subsection 3.5. Finally, we show the correctness of the simulation in Subsection 3.6.

3.1 Off-Line Turing Machines

In this paper, we consider deterministic off-line Turing machines in order to exhibit that space compression can happen with sublinear space complexities.

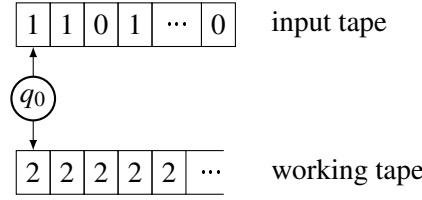


Figure 4: Initial configuration of an off-line Turing machine

The input tape is finite and uses a binary encoding, i.e. the input alphabet is $\Sigma = \{0, 1\}$. It is read-only (symbols cannot be modified) with a tape head pointing towards the tape square being read. The input tape contains the input of the Turing machine, and its length is the length of the input.

Off-line Turing machines additionally have a working tape, a one-way infinite tape that can be read and written on (symbols can be modified). The tape heads on both tapes move simultaneously, but independently. We allow for an additional symbol 2 as a blank symbol or separator, i.e. the working tape alphabet is $\Gamma = \{0, 1, 2\}$. We use 2 instead of another symbol because it is useful for the working tape to be encoded in base 3 for the purposes of our simulation. Initially, the working tape head contains only blank symbols, i.e. the symbol 2. We define *working tape contents* to be the section of the working tape starting at the beginning of the tape, and ending with the final nonblank (not a 2) tape square that is followed by only blanks (only 2's).

We consider the finite state set Q to consist of integers. There is an initial state q_0 , an accepting state h_a , and a transition function $\delta: Q \times \Sigma \times \Gamma \rightarrow Q \times \Gamma \times \{L, R, S\} \times \{L, R, S\}$. Consider the transition $\delta(q, a, b) = \langle q', b', D_1, D_2 \rangle$. This means in state q , with a on the input tape, and b on the working tape, the machine goes into state q' , replaces the working tape symbol with b' , moves the input tape head in direction D_1 , and moves the working tape head in direction D_2 . Transitions and computations are defined as usual.

The time complexity of a Turing machine is the function which, to an integer n , associates the maximum number of transitions (or length of computation) starting from an input of size n . The space complexity is the function which associates to n the maximum number of used working tape squares for an input of size n .

3.2 The Basic Idea of the Simulation

We simulate a Turing machine M with an SGP 2 program called $\text{Sim}(M)$, as specified in Figures 7 to 12. Input symbols, tape symbols, and state names are encoded as integers. The compression is achieved by representing blocks of Turing machine tape squares as single nodes. This idea is shared with van Emde Boas' representation [17]. By representing graphs in SGP 2, we can have an edge that provides a direct link to a distant list element, which is impossible in a Turing machine.

The space compression comes from how we represent the working tape, which is illustrated in Figure 5 (note that the text not contained within nodes or squares is explanatory). The working tape is divided into blocks of size $c = 4$. Each block can be interpreted as a 4-digit number in base 3. Block 1 for instance contains the ternary number 1022. In base 10, that is $1 \cdot 3^3 + 0 \cdot 3^2 + 2 \cdot 3^1 + 2 \cdot 3^0 = 27 + 6 + 2 = 35$. Hence we say the content of block 1 is 35. The graph CACHE represents the block currently containing the working tape head, which is block 1 in this case. We call this block the *active block*. CACHE is a doubly-linked list whose 4 nodes are labelled with the symbol of the corresponding tape square. The graph BLOCKSET is a less direct representation of the working tape. Each of its nodes

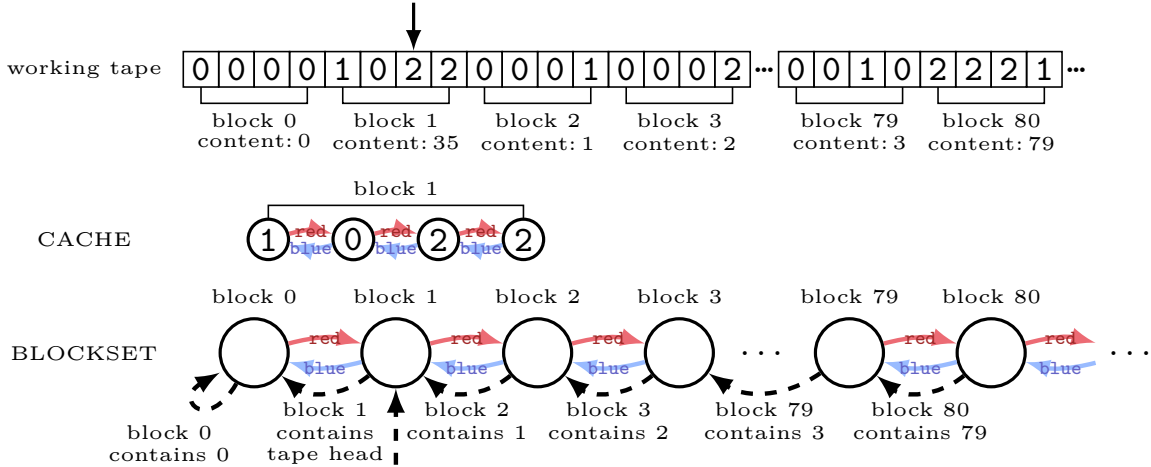


Figure 5: Graph representation of the working tape of a Turing machine.

represents a block of working tape squares instead of a single square. Instead of storing the content of a block as a spacious label, we store it as a single dashed edge. (This is the crucial idea making compression possible.) Block 79 for instance has content 3, so there is a dashed edge from node “block 79” to node “block 3”. The content of a block is converted into the index of a node in BLOCKSET (seen as a doubly-linked list). We call the nodes of BLOCKSET *block nodes*.

In general, if block i has content j , there is dashed edge in BLOCKSET from node “block i ” to node “block j ”. The only exception is the node representing the active block, in this case node “block 1”, which has a dashed edge pointing towards node “block 0”. This is by convention for the following reason. The content of “block 1” is stored in CACHE, so storing it in BLOCKSET as well would be redundant. So instead of using the dashed edge to store the content of “block 1”, the dashed edge is directed to the leftmost node (“block 0”) by convention. Since that dashed edge now holds no information, when labels in CACHE are changed, the dashed edge does not need to be updated. Once the tape head moves to a different block, and the active block changes, that dashed edge is redirected to represent to content of its block again, and the outgoing dashed edge of the new active block node is redirected to node “block 0”.

The only part of the tape we need to represent is the working tape contents, since only blanks follow. For a given working tape, let b be the minimum number of blocks that contain the working tape contents. So we only need b nodes in BLOCKSET. Remember $c = 4$ is the number of tape squares per block. Since we are converting the content of a block into a list index in BLOCKSET, the list needs to be long enough to have that index. In other words, the number of nodes b in BLOCKSET needs to be greater than the largest possible block content. A 4-digit ternary number can have $3^c = 3^4 = 81$ different values, so we need $b \geq 3^c = 81$. Since our aim is compressing space, we want to use the lowest number of nodes, so we pick $b = 3^c = 81$. If the tape contents exceed the size b blocks, we extend our representation by incrementing c , and updating $b = 3^c$. These new values become $c = 5$ and $b = 243$.

After incrementation, the number of nodes in CACHE is $c = 5$, and the number of edges $2 \cdot (c - 1) = 8$, so 13 in total. The number of nodes in BLOCKSET is $b = 243$, and the number of edges $b + 2 \cdot (b - 1) = 727$, so 970 in total. However the number of squares on the working tape is $b \cdot c = 19683$, which is significantly larger. The graph representation uses less space for $c \geq 5$. Let us describe this behaviour asymptotically. We have $c = \log_3 b$. The number of outgoing edges of a node is bounded by

3, hence the number of nodes and edges is bounded by $4(b + c)$, which is $O(b + \log_3 b) = O(b)$. The number of working tape squares on the other hand is $O(b \log_3 b)$.

3.3 Turing Machine Configurations as Graphs

Figure 4 shows the initial configuration of a Turing machine with its input and working tapes and tape heads, as well as the initial state q_0 . The corresponding graph can be seen in Figure 6.

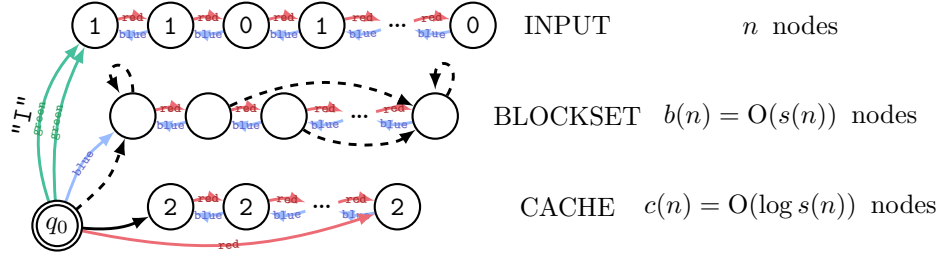


Figure 6: Graph representation of the initial configuration of a Turing machine of space complexity $O(s(n) \log s(n))$.

Let $O(s(n) \log s(n))$ be the space complexity of a Turing machine, where n is the size of the input tape, and s some function. For the rest of this paper, we omit the base of the logarithm since $O(\log_3 n) = O(\log_c n)$ for any constant $c > 1$. If we choose the number of blocks to be $b = O(s(n))$, we can represent $O(s(n) \log s(n))$ tape squares in space $O(s(n))$, as outlined by the previous subsection. This will be proven in Theorem 4.

The graph representation consists of the *central node*, labelled by the initial state q_0 as represented by some integer, and the subgraphs INPUT, BLOCKSET, and CACHE. We represent ‘left’ and ‘right’ positioning by blue and red edges respectively. These edges will not be modified by the program. The unmarked and unlabelled green edges mark the positions of the tape heads. Dashed edges serve to encode the state of BLOCKSET and how it relates to CACHE, as described in the previous subsection. The dashed edge from the central node points towards the block node whose block currently contains the working tape head (the active block node). By convention, there is a dashed edge from the active block node to the first block node, as stated in the previous subsection. Hence there is always a path of dashed edges from the central node to the active block node to the first block node.

The central node is a root that is never moved. Parts of the graph such as tape head positions and tape ends can be accessed efficiently via the central node’s outgoing edges. The only other time roots are used in the simulation is when extending the tape. We use roots to keep track of the end of the tape, and later to traverse the tape. For a more in-depth explanation of how the roots move, see Subsection 3.4.

INPUT represents the input tape of the Turing machine, and the position of the tape head is represented by an unmarked green edge from the central node. The green edge labelled “I” always points to the leftmost square of the input tape. The working tape is represented by BLOCKSET and CACHE. These subgraphs are implemented as doubly linked lists in order to enable fast rule matching.

Each node in BLOCKSET represents a block of c nodes of the working tape. The blue edge from the central node points towards the leftmost node, and the dashed edge towards the block that contains the position of the tape head.

CACHE is the block of the working tape that contains the tape head. Its position within BLOCKSET is marked with a dashed edge from the central node. The red edge from the central node points towards

the rightmost node of the cache, and the unmarked edge towards the current position of the tape head. The node labels represent the content of each square.

The operations of the Turing machine happen within CACHE. If the tape head moves out of bounds of CACHE, we encode the content of CACHE within BLOCKSET, move on to another block, and decode it into CACHE.

Initially, the tape heads are on the first square of their tapes, and the working tape is blank (i.e. filled with 2's). Remember CACHE has c nodes and BLOCKSET 3^c . We start simulation with $c = 2$ because it is the smallest nontrivial value for c . If $c = 1$, then each block would represent a single tape square, and no space would be gained. If we run out of nodes in BLOCKSET, c is incremented and BLOCKSET adjusted accordingly.

3.4 The Simulation $\text{Sim}(M)$

In this subsection, we assume we have a Turing machine M and describe $\text{Sim}(M)$, the SGP 2 program that simulates M . The program $\text{Sim}(M)$ takes inputs that consist of graphs such as INPUT together with the central node and green edges as represented in Figure 6. The overall behaviour of the program is that it starts the simulation with a small BLOCKSET and CACHE. If we run out of tape squares, we reset to the initial configuration, extend BLOCKSET and CACHE, and restart the simulation.

Figure 7 contains the main control sequence of $\text{Sim}(M)$. First the rule `setup` is called. It matches the root and constructs CACHE with two nodes and BLOCKSET with nine nodes, as seen in Figure 6. We omit the definition of that rule since it is straightforward.

Next we have the loop `Simulate!`, which means the procedure `Simulate` is applied until no longer possible. The loop body first calls `Transitions` (Figure 8), a procedure that applies the transition function to the current state of M . If `Transitions` cannot apply the transition function, the procedure results in failure, and the loop terminates. Next, try `MoveLeft` is called, which means we attempt to apply `MoveLeft` and if it fails, we skip this instruction. The rules in `MoveLeft` detect a label "L" on the unmarked edge adjacent to the central node, which is created in `Transitions` if the working tape head needs to move to the left. It then executes that move in CACHE. Analogously, try `MoveRight` attempts to move the working tape head to the right. These rule sets (which are nondeterministic calls of the rules they contain) will fail however if the tape head moves out of bounds of CACHE, which is detected by `Left` or `Right` since the continued presence of the "L" or "R" label indicates that tape head movement still needs to be done. In that case, `PrevBlock` or `NextBlock` (Figure 9) are called in order to move towards the relevant block.

Figure 8 contains `Transitions`, the rule set that encodes the transition functions of M . For each entry in the transition table, there is a corresponding rule in `Transitions`. The rules are divided into nine categories `transitionXY`, where X represents the movement of the input tape head, and Y that of the working tape head. Each rule implements the change of labels and the movement of the input tape head directly. For the movement of the working tape head, a label is left behind on the unmarked edge, and the actual movement is handled by `Simulate` in Figure 7.

In Figure 9, we show the procedures `NextBlock` and `PrevBlock`, which handle movement to another block. We call the node pointed at by a dashed edge from the central node the *active block node*. The procedure `Encode` (Figure 10) saves the information from CACHE into BLOCKSET. This process is further described in the example in Subsection 3.5. The procedure `CacheInc` increments CACHE as a ternary number. First the rightmost digit is turned into a red root. Then the digit is attempted to be incremented with the rule set `Inc`. If the digit is a 0 or a 1, incrementation succeeds, the process is marked as finished by turning the red root into an unmarked node, and `CacheInc` terminates. If the digit

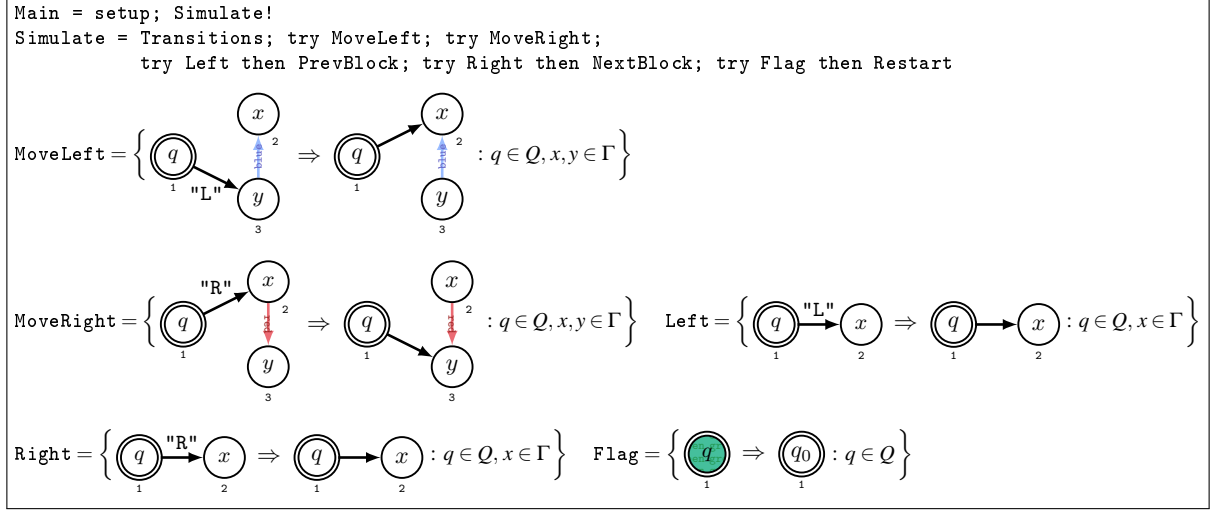
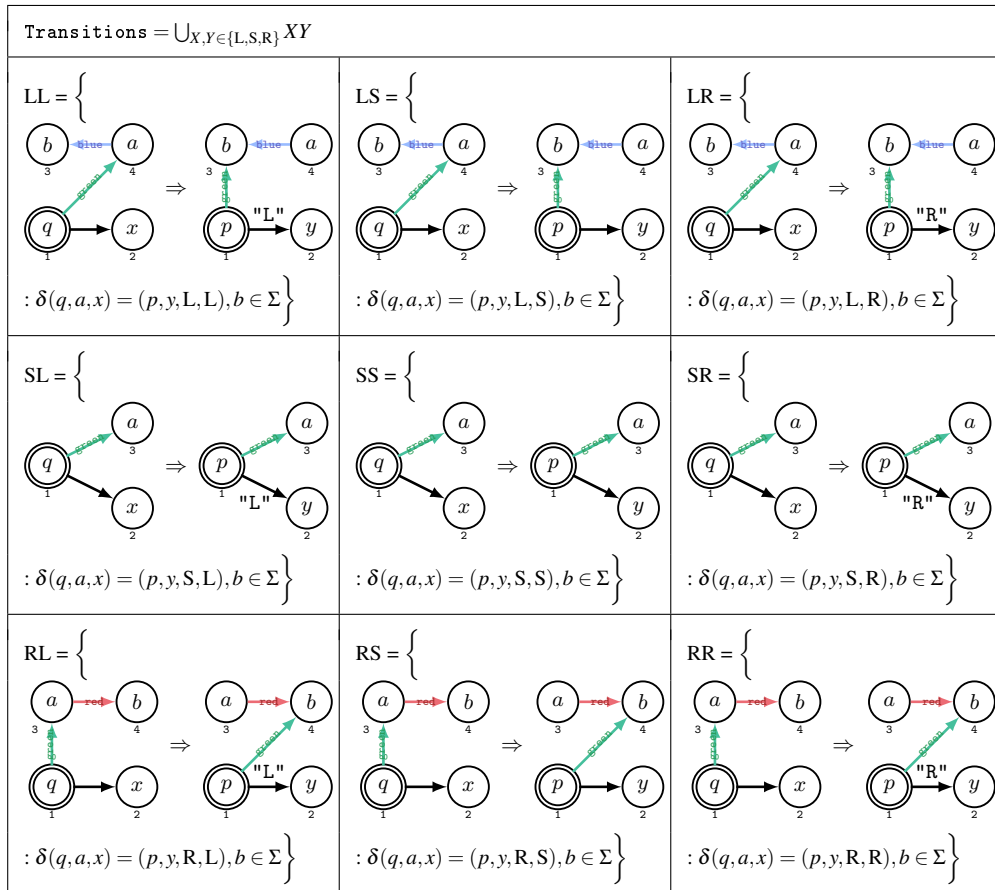
Figure 7: SGP 2 program $\text{Sim}(M)$ that simulates Turing machine M .

Figure 8: Rule set Transitions that models the transitions of a Turing Machine.

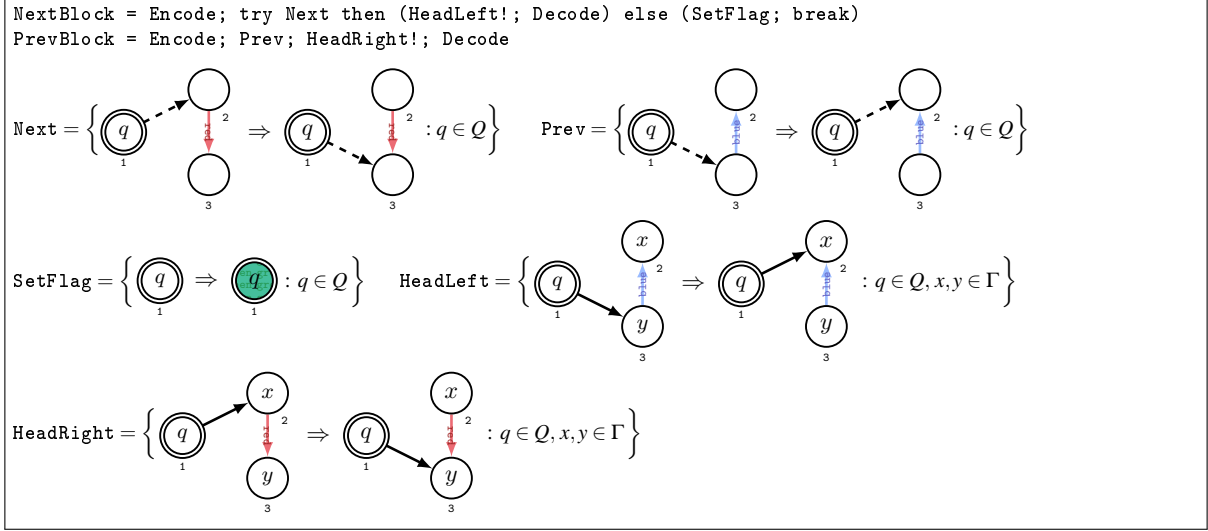


Figure 9: Procedures NextBlock and PrevBlock that change the active block.

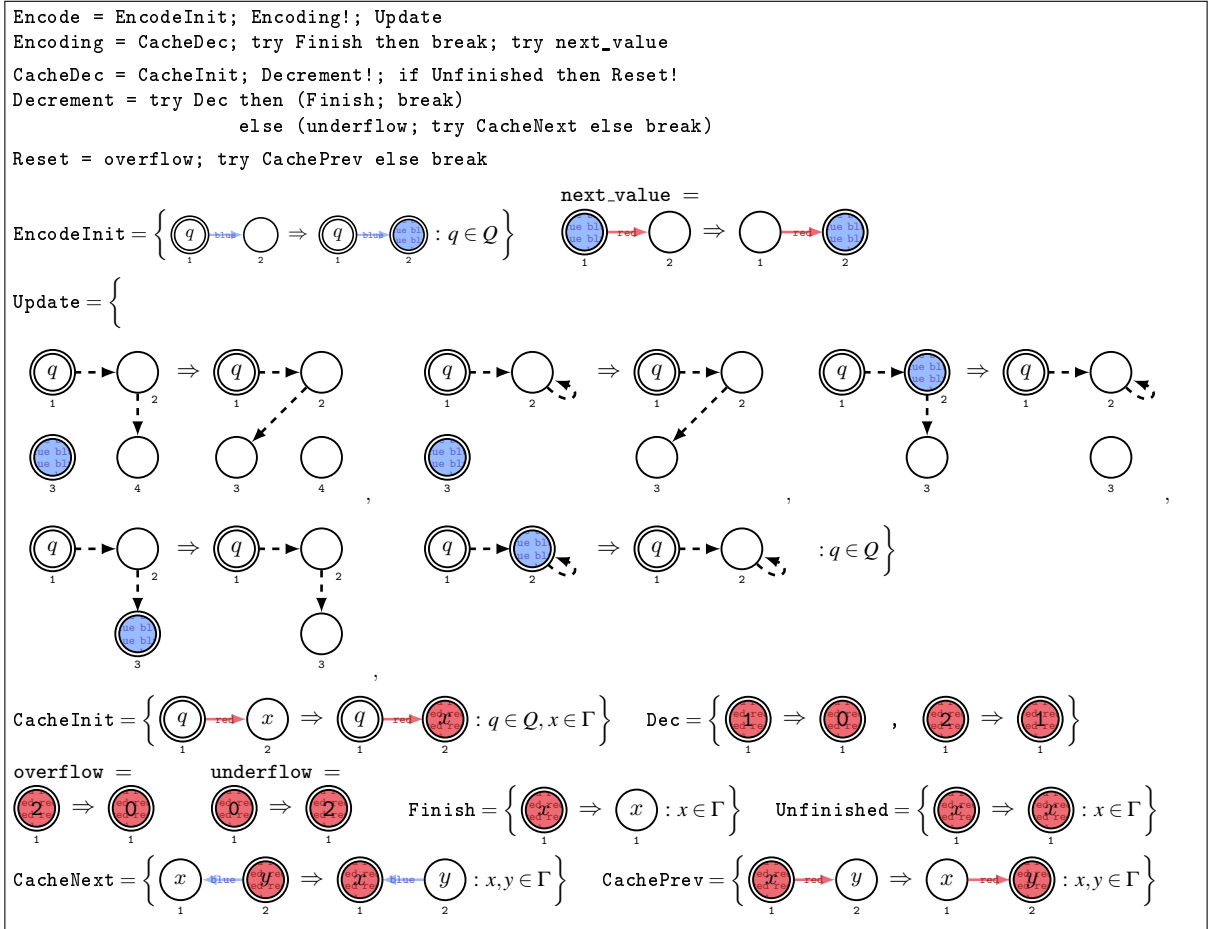


Figure 10: Procedure Encode that encodes the current block into BLOCKSET.

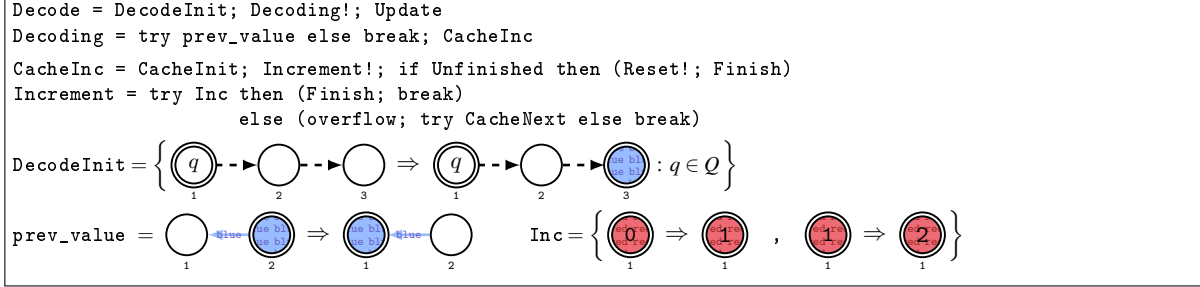


Figure 11: Procedure Decode that decodes the current block into CACHE.

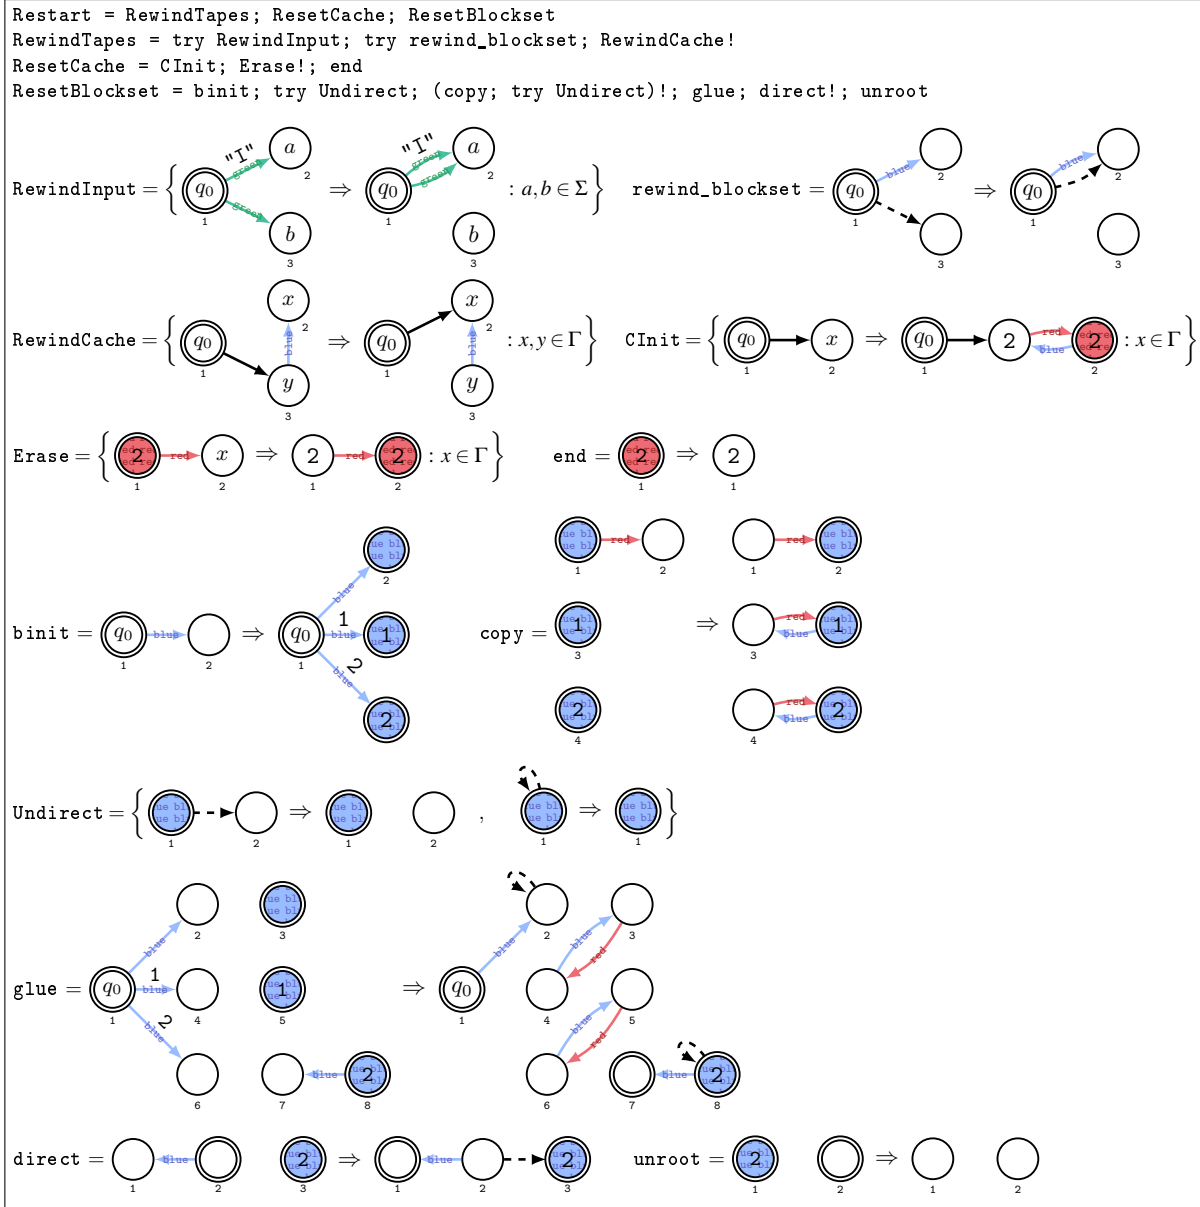


Figure 12: Procedure Restart that resets the simulation and enlarges the tape.

is a 2 however, Inc fails, the 2 is turned into a 0, and the red root is moved along to the next digit which is attempted to be incremented. If Increment! ends without Inc ever succeeding, it means that all digits were 2, and that the red root is still present. In that case, Reset! turns all digits into 0.

The procedure Decode in Figure 11 decodes the active block in BLOCKSET and stores the information in CACHE. It is analogous to Encode (note that some of Decode's rules and procedures appear in Figure 10). Similarly to CacheInc, CacheDec decrements CACHE as a ternary number if that number is not 0, and keeps it as 0 if it already is.

Figure 12 shows the procedure Restart which, when the machine runs out of tape, resets the simulation and extends the tape. CACHE and hence the size of a block is extended by one square, and the size of BLOCKSET is tripled. The procedure RewindTapes rewinds the tape heads to the beginning. Next, ResetCache resets the content of CACHE back to 2's (the blank symbol). Then, ResetBlockset extends BLOCKSET and directs dashed edges to their initial state. In rules binit and copy, an unlabelled blue root traverses the existing BLOCKSET, while blue roots with labels 1 and 2 traverse two copies of that BLOCKSET as they are being created. In the meantime, Undirect deletes dashed edges that are outgoing from block nodes. The rule glue appends the copies to the original BLOCKSET, tripling its length. The blue nodes are unrooted except for the final node, which remains a blue root labelled 2. The penultimate node becomes an unmarked root, which is used in direct to traverse BLOCKSET from right to left, creating dashed edges where needed. The rule unroot then removes the blue and unmarked roots.

3.5 Example

In this subsection, we give an example of a Turing machine simulation, and show how we move from one block to another. Consider a Turing machine that takes as input the number n represented in unary, and writes n in binary on its working tape n times. It is reasonable to assume the machine has a space complexity of $O(n \log n)$. If $n = 6$, the machine uses 18 squares, which are filled with 6 copies of the string 110 (6 in binary). A CACHE size of 2 and a BLOCKSET size of 9 are enough to represent $2 \cdot 9 = 18$ tape squares. Their representation in the simulation has only $2 + 9 = 11$ nodes, which is $O(n + \log n) = O(n)$. The initial state of the machine on input 6 is shown in Figure 13, and the final one in Figure 14.

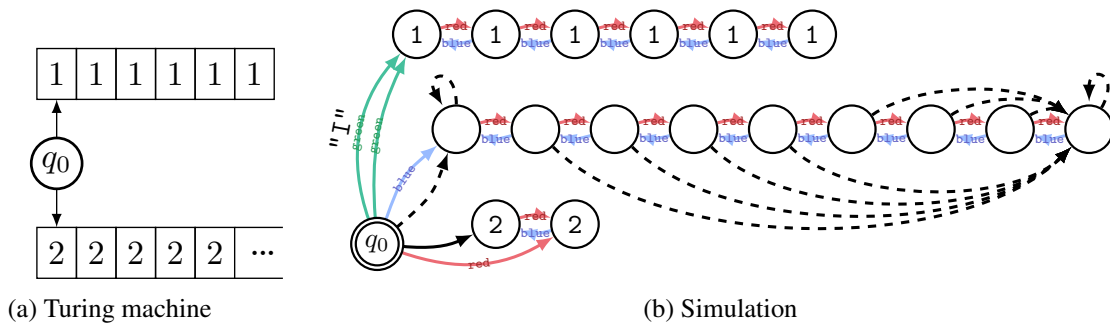


Figure 13: Initial state of the example machine

Initially, the working tape is blank (all 2's). Hence most blocks consist of two 2's, i.e. their content is 8 (22 in ternary is 8 in base 10). So most block nodes point towards the block node of index 8 (final block node). The exception is the first node, which points to itself. This is because it represents the block that contains the tape head, which is set to index 0 by convention. In the final state, the working tape is

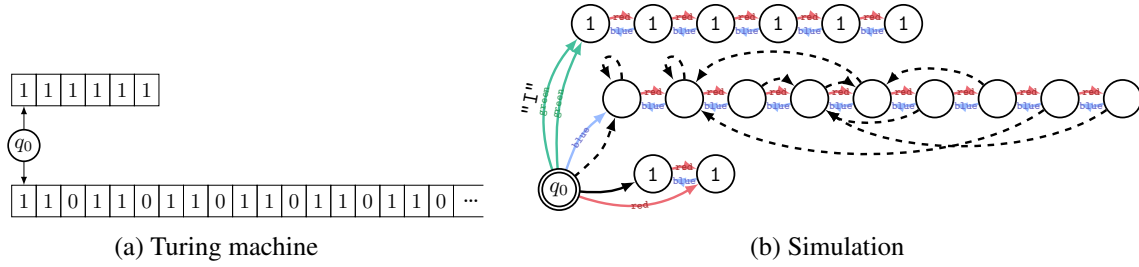


Figure 14: Final state of the example machine

represented in the same manner. Note that even though the working tape content looks binary, 2's (blank symbols) could be present, and hence tape blocks are considered a ternary number. For instance consider the third block (of size 2). It contains 10, so the third block node points towards the block node of index 3 (fourth block node).

Let us sketch the behaviour of the machine. First, the input is copied in unary onto the working tape for use as a counter. Then, a binary number to the right of the counter is incremented while traversing the input. The previous step is repeated while decrementing the counter until it reaches 0. Tape contents need to be shifted. The symbol 2 can be used as a separator.

Now let us show what happens when the block has to be changed. Consider the situation where 0100 are the first 4 squares of the working tape and the tape head is on the third square. Let the next transition write symbol 1 and move the tape head to the left, and change state q_i to state q_j .

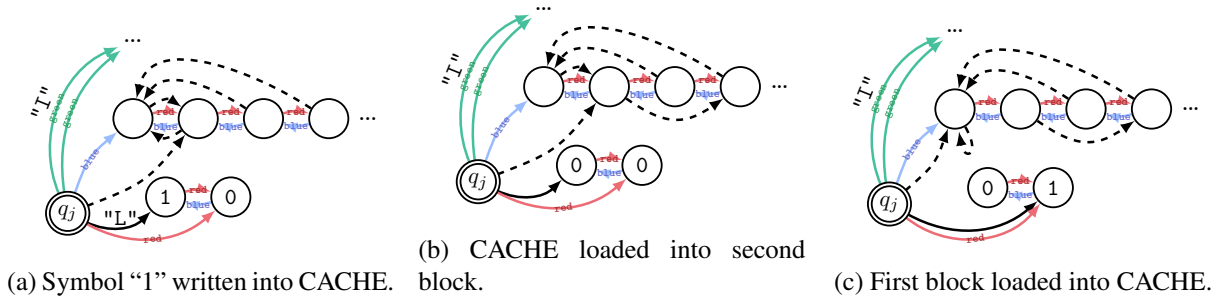


Figure 15: Changing blocks

Figure 15 shows the process of changing the active block. In Subfigure 15a, a rule from Transitions has just been applied, labelling the left node of CACHE 1, and the unmarked edge “L”. The working tape head is yet to be moved. Next, in Simulate, Left matches and PrevBlock is called. The procedure Encode loads the content of CACHE into a dashed edge within BLOCKSET. While CACHE is decremented as a ternary number, the outgoing dashed edge of the current (second) BLOCKSET node is shifted to the right. This produces the graph from Subfigure 15b. CACHE does not represent the content of the second block anymore, but the dashed edge from the second node in BLOCKSET to the fourth does. Next we move one block to the left with Prev, and reposition the tape head in CACHE with HeadRight!. We then load the new block into CACHE. While moving the outgoing dashed edge of the current (first) BLOCKSET node to the left, the ternary number in CACHE is incremented. This results in the graph from Subfigure 15c. We have now loaded the content of the first block into CACHE, so there is no need to store it with a dashed edge anymore. Hence the dashed edge goes from the first node

in BLOCKSET to itself by convention.

3.6 Correctness

We define the *configuration* of a Turing machine to consist of the input and working tapes, the position of the tape heads on those tapes, and the current state.

Let us define what graphs programs in *Sim* operate on. We call graphs as in Figure 6 *configuration graphs*. These graphs can vary from the depicted graph in the following ways. The targets of the unlabelled green edge, dashed edges, and the unmarked edge can be any node in INPUT, BLOCKSET, and CACHE respectively. An exception to that is the outgoing dashed edge of the active node in BLOCKSET (target of dashed edge from the root), which is targeted towards the leftmost node in BLOCKSET since the content of that block is currently stored in CACHE. This is a convention that allows the graph representation of a configuration to be unique. The labels of nodes in INPUT, CACHE, and the root node can be any element from the input alphabet, tape alphabet, or set of states respectively, encoded as SGP 2 labels.

Note that the graphs' capacity to represent tape squares is limited by the sizes of BLOCKSET and CACHE. Hence for a non-negative integer k , we define a function enc_k that encodes the configuration S of a Turing machine as the configuration graph $\text{enc}_k(S)$ with 3^{k+2} nodes in BLOCKSET and $k+2$ in CACHE. The integer k corresponds to the number of times the tape is extended by the procedure *Restart*. For a given value of k that is large enough, $\text{enc}_k(S)$ exists and is unique.

For a Turing machine M we denote a single transition from configuration S to S' by $S \Rightarrow_M S'$, and the transitive reflexive closure of that relation using \Rightarrow_M^* . Similarly, for an SGP 2 rule r we use $S \Rightarrow_r S'$ to denote a change of configuration, and $S \Rightarrow_P^* S'$ for a program P .

During the execution of *Sim*(M), graphs that are not configuration graphs are generated. They only differ in terms of edge labels, node marks, roots, and dashed edge targets. However, these variations are temporary. Once the procedure *Simulate* has successfully terminated, the host graph is once again a configuration graph.

We say an SGP 2 program $P!$ (i.e. a loop where procedure P is applied as long as possible) *simulates* a Turing machine M of initial configuration I if, for each two configurations S and S' of M such that $I \Rightarrow_M^* S \Rightarrow_M S'$, we have $\text{enc}_0(I) \Rightarrow_{P!}^* \text{enc}_k(S) \Rightarrow_P^* \text{enc}_k(S')$ for some integer k .

The following theorem shows that this is indeed the case for the class of programs presented in this paper.

Theorem 2 (Correctness). *Let M be a Turing machine and $\text{Sim}(M)$ the corresponding SGP 2 program. Then the subprocedure *Simulate!* simulates M .*

Proof. Let S and S' be configurations of M such that $I \Rightarrow_M^* S \Rightarrow_M S'$. Then some transition $\delta(q, a, x) = (p, y, X, Y)$ happened, where $p, q \in Q$, $a \in \Sigma$, $x, y \in \Gamma$, and $X, Y \in \{L, S, R\}$. So the only changes from S to S' are that q and x have been updated to p and y , and that the input and working tape heads have been shifted into direction X and Y respectively.

We proceed by induction on the length of $I \Rightarrow_M^* S$. For the base case we only need to show $\text{enc}_k(S) \Rightarrow_P^* \text{enc}_k(S')$ for some integer k . This also happens to be what we need to show for the induction step since $\text{enc}_0(I) \Rightarrow_{P!}^* \text{enc}_k(S)$ is given by the induction hypothesis, and the existence of a large enough k by part of the proof of Theorem 4.

Now it remains to show that $\text{enc}_k(S) \Rightarrow_P^* \text{enc}_k(S')$. The procedure *Reset* can be ignored because it is never called since we already have a large enough k .

When running `Simulate` on the graph $\text{enc}_k(S)$, the first call is the rule set `Transitions`. A rule in that set that is guaranteed to be applicable is the one corresponding to the aforementioned Turing machine transition. It correctly updates the q and x to be p and y to match $\text{enc}_k(S')$. It also moves the input tape head to the correct position. For the rest of this proof, we will show that the working tape head moves to the correct position.

If $Y = S$, the working tape head is not moved, corresponding to $\text{enc}_k(S)'$. Due to lack of an unmarked edge label, none of the `try` conditions in `Simulate` are applicable, and we terminate with $\text{enc}_k(S')$.

The cases for $Y = L$ and $Y = R$ are analogous, so we will only argue for the former. When attempting to apply `MoveLeft`, if the target of the unmarked edge is not the leftmost node of `CACHE`, the rule set succeeds, the unmarked edge is in the correct position, and no other conditions of `try` statements in `Simulate` are applicable. Hence we terminate with $\text{enc}_k(S')$.

If target of the unmarked edge is the leftmost node of `CACHE` however, the only `try` condition that can match is `Left`. This removes the edge label and calls `PrevBlock`. So it remains to argue that `PrevBlock` turns the host graph into $\text{enc}_k(S')$.

The commands `Prev`; `HeadRight!` correctly position the working tape head on the rightmost node of the previous block. It remains to show that `Encode` and `Decode` preserve the working tape content according to the encoding enc_k .

The procedure `Encode` contains a loop that, whenever it decrements the content of `CACHE`, it increments the content of the active block (represented by an outgoing dashed edge). This happens until the `CACHE` content reaches 0, meaning the content of `CACHE` is correctly store in `BLOCKSET`. After moving to the new block, the procedure `Decode` is called. It contains a loop that, whenever it decrements the content of the active block, increments the content of `CACHE`. This terminates once the content of the active block has reached 0, meaning the content of the active block has been correctly stored in `CACHE`. Hence the host graph is now $\text{enc}_k(S')$. \square

4 From the Simulation to Computation Models

Based on the `Sim` programs, we can generalise to computation models that are subsets of GP2, for which we can define a time measure. One obstacle is time-intensive rule matching. A second one is that when *critical subprograms* (bodies of loops and conditions of `if` and `try`) fail, they need to be undone, which can be handled in different ways. The current implementation of GP2 reverses the relevant rule applications. The formal semantics uses a stack of graphs to keep track of the state of the host graph before entering a critical subprogram. Reversion can then be done by a simple `pop` operation. The downside to this is that the host graph has to be duplicated every time a loop or a branching statement is entered, which can add a polynomial factor to the time complexity of the overall program. It is also space-intensive, but only by a polynomial factor since only finitely many loop and branching statements can be nested in a command sequence. To avoid undoing, one wants failed critical subprograms to be *null*, i.e. they do not change the graph state, meaning they do not need to be undone.

Models that overcome these obstacles can be implemented to match rules in constant time, which allows us to assign unit time cost to that operation. They also make undoing redundant, allowing us to charge no time cost for that.

Another obstacle is that matching a set of rules is nondeterministic, and backtracking that nondeterminism can be costly in time or space. This can be addressed at the implementation level. The current GP2 compiler for instance picks the first match it finds and does not backtrack. A more general approach however is to show that within the scope of running a program on an input graph, every rule and rule set

(nondeterministic call of a list of rules) can only have at most one match, effectively making the program deterministic. This property applies to the class of programs in this paper and is shown in Proposition 2.

We define an *efficient model* as $\mathcal{M} = \langle \mathcal{P}, \mathcal{I} \rangle$, where \mathcal{P} is a set of GP 2 programs and \mathcal{I} a set of GP 2 graphs such that the following two properties are satisfied within the scope of derivation sequences starting with any $\langle P, I \rangle \in \mathcal{M}$.

- (1) *Constant Matching*: Every rule matches in constant time.
- (2) *Critical Subprogram*: Every critical subprogram that fails is null.

Let us define the *graph space* measure of a graph as the number of nodes plus the number of edges. We do not consider the size of labels in this paper since all labels are of constant size. This measure is uniform in that it gives unit cost to each node and edge. A discussion about this uniformity can be found at the end of Section 5.

For an efficient model \mathcal{M} , we define the time complexity of a program P of \mathcal{M} as the maximum number of rule calls in terminating derivation sequences starting with P on graphs of a given size.

We show in this section that $\langle P, I \rangle$, where $P = \{\text{Sim}(M) \mid M \text{ is a Turing machine}\}$ and $I = \{\text{CG} \mid \text{CG is a configuration graph}\}$, is an efficient model. The constant matching property is given by Theorem 1, which holds because all rules are fast, input graphs have bounded outdegree and a bounded number of roots, and the programs preserve this boundedness. Moreover, when extending the tape with `Restart`, outgoing dashed edges of existing block nodes are removed, and one outgoing dashed edge is added for each block node, preserving bounded outdegree. The critical subprogram property is shown by Proposition 1.

The remainder of this section consists of aforementioned propositions and a lemma needed to prove one of them.

Lemma 1. *Let P be a critical subprogram in a graph state S . Assume either P cannot fail from state S , or P can only fail from state S due its first component being a rule or rule set call that fails to match. Then if P fails from state S , P is null.*

Proof. If P cannot fail from state S , the lemma is trivially satisfied. Now assume P fails from S due to its first component being a rule or rule set failing to match. Since it failed to match, the first component cannot have changed the host graph, and since it is the first, no other component can have changed the host graph. \square

Proposition 1 (Critical Subprogram Property). *In $\text{Sim}(M)$, given configuration graphs as inputs, every critical subprogram that fails is null.*

Proof. We will argue for each critical subprogram of `Sim` that Lemma 1 applies.

The conditions of all `if` and `try` statements are either rules or rule set calls, satisfying Lemma 1.

Let us now argue for the loop bodies. The procedures `Erase`, `RewindCache`, `HeadLeft` as well as `HeadRight` are rule sets, and `rewind_blockset` and `direct` are rules, satisfying Lemma 1.

Reset: The rule `overflow` can fail, but the other component cannot since it is a `try` statement whose branches only contain `break` which cannot fail.

Increment: This consists of a `try` statement, so only the branches can fail. `Finish` cannot fail because `CacheInit` is always called before this loop, providing a match. The rule `overflow` cannot fail because it is only called when `Inc` fails and the labels of nodes in `CACHE` are either 0, 1, or 2. The `try` statement cannot fail because its branches cannot fail.

Decrement: The reasoning is analogous to that in the previous paragraph.

Decoding: The `try` statement cannot fail because its branches cannot fail. For the remainder of this paragraph, we argue that `CacheInc` cannot fail either. `Increment!` cannot fail because it is a loop. `CacheInit` always succeeds because the target of the red edge originating from the unmarked root is unmarked because of the structure of the input graph, and because previous calls of `CacheInc` and `CacheDec` turn the only marked rooted node in this part of the graph unmarked and unrooted with `Finish`.

Encoding: The reasoning is analogous to that in the previous paragraph.

The body of the loop `(copy; try Undirect)!` contained in `ResetBlockset`: The rule `copy` is allowed to fail, and `try Undirect` cannot fail since it is a `try` statement.

Simulate: `Transition` is a rule set at the start of the body and hence allowed to fail. The rest consists of `try` statements whose branches call `PrevBlock` and `NextBlock`. `Prev` and `Next` cannot fail because we assume that `BLOCKSET` in the input graph is large enough to accommodate the execution of the Turing machine. `Encode` and `Decode` do not fail, which we will argue for the rest of this paragraph. `EncodeInit` and `DecodeInit` find a match because of the structure of the input graph and because previous calls of `Encode` and `Decode` leave `BLOCKSET` unmarked and unrooted. `Encoding!` and `Decoding!` are loops and hence cannot fail. `Update` has a match regardless of the blue root's location within `BLOCKSET`. \square

Proposition 2 (Unique Matches). *In $\text{Sim}(M)$, given configuration graphs as inputs, whenever a rule or rule set is called, there is at most one match.*

Proof. In each rule set, if the left-hand side of two rules have the same structure, their labels are different. In `Transitions` in particular, no two rules have the same left-hand side since we consider deterministic Turing machines. Among rules with different labels but the same structure, at most one can match, namely the one that has the same labels as the host graph. Among rules with different structures and the right labels, at most one can match because they differ by an edge or node that is unique in host graphs. Furthermore, it is easy to check that each rule can only have at most one match in the host graph. Each rule contains a root with a unique combination of mark and label. And for each node, outgoing edges can be distinguished in the same way. \square

5 Results on Time and Space Complexity

In this section, we present theorems on the time and space complexities of the simulation.

Theorem 3 (Time Complexity). *Every Turing machine M of time complexity $t(n)$ is simulated by $\text{Sim}(M)$ in time $O(t^2(n))$, where n is the size of the input.*

Proof. Given the discussion in Section 4, we can assign unit time to rule and rule set calls and argue time complexity to be the number of such calls.

First, we will show that simulating one step of M (not counting restarting the simulation) takes $\text{sim}(n) = O(s(n))$ time. We consider the size of `CACHE` from the final simulation since it provides an upper bound. The only sources of non-constant time are `PrevBlock` and `NextBlock`. Their complexity is the worst of the loops `Encoding!`, `Decoding!`, `HeadLeft!`, and `HeadRight!`. The latter two simply traverse `CACHE`, which takes $O(s(n))$ time. The former two decrement/increment `CACHE` as a ternary counter. Their time complexity is proportional to the number of digit operations it takes to decrement the counter from the number $s(n)$ all the way to 0. The rightmost digit is modified $s(n)$ times, the next

one $\frac{1}{3}s(n)$ times, the one after $\frac{1}{3} \cdot \frac{1}{3}s(n)$ times, and so on. So the total number of digit operations is $\sum_{k=0}^{\log s(n)} s(n) \left(\frac{1}{3}\right)^k$. Using properties of the geometric series, one can see that this is $O(s(n))$.

In this paragraph, we will show that resetting the simulation and extending the tape takes $r(n) = O(s(n))$ time. Consider the loops of `Restart`. Both `RewindCache!` and `Erase!` traverse `CACHE` and hence take time $O(\log s(n))$. The loop `(copy; try Undirect)!` in `ResetBlockset` traverses `BLOCKSET` and thus takes $O(s(n))$ time. And `direct!` traverses the extended `BLOCKSET` in $O(3 \cdot s(n)) = O(s(n))$ time.

Next, we will show that the number of times the simulation is restarted is $l(n) = O(\log s(n))$. The final size of the tape of M is $O(s(n) \log s(n))$. Since the number of represented squares is tripled in each step, the number of steps is $O(\log(s(n) \log s(n)))$. Using the formula for the logarithm of a product, one can simplify this to $O(\log s(n))$.

The total time complexity can be bounded by $\text{reset}(n) + \text{simulation}(n)$, where $\text{reset}(n) = l(n) \cdot r(n)$ is the total cost of all resets, and $\text{simulation}(n) = l(n) \cdot t(n) \cdot \text{sim}(n)$ the total cost of simulating M across all resets. Using results from previous paragraphs, we get $\text{reset}(n) = O(s(n) \log s(n))$ and $\text{simulation}(n) = O(t(n) \cdot s(n) \log s(n))$. Since space complexity $s(n) \log s(n)$ can be bounded by time complexity $t(n)$, the entire simulation takes $O(t^2(n))$ time. □

Theorem 4 (Space Complexity). *Every Turing machine M of space complexity $O(s(n) \log s(n))$ is simulated by $\text{Sim}(M)$ in graph space $O(s(n))$, where n is the size of the input.*

Proof. During the execution of $\text{Sim}(M)$, nodes and edges are only created by `setup` and `Restart`, and none are ever deleted. The numbers of nodes and edges only differ by a constant factor since the outdegree is bounded, so we will argue for space complexity using number of nodes only.

Initially, after application of `setup`, `BLOCKSET` has $b(n) = 3^2$ nodes, and `CACHE` $c(n) = 2$. Then, each application of `Restart` adds one to $c(n)$ and triples $b(n)$. So after k iterations, we have $b(n) = 3^{2+k}$ and $c(n) = 2 + k$.

The Turing machine needs $S(n) = O(s(n) \log s(n))$ tape squares. This means that there are positive integers n_0 and c such that $S(n) \leq cs(n) \log s(n)$ for all $n \geq n_0$. So for all n we can say $S(n) \leq cs(n) \log s(n) + m$, where $m = \max_{n \in \{0, \dots, n_0\}} S(n)$, a constant.

Assume `Restart` is called $k = \log s(n) - 2 + d$ times, where $d = \max(m, \log c)$. For that value of k , we have $c(n)b(n) = 3^d s(n) \log s(n) + 3^d s(n) d$. By definition we have $d \geq m$ and $3^d \geq c$. Furthermore, we have $3^d s(n) \geq 1$. Hence we get $c(n) \cdot b(n) \geq cs(n) \log s(n) + m \geq S(n)$. So for this value of k , the graph can store enough tape squares to execute M .

With the aforementioned value of k , the number of nodes in this graph is $c(n) + b(n) = 3^d s(n) + \log s(n) + d$, which is $O(\log s(n) + s(n)) = O(s(n))$.

Hence the number of nodes of the graph that is created is bounded by $c(n) + b(n)$, and hence in $O(s(n))$. □

One might wonder why this space compression is not possible on random access machines. GP2 does have a C implementation after all. The reason for this is related to how graphs are represented in random access machines (RAMs), and how much space that takes. Graph edges are usually implemented as pointers. However the size of pointer addresses grows logarithmically with the number of nodes, since these addresses are usually stored as binary numbers. So in the context of RAMs, it does not seem very accurate to assign unit cost to edges. To take this into account, one can use *logarithmic space*, in which graphs of $s(n)$ nodes are assigned a cost of $s(n) \log s(n)$.

If we use logarithmic space on our model, space compression is nullified since the simulation then has the same asymptotic space complexity as the machine it simulates. This puts into question whether uniform or logarithmic space should be used, which is discussed by van Emde Boas [17]. One may want to charge more than unit space since in RAMs, edges are represented by pointers whose size grows with the number of nodes. A related issue can be found in the time measure of RAM models when programs have to deal with large integers.

6 Conclusion

In GP 2, we have found the same space compression phenomenon that SMMs and KUMs exhibit. Unlike the other computation models, GP 2 uses rule-based graph transformation. The compression happens due to use of a graph-based data structure to encode the tape of a Turing machine, which is more space-efficient with respect to uniform measures.

Using this simulation, we have identified efficient computation models that are subsets of GP 2. Requiring these models to have certain properties allows us to define rigorous complexity measures.

In future work, the simulation could be refined in the way Luginbuhl [12] refines the one by van Emde Boas [17], namely with a more complex data structure that allows for real-time simulation with some form of space compression.

Furthermore, the simulation could be extended to nondeterministic Turing machines. Extending the simulation would be straightforward, but the complexity results would need further proofs.

One could also investigate what measurable impact this space compression has on a practical example. Graph programs working on linked lists are the most likely to benefit from this paper's space compression. As a case study, one could pick a list sorting algorithm, implement it on integer lists both with and without space compression, and compare their memory usage.

References

- [1] *GP 2 Rule Matching Code*. <https://github.com/UoYCS-plasma/GP2/blob/master/src/genRule.c>. Accessed: 2022-08-08.
- [2] Alfred Aho, John Hopcroft & Jeffrey Ullman (1974): *The Design and Analysis of Computer Algorithms*. Addison-Wesley.
- [3] Christopher Bak (2015): *GP 2: Efficient Implementation of a Graph Programming Language*. Ph.D. thesis, Department of Computer Science, University of York, UK. Available at <https://etheses.whiterose.ac.uk/12586/>.
- [4] Christopher Bak & Detlef Plump (2012): *Rooted Graph Programs*. In: *Proc. 7th International Workshop on Graph Based Tools (GraBaTs 2012)*, *Electronic Communications of the EASST* 54, doi:10.14279/tuj.eceasst.54.780.
- [5] Christopher Bak & Detlef Plump (2016): *Compiling Graph Programs to C*. In: *Proc. International Conference on Graph Transformation (ICGT 2016)*, *Lecture Notes in Computer Science* 9761, Springer, pp. 102–117, doi:10.1007/978-3-319-40530-8_7.
- [6] Graham Campbell, Brian Courtehoue & Detlef Plump (2022): *Fast Rule-Based Graph Programs*. *Science of Computer Programming* 214, doi:10.1016/j.scico.2021.102727. 32 pages.
- [7] Brian Courtehoue & Detlef Plump (2021): *A Small-Step Operational Semantics for GP 2*. *Graph Computation Models (GCM 2021)* 214, pp. 89–110, doi:10.4204/EPTCS.350.6.

- [8] Edsger W Dijkstra (1968): *Letters to the editor: go to statement considered harmful*. *Communications of the ACM* 11(3), pp. 147–148, doi:10.1145/362929.362947.
- [9] Joseph Y. Halpern, Michael C. Loui, Albert R. Meyer & Daniel Weise (1986): *On time versus space III*. *Mathematical systems theory* 19, pp. 13–28, doi:10.1007/BF01704903.
- [10] Ivaylo Hristakiev & Detlef Plump (2016): *Attributed Graph Transformation via Rule Schemata: Church-Rosser Theorem*. In: *Software Technologies: Applications and Foundations – STAF 2016 Collocated Workshops, Revised Selected Papers, Lecture Notes in Computer Science* 9946, Springer, pp. 145–160, doi:10.1007/978-3-319-50230-4_11.
- [11] A. N. Kolmogorov & V. A. Uspenskii (217–245): *On the definition of an algorithm*. *American Mathematical Society Translations: Series 2* 29, pp. 217–245, doi:10.1090/trans2/029.
- [12] David R. Luginbuhl & Michael C. Loui (1993): *Hierarchies and Space Measures for Pointer Machines*. *Information and Computation* 104(2), pp. 253–270, doi:10.1006/inco.1993.1032.
- [13] Detlef Plump (2012): *The Design of GP 2*. In: *Proc. 10th International Workshop on Reduction Strategies in Rewriting and Programming (WRS 2011), Electronic Proceedings in Theoretical Computer Science* 82, pp. 1–16, doi:10.4204/EPTCS.82.1.
- [14] Detlef Plump (2017): *From Imperative to Rule-based Graph Programs*. *Journal of Logical and Algebraic Methods in Programming* 88, pp. 154–173, doi:10.1016/j.jlamp.2016.12.001.
- [15] A. Schönhage (1980): *Storage Modification Machines*. *SIAM Journal on Computing* 9(3), pp. 490–508, doi:10.1137/0209036.
- [16] Steven Skiena (2020): *The Algorithm Design Manual*, 3rd ed. edition. *Texts in Computer Science*, Springer, doi:10.1007/978-3-030-54256-6.
- [17] Peter van Emde Boas (1989): *Space measures for storage modification machines*. *Information Processing Letters* 30(2), pp. 103–110, doi:10.1016/0020-0190(89)90117-8.