

Speeding up Generalized PSR Parsers by Memoization Techniques

Mark Minas

Universität der Bundeswehr München
Neubiberg, Germany

mark.minas@unibw.de

Predictive shift-reduce (PSR) parsing for hyperedge replacement (HR) grammars is very efficient, but restricted to a subclass of unambiguous HR grammars. To overcome this restriction, we have recently extended PSR parsing to generalized PSR (GPSR) parsing along the lines of Tomita-style generalized LR parsing. Unfortunately, GPSR parsers turned out to be too inefficient without manual tuning. This paper proposes to use memoization techniques to speed up GPSR parsers without any need of manual tuning, and which has been realized within the graph parser distiller GRAPPA. We present running time measurements for some example languages; they show a significant speed up by some orders of magnitude when parsing valid graphs. But memoization techniques do not help when parsing invalid graphs or if all parses of an ambiguous input graph shall be determined.

1 Introduction

In earlier work [5], we have devised *predictive shift-reduce* parsing (PSR), which lifts D.E. Knuth's LR string parsing [11] to graphs and runs in at most expected linear time in the size of the input graph. However, parsing for graph grammars based on hyperedge replacement (HR) is in general NP-hard, even for a particular grammar [3, sect. 7.1]. Therefore, PSR parsing is restricted to a subclass of HR grammars, which particularly must be unambiguous. We have recently extended PSR parsing to *generalized PSR* (GPSR) parsing [9], which can be applied to every HR grammar.

GPSR parsing has been motivated by *generalized LR* (GLR) parsing for strings, originally devised by M. Tomita [17], and extended and improved by several authors (for an overview see [14]). The original GLR parsing algorithm by Tomita runs in $O(n^{k+1})$ where k is the length of the longest rule, whereas improved versions like *Binary Right Nulled GLR* (BRNGLR) parsers run in worst-case cubic time [15].

GPSR parsing cannot be efficient in general because GPSR parsers can be applied to every HR grammar. But our experiments [9] have shown that GPSR parsers are even slower than simple graph parsers that extend the *Cocke-Younger-Kasami* (CYK) algorithm to graphs [12, 13]. Manual tuning of GPSR parsers by using language specific strategies (see Sect. 4) helped to improve their efficiency, but even those tailored parsers have not always been faster than the corresponding CYK parsers.

GPSR parsers identify parses of an input graph in a search process that may run into dead ends. They are inefficient because they waste time in this process and because they discard all information collected in these dead ends, even if it could be used later. This paper proposes to use memoization techniques to keep the information and to reuse it later. Reuse allows to skip long sequences of parsing operations that would just recreate information that has already been collected earlier.

GPSR parsing with memoization has been implemented in the graph-parser distiller GRAPPA¹. Experiments with generated parsers for different example languages demonstrate that memoization substantially improves parsing speed.

¹Available under www.unibw.de/inf2/grappa.

Note that defining graph languages by graph grammars and using graph parsing is not the only way to check the validity of graphs. A different widespread approach (e.g., [16]) is to use meta-models with additional constraints (e.g., OCL constraints). Checking whether a graph conforms to a given meta-model and the constraints can be easier than graph parsing. But it is generally considered more difficult to come up with a complete set of constraints that are accepted by all valid graphs, but violated by all invalid graphs. For instance, defining just the set of all series-parallel graphs or all flowchart graphs by constraints is non-trivial, but straight-forward with graph grammars (see Sect. 6). In those cases, efficient graph parsing may be favored over meta-model and constraint checking.

The remainder of this paper is structured as follows. After recalling HR grammars in Sect. 2, PSR parsing in Sect. 3, and GPSR parsing in Sect. 4, we describe how memoization can speed up GPSR parsing in Sect. 5. We compare its performance with plain GPSR parsing and with CYK parsing using three example graph languages in Sect. 6: Sierpinski graphs, series-parallel graphs, and structured flowcharts, where GPSR parsing with memoization substantially improves plain GPSR parsing. Sect. 7 concludes the paper.

2 Graph Grammars Based on Hyperedge Replacement

Throughout the paper, we assume that X is a global, countably infinite supply of *nodes*, and that Σ is a finite set of *symbols* that comes with an *arity function* $\text{arity}: \Sigma \rightarrow \mathbb{N}$, and is partitioned into disjoint subsets \mathcal{N} of *nonterminals* and \mathcal{T} of *terminals*.

We write hyperedges with their attached nodes as literals and hypergraphs as ordered sequences of literals. This first may seem unusual, but it will turn out to be beneficial as parsers will read hyperedges of the input hypergraph in a certain order.

Definition 1 (Hypergraph) For a symbol $a \in \Sigma$ and $k = \text{arity}(a)$ pairwise distinct nodes $x_1, \dots, x_k \in X$, a literal $\mathbf{a} = a^{x_1 \dots x_k}$ is a hyperedge that is labeled with a and attached to x_1, \dots, x_k . \mathcal{L}_Σ denotes the set of all literals (over Σ).

A hypergraph $\gamma = \langle V, \phi \rangle$ over Σ consists of a finite set $V \subseteq X$ of nodes and a sequence $\phi = \mathbf{e}_1 \dots \mathbf{e}_n \in \mathcal{L}_\Sigma^*$ of literals such that all nodes in these literals are in V . \mathcal{G}_Σ denotes the set of all hypergraphs over Σ .

We say that two hypergraphs $\gamma = \langle V, \phi \rangle$ and $\gamma' = \langle V', \phi' \rangle$ are equivalent, written $\gamma \bowtie \gamma'$, if $V = V'$ and ϕ is a permutation of ϕ' .

In the following, we usually call hypergraphs just *graphs* and hyperedges just *edges* or *literals*. For a graph $\gamma = \langle V, \phi \rangle$, we use the notation $V_\gamma = V$.

Note that a graph $\langle V, \phi \rangle$ may contain the same literal more than once in ϕ , representing indistinguishable, i.e., parallel edges. Note also that graphs are sequences rather than multisets of literals, i.e., two graphs $\langle V, \phi \rangle$ and $\langle V', \phi' \rangle$ with the same set of nodes, but with different sequences of literals are considered to differ, even if $V = V'$ and ϕ' is just a permutation of ϕ . However, such graphs are equivalent, denoted by the equivalence relation \bowtie . In contrast, “ordinary” graphs would rather be represented using multisets of literals instead of sequences. The equivalence classes of graphs, therefore, correspond to conventional graphs. The ordering of literals is technically convenient for the constructions in this paper. However, input graphs to be parsed should of course be considered up to equivalence. Thus, we will make sure that the developed parsers yield identical results on graphs g, g' with $g \bowtie g'$.

An injective function $\rho: X \rightarrow X$ is called a *renaming*, and γ^ρ denotes the graph obtained by replacing all nodes in γ according to ρ . Although renamings are, for technical simplicity, defined as functions on the whole of X , only the finite subset $V_\gamma \subseteq X$ will be relevant. We define the “concatenation” of two graphs $\gamma = \langle V, \phi \rangle, \gamma' = \langle V', \phi' \rangle \in \mathcal{G}_\Sigma$ as $\gamma\gamma' = \langle V \cup V', \phi\phi' \rangle$. If a graph $\gamma = \langle V, \phi \rangle$ is completely

determined by its sequence ϕ of literals, i.e., if each node in V also occurs in some literal in ϕ , we simply use ϕ as a shorthand for γ . In particular, a literal $\mathbf{a} = \mathbf{a}^{x_1 \dots x_k} \in \mathcal{E}_\Sigma$ is identified with the graph $\langle \{x_1, \dots, x_k\}, \mathbf{a} \rangle$.

A *hyperedge replacement rule* $r = (\mathbf{A} \rightarrow \alpha)$ (rule for short) has a nonterminal edge $\mathbf{A} \in \mathcal{E}_{\mathcal{N}}$ as its *left-hand side*, and a graph $\alpha \in \mathcal{G}_\Sigma$ with $V_{\mathbf{A}} \subseteq V_\alpha$ as its *right-hand side*.

Consider a graph $\gamma = \beta \bar{\mathbf{A}} \bar{\beta} \in \mathcal{G}_\Sigma$ with a nonterminal edge $\bar{\mathbf{A}}$ and a rule $r = (\mathbf{A} \rightarrow \alpha)$. A renaming $\mu: X \rightarrow X$ is a *match* (of r in γ) if $\mathbf{A}^\mu = \bar{\mathbf{A}}$ and if $V_\gamma \cap V_{\alpha^\mu} \subseteq V_{\mathbf{A}^\mu}$.² A match μ of r *derives* γ to the graph $\gamma' = \beta \alpha^\mu \bar{\beta}$. This is denoted as $\gamma \Rightarrow_{r, \mu} \gamma'$. If \mathcal{R} is a finite set of rules, we write $\gamma \Rightarrow_{\mathcal{R}} \gamma'$ if $\gamma \Rightarrow_{r, \mu} \gamma'$ for some match μ of some rule $r \in \mathcal{R}$.

Definition 2 (HR Grammar) A hyperedge replacement grammar $\Gamma = (\Sigma, \mathcal{T}, \mathcal{R}, Z)$ (HR grammar for short) consists of symbols Σ with terminals $\mathcal{T} \subseteq \Sigma$ as assumed above, a finite set \mathcal{R} of rules, and a start graph $Z = Z^\varepsilon$ with $Z \in \mathcal{N}$ of arity 0. Γ generates the language $\mathcal{L}(\Gamma) = \{g \in \mathcal{G}_\Sigma \mid Z \Rightarrow_{\mathcal{R}}^* g\}$.

In the following, we simply write \Rightarrow and \Rightarrow^* because the rule set \mathcal{R} in question will always be clear from the context.

Example 1 (A HR Grammar for Sierpinski Triangles) *The following rules*

$$Z^\varepsilon \xrightarrow[0]{} D^{xyz} \quad D^{xyz} \xrightarrow[1]{} D^{xuw} D^{uyv} D^{wvz} \quad D^{xyz} \xrightarrow[2]{} t^{xyz}$$

generate Sierpinski triangles as graphs where triangles are represented by ternary edges of type t . This grammar is in fact a slightly modified version of [8, p. 189] where edges of triangles are represented by binary edges.

Fig. 1 shows a derivation with graphs as diagrams, in particular with t -edges drawn as triangles. This corresponds to the following derivation. Underlines indicate rewritten nonterminal edges:

$$\begin{aligned} Z^\varepsilon &\xrightarrow[0]{} \underline{D^{ahl}} \xrightarrow[1]{} D^{abc} \underline{D^{bjh}} \underline{D^{cjl}} \xrightarrow[1]{} D^{abc} D^{bde} D^{dhi} \underline{D^{eij}} \underline{D^{cjl}} \\ &\xrightarrow[1]{} D^{abc} D^{bde} D^{dhi} \underline{D^{eij}} \underline{D^{cfg}} \underline{D^{fje}} \underline{D^{gfl}} \xrightarrow[2]{} t^{abc} t^{bde} t^{dhi} t^{eij} t^{cfg} t^{fje} t^{gfl} \end{aligned}$$

3 Predictive Shift-Reduce Parsing

The article [5] gives detailed definitions and correctness proofs for PSR parsing. Here we recall the concepts only so far that we can describe its generalization in the next section.

A PSR parser attempts to construct a derivation by reading the edges of a given input graph one after the other.³ However, the parser must not assume that the edges of the input graph come in the same order as in a derivation. E.g., when constructing the derivation in Fig. 1, it must also accept an input graph $t^{abc} t^{bde} t^{cfg} t^{dhi} t^{eij} t^{fje} t^{gfl}$ where the edges are permuted.

Before parsing starts, a procedure described in [4, Sect. 4] analyzes the grammar for the *unique start node* property, by computing the possible incidences of all nodes created by a grammar. The unique start nodes have to be matched by some nodes in the right-hand side of the start rule of the grammar, thus determining where parsing begins. For our example, the procedure detects that every Sierpinski graph

²I.e., a match μ makes sure that the nodes of α^μ that do not occur in $\bar{\mathbf{A}} = \mathbf{A}^\mu$ do not collide with the other nodes in γ .

³We silently assume that input graphs do not have isolated nodes. This is no real restriction as one can add special edges to such nodes.




Figure 1: A derivation of the graph $t^{abc}t^{bdc}t^{dhi}t^{cij}t^{cfa}t^{fjt}t^{gfi}$ in diagram notation. Nodes are drawn as circles with their names inscribed, nonterminal edges as boxes or triangles around their label, with lines to their attached nodes, and terminal edges as triangles visiting their attached nodes counter-clockwise, starting at the top corner. The numbers inside the terminal triangles are used later to refer to the corresponding terminal edges.

has a unique topmost node. That is a node with a single t -edge attached where the node is the first in the edge's attachments. The node x in the start rule $Z^E \rightarrow D^{xyz}$ must be bound to the topmost node of any input graph.⁴ If the input graph has no topmost node, or more than one, it cannot be a Sierpinski graph, so that parsing fails immediately.

A PSR parser is a push-down automaton that is controlled by a *characteristic finite automaton* (CFA). The stack of the PSR parser consists of states of the CFA. The parser makes sure that the sequence of states on its stack always describes a valid walk through its CFA.

Fig. 2 shows the CFA for our example of Sierpinski graphs. It has been generated by the graph parser distiller GRAPPA¹, using the constructions described in [5], and consists of eight states. Each state has a unique state number and a number of *parameters*, which are written as subscript and superscript, respectively. Parameters are placeholders for nodes of the input graph, which have already been read by the parser. The initial state is q_0^a . Its parameter a is bound to the start node of the input graph, i.e., the topmost node, when parsing starts. Transitions between states are labeled by pairs with a slash as a separator. The first part of a label is the *trigger* of the transition whereas the second part of a label determines the parameters of the target state of the transition. Note that the latter is in fact the target state of the transition with its parameters set to the values used in the label. The trigger is a placeholder for an edge whose attached nodes are either parameters of the source state of the transition, or placeholders \underline{x} or \underline{y} , which stand for nodes of the input graph that have not yet been read by the parser. Note that some transitions have multiple labels. This is in fact a shortcut for different transitions, each

⁴The other two nodes of the start rule, in fact, can be uniquely identified, too, which could be used as a second and a third start node bound to y and z , respectively. However, the corresponding CFA is too complicated for a presentation in this paper.

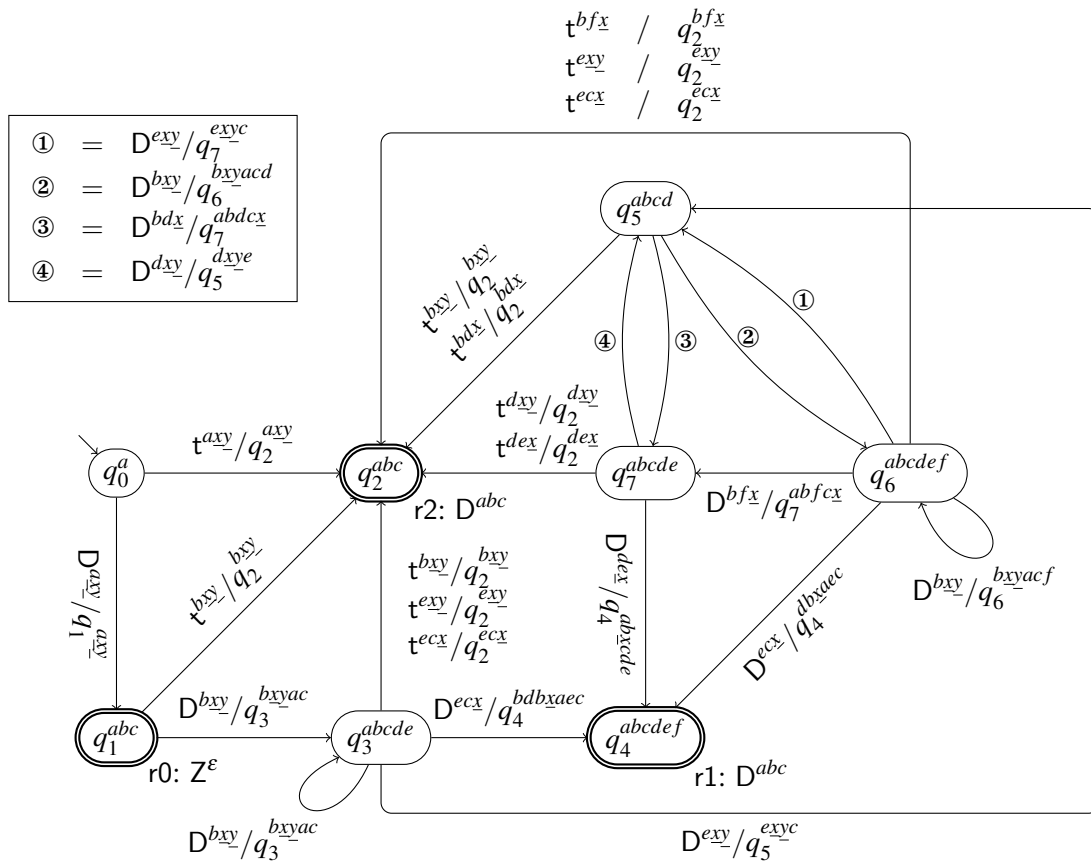


Figure 2: The characteristic finite automaton for the HR grammar of Sierpinski triangles.

with one of these labels. We are going to describe the meaning of labels in the following and shall use $t^{abc}t^{bde}t^{cfg}t^{dhi}t^{ej}t^{fk}t^{gl}$ as an input graph.

A PSR parser starts with a stack that contains the input state with its parameters bound to the start nodes. In our example this is q_0^a with a being bound to node a , written as q_0^a . We call such a state with all its parameters being bound to input graph nodes a *concrete state*. The next action of the parser is always determined by the topmost state on the stack, which is concrete, and by consulting the corresponding state in the CFA. Three different types of actions are distinguished:

A *shift* action reads a yet unread edge of the input graph. This corresponds to an outgoing transition with a terminal trigger. The trigger fits if the input graph contains an unread edge labeled with the trigger label and being attached to input graph nodes as specified by the node placeholders of the trigger. If the topmost state is q_0^a , there is an outgoing transition to state q_2^{abc} with a trigger t^{axy} . Parameter a is bound to a , and its second and third attached nodes must be unread nodes, indicated by x and y . Edge t^{abc} fits this trigger because t^{abc} and b as well as c are yet unread. The shift action marks this edge as well as its attached nodes as read, and pushes the target state of the transition on the stack. The second part of the label determines the binding of this state. In our example, this is q_2^{axy} where a, x, y are bound to a, b, c , respectively. As a consequence, the stack will now contain q_0^a and q_2^{abc} with the latter being the new

topmost state.

A **reduce** action is performed when the top state of the stack corresponds to the right-hand side of a rule which is then replaced by the corresponding left-hand side. The parser recognizes this situation by inspecting just the topmost state of the stack; states that allow a reduce action are marked accordingly. In Fig. 2, these states are drawn with a thick border and additionally labeled by r_0 , r_1 , and r_2 together with a placeholder for a nonterminal edge. For instance, q_2^{abc} is labeled by $r_2:D^{abc}$ where r_2 means a reduction using rule 2 of the grammar. The reduce action in fact consists of three consecutive steps. In the first step, the parser creates a nonterminal as indicated by the state label. In our example, it is $r_2:D^{abc}$. With a topmost state q_2^{abc} , a, b, c are bound to a, b, c , which produces a nonterminal D^{abc} . In the second step of the reduce action, the parser pops as many states off the stack as this rule's right-hand side contains edges, i.e., just one state for rule 2. For instance, when starting with stack contents $q_0^a q_2^{abc}$, q_2^{abc} is popped off, yielding a stack just containing q_0^a . The third step is called a *goto step*. It inspects the new topmost state, i.e., q_0^a here, and selects an outgoing transition whose trigger fits the nonterminal edge produced in the first step, i.e., D^{abc} and the transition to q_1^{abc} . The parser then pushes the target state with its parameters bound according to the transition label. In our example, the stack is then $q_0^a q_1^{abc}$.

An **accept** action is in fact a particular reduce action for the start rule, i.e., rule 0 in our example. The input graph is accepted if the topmost state of the stack is labeled with $r_0:Z^\epsilon$, i.e., state q_1^{abc} in our example, and if all nodes and edges of the input graph are marked as read. In our example with stack contents $q_0^a q_1^{abc}$, the parser has rather reached the accepting state, but there are some unread edges and nodes, i.e., the input graph cannot be accepted yet.

The parser fails if neither a shift, reduce, nor accept action is possible.

As described in [5], such a CFA can be computed for every HR grammar.⁵ But it can control a PSR parser as described above only if its states do not have *conflicts*. A conflict is a situation where the parser must choose between different actions. It is clear that the parser cannot run into a dead end if no state of the CFA has a conflict; the parser can then always predict the correct action which avoids a dead end for valid graphs.⁶ But in the case of conflicts, the parser must choose between several actions; it cannot predict the correct next action. A grammar with such a CFA is not PSR parseable.

Our example grammar for Sierpinski graphs is not PSR parseable because states q_3^{abcde} and q_6^{abcdef} have conflicts. When the parser reaches q_3^{abcde} , for instance, it must read a t-edge in the next shift step, and it must choose between an edge being attached to b (or rather the node that b is bound to) or e , indicated by the transition to q_2^{abc} .

4 Generalized Predictive Shift-Reduce Parsing

In [9] we have proposed *generalized PSR* (GPSR) parsing for grammars that are not PSR parseable. A GPSR parser is primarily a PSR parser that follows all different choices if a state has conflicts. It tries to save time and space in a similar way as Tomita-style GLR parsers for context-free string grammars. Let us briefly summarize how GPSR parsing works.

Whereas a PSR parser maintains a single stack for parsing, a GPSR parser in fact maintains a set of stacks. This set is stored as a so-called *graph-structured stack* (GSS), which is described in the next

⁵[5] describes a simplified algorithm for computing the CFA, which may fail to terminate for some HR grammars. *Grappa*, however, employs a more sophisticated algorithm, which can handle these grammars, too.

⁶This does not necessarily mean that PSR parsers are deterministic; different edges may be chosen for the same shift action. This does not occur in our example of Sierpinski graphs. In general, a grammar can only be PSR parseable if it additionally satisfies the *free edge choice* property [5].

paragraph. For each stack, the parser determines all possible actions based on the CFA as described for the PSR parser. The parser has found a successful parse if the action is *accept* and the entire input graph has been read. (It may proceed if further parses shall be found.) If the parser fails for a specific stack, the parser just discards this stack, stops if this has been the last remaining stack, and fails altogether if it has not found a successful parse previously. If the CFA, however, indicates more than one possible action, the parser duplicates the stack for each of them, and performs each action on one of the copies.

In fact, a GPSR parser does not store complete copies of stacks, but shares their common prefixes and suffixes. The resulting structure is a DAG known as a *graph-structured stack* (GSS) as proposed by M. Tomita [17]. Each node of this DAG (called *GSS node* in the following) is a state. An individual stack is represented as a path in the GSS, from some topmost state to the unique initial state. Working on the GSS instead of on a set of complete copies of different stacks does not only save space, but also time: instead of repeating the same operations on different stacks that share the same suffix, the parser has to perform these actions only once. Furthermore, maintaining the GSS simplifies the construction of all parse trees (the so-called parse forest) of an ambiguous input. But we ignore this aspect in this paper.

Remember that we represent graphs as permutations of edges. By trying out every action offered by the CFA in each step, the GPSR parser effectively performs an exhaustive search in the set of all permutations of the input graph edges permitted by the CFA. This has two immediate effects for a GPSR parser:

1. Consider two different stacks reached by the GPSR parser. Each stack represents a different history of choices the parser has made. In particular, different input graph edges may have been read in these histories. The parser, therefore, cannot globally mark edges as read, but it must store, for each stack separately, which edges of the input graph have been read. In fact, each GSS node keeps track of the set of input graph edges that have been read so far. Note that GSS nodes may be shared only if both their concrete states and their sets of read edges coincide.
2. Whenever the parser has a GSS that represents at least two different stacks, it must choose the stack that it considers next for its actions. It may, for instance, employ a breadth-first strategy or a depth-first strategy. This is in fact the major difference between GLR parsers for context-free string grammars and GPSR parsers for HR grammars: Whenever a GLR parser executes a shift action, this is done for all top-level GSS nodes “simultaneously”. And it then performs all possible reduce actions before the next shift action is executed. As a consequence, each stack encoded in the GSS represents a parse of the same substring of the input string. This is not the case for GPSR parsers. They may be rather forced to try out several reading sequences of the input graph, which may result in exponential complexity.

In [9], we have shown for two example languages (series-parallel graphs and structured flowcharts; see Sect. 6) that the chosen strategy strongly affects the parser speed. In fact, a standard strategy was always too slow, even slower than a simple CYK parser. Instead, specifically tailored strategies have been used that give certain grammar rules preference over others. This requires extra manual work when building a parser and was the motivation for this paper, in particular because even this does not always help in creating a GPSR parser that is faster than a CYK parser.

As a matter of fact, breadth-first and depth-first produce slow parsers for the language of Sierpinski graphs, too. We shall demonstrate this by describing the steps of the GPSR parser for the input graph $t^{abc}t^{bde}t^{cfg}t^{dhi}t^{eij}t^{fj^k}t^{g^{\ell}}$ (see Fig. 1). To save space, we refer to these edges by numbers $1 = t^{abc}$, $2 = t^{bde}$, $3 = t^{cfg}$, $4 = t^{dhi}$, $5 = t^{eij}$, $6 = t^{fj^k}$, $7 = t^{g^{\ell}}$. These numbers correspond to the numbers within the triangles in Fig. 1. And we write GSS nodes in compact form: e.g., 2_{124}^{dhi} refers to the concrete state q_2^{dhi}

0	0_{\emptyset}^a	
1	$0_{\emptyset}^a - 2_1^{abc}$	
2	$0_{\emptyset}^a - 1_1^{abc}$	
3	$0_{\emptyset}^a - 1_1^{abc} - 2_{12}^{bde}$	
4	$0_{\emptyset}^a - 1_1^{abc} - 3_{12}^{bdeac}$	
5	$0_{\emptyset}^a - 1_1^{abc} - 3_{12}^{bdeac} - 2_{123}^{cfe}$ 2_{124}^{dhi}	
6	$0_{\emptyset}^a - 1_1^{abc} - 3_{12}^{bdeac} - 5_{123}^{cfe}$ 2_{124}^{dhi}	
7	$0_{\emptyset}^a - 1_1^{abc} - 3_{12}^{bdeac} - 5_{123}^{cfe} - 2_{1236}^{fjt}$ 2_{124}^{dhi}	
8	$0_{\emptyset}^a - 1_1^{abc} - 3_{12}^{bdeac} - 5_{123}^{cfe} - 6_{1236}^{fjtcge}$ 2_{124}^{dhi}	
9	$0_{\emptyset}^a - 1_1^{abc} - 3_{12}^{bdeac} - 5_{123}^{cfe} - 6_{1236}^{fjtcge} - 2_{12367}^{gtl}$ 2_{124}^{dhi}	
10	$0_{\emptyset}^a - 1_1^{abc} - 3_{12}^{bdeac} - 5_{123}^{cfe} - 6_{1236}^{fjtcge} - 4_{12367}^{cflfgt}$ 2_{124}^{dhi}	
11	$0_{\emptyset}^a - 1_1^{abc} - 3_{12}^{bdeac} - 5_{12367}^{cile}$ 2_{124}^{dhi}	
12	$0_{\emptyset}^a - 1_1^{abc} - 3_{12}^{bdeac} - 2_{124}^{dhi}$	
13	$0_{\emptyset}^a - 1_1^{abc} - 3_{12}^{bdeac} - 3_{124}^{dhibe}$	
14	$0_{\emptyset}^a - 1_1^{abc} - 3_{12}^{bdeac} - 3_{124}^{dhibe} - 2_{1245}^{eij}$	
15	$0_{\emptyset}^a - 1_1^{abc} - 3_{12}^{bdeac} - 3_{124}^{dhibe} - 4_{1245}^{bhjdei}$	
16	$0_{\emptyset}^a - 1_1^{abc} - 3_{1245}^{bhjac}$	
17	$0_{\emptyset}^a - 1_1^{abc} - 3_{1245}^{bhjac} - 2_{12345}^{cfe}$	
18	$0_{\emptyset}^a - 1_1^{abc} - 3_{1245}^{bhjac} - 5_{12345}^{cfe}$	
19	$0_{\emptyset}^a - 1_1^{abc} - 3_{1245}^{bhjac} - 5_{12345}^{cfe} - 2_{123456}^{fjt}$	
20	$0_{\emptyset}^a - 1_1^{abc} - 3_{1245}^{bhjac} - 5_{12345}^{cfe} - 7_{123456}^{cflfgt}$	
21	$0_{\emptyset}^a - 1_1^{abc} - 3_{1245}^{bhjac} - 5_{12345}^{cfe} - 7_{123456}^{cflfgt} - 2_{1234567}^{gtl}$	
22	$0_{\emptyset}^a - 1_1^{abc} - 3_{1245}^{bhjac} - 5_{12345}^{cfe} - 7_{123456}^{cflfgt} - 4_{1234567}^{cflfgt}$	
23	$0_{\emptyset}^a - 1_1^{abc} - 3_{1245}^{bhjac} - 4_{1234567}^{ahlbci}$	
24	$0_{\emptyset}^a - 1_{1234567}^{ahl}$	

Figure 3: Graph-structured stacks and steps of the GPSR parser when parsing the Sierpinski graph with the edges $1 = t^{abc}$, $2 = t^{bde}$, $3 = t^{cfe}$, $4 = t^{dhi}$, $5 = t^{eij}$, $6 = t^{fjt}$, $7 = t^{gtl}$.

and indicates that the edges $1 = t^{abc}$, $2 = t^{bde}$, and $4 = t^{dhi}$ have been read already. Fig. 3 shows the graph-structured stacks after each step of the GPSR parser where a step consists of all actions performed by the parser when working on a specific state. Stacks grow to the right, i.e., the initial state is at the left end whereas topmost states are at the right ends. The steps in fact follow the depth-first strategy which turned out to be a bit faster than the breadth-first strategy.

The parser starts (step 0) with a single stack that contains just 0_{\emptyset}^a , i.e., the initial (concrete) state q_0^a where no edge has been read yet. The first four steps are just PSR steps as described in the previous section: edge $1 = t^{abc}$ is shifted in step 1, a reduce action for rule 2 happens in step 2. Edge $2 = t^{bde}$ is shifted in step 3, and this edge is reduced using rule 2 in step 4, reaching state q_3^{bdeac} . This state allows to shift $3 = t^{cfe}$ as well as $4 = t^{dhi}$ (see Fig. 2), producing two stacks, represented by the GSS after step 5. Note that the topmost states of these stacks are both q_2 -states, but with different parameter bindings and differing sets of read input graph edges. 2_{123}^{cfe} is reduced in step 6, resulting in 5_{123}^{cfe} , which is considered next in step 7 because of the depth-first strategy. In fact, the parser continues working on this stack until it fails in step 12: the stack has the topmost state 5_{12367}^{cile} when step 12 starts, i.e., only $4 = t^{dhi}$ and $5 = t^{eij}$ are yet unread, but they do not fit to any outgoing transition of q_5^{cile} . The parser continues working on the remaining stack, i.e., with topmost state 2_{124}^{dhi} . The remaining steps are again plain PSR steps because the parser does not need to choose between different actions until it accepts the input graph in step 24 in state $1_{1234567}^{ahl}$, i.e., the accepting state with all edges having been read.




Figure 4: Visual representation of the edges that are marked as read after the different steps in Fig. 3.

The parser has in fact wasted time by choosing the topmost state $2_{123}^{c_jg}$ for the next stack to work on in step 6. If it had chosen $2_{124}^{d_hi}$ instead, it would have eventually reached the following GSS in step 17:

$$0_{\emptyset}^a \begin{cases} 1_1^{abc} - 3_{12}^{bdac} - 2_{124}^{dhi} \\ 1_{1234567}^{ahf} \end{cases}$$

i.e., the parser would have found a successful parse after 17 instead of 24 steps. Of course, the parser could then continue with the remaining stack, which would correspond to the steps 6–12 in Fig. 3., i.e., it would not produce further results. The parser can terminate as soon as it has found a parse because the grammar is unambiguous. But even if the grammar were ambiguous, the parser could terminate after the first parse being found if one is interested in just one parse.

So the question remains whether the parser could be improved by more carefully choosing the stack where the parser continues. For this purpose, consider Fig. 4, which shows the diagram of the input graph after steps 1–20 in Fig. 3. It highlights those edges that are marked as read in the state that has just been pushed to the GSS in the corresponding step. In steps 1 and 2, for instance, it is the topmost triangle $1 = t^{abc}$, in steps 3 and 4 triangles $1 = t^{abc}$ as well as $2 = t^{bdc}$, and so on. Fig. 4 does not show the situation for steps 21–24 where all edges are marked as read. As one can see, the parser erroneously “walks down” to the lower right triangles 3, 6, and 7 in steps 6–11, discards the corresponding stack in




Figure 5: A Sierpinski graph where “walking down right first” is faster than “walking down left first”.

step 12, and again walks down the same path in steps 17–20. So one could assume that a strategy that chooses the left walk first (which corresponds to state 2_{124}^{bhi}) instead of the right walk (which corresponds to state 2_{123}^{cig}) would improve the parser behavior. However, this is in general not the case. Fig. 5 shows a Sierpinski graph where walking down right first finds the parse faster than walking down left first (25 vs. 30 steps): Walking left down first reduces triangles 6–8 twice, once before reducing triangles 3–5 and once after that, which is avoided when walking right down first. Apparently, there is not an easy strategy to always find the parse fast. But memoization solves this problem.

5 Memoization

The GPSR parser finds a parse for a valid input graph faster if it either avoids dead ends like the erroneous walk right down in steps 6–11 (Fig. 4) or if the effort spent in such a dead end is not wasted, but is reused later. To see this, let us consider the parsing steps in Fig. 3 more closely. In step 11, it performs a reduce action for rule 1 on state 4_{12367}^{cjlfgt} producing a nonterminal D^{cjl} (see Fig. 2), popping three states off the stack yielding 3_{12}^{bdac} as a (temporary) topmost state and then pushes state 5_{12367}^{cjle} . Moreover, it is known, by comparing the set of read edges of this new topmost state with the set of its predecessor on the stack, that D^{cjl} represents the subgraph consisting of the triangles 3, 6, and 7. But the same nonterminal D^{cjl} representing the same subgraph is again produced in step 23 where the parser performs a reduce action for rule 1 on state $4_{1234567}^{cjlfgt}$, pops three states off the stack yielding state 3_{1245}^{bhjac} , and performs a goto step to $4_{1234567}^{ahlbj}$ triggered by D^{cjl} (see Fig. 2). Note, however, that 3_{1245}^{bhjac} was already the topmost state after step 16. So if the parser remembered that it has produced a D^{cjl} earlier, it could reuse it in step 17 and perform a goto step to $4_{1234567}^{ahlbj}$ right away. As a result, the parser would immediately reach the GSS that Fig. 3 shows after step 23, i.e., the parser would skip six steps and accept the input graph in 18 instead of 24 steps.

In the following, we describe how this observation leads to a systematic approach that allows to skip entire sequences of parsing steps by reusing nonterminal edges that have been produced earlier. This is a *memoization* approach because it depends on memorizing these nonterminal edges.

The main idea is to store a nonterminal edge in a *memo store* whenever it is produced in a reduce action and to look up nonterminals in the memo store whenever the parser reaches a state with an outgoing transition triggered by nonterminal edges. The memo store in fact must store nonterminal edges together with the set of terminal edges that are represented by them. To be more precise, let us assume that the parser analyzes the input graph $h \in \mathcal{G}_{\mathcal{T}}$. The memo store then contains pairs $\langle \mathbf{A}, g \rangle$ where \mathbf{A} is a nonterminal edge, $g \in \mathcal{G}_{\mathcal{T}}$ is a terminal graph with $\mathbf{A} \Rightarrow^* g$ and $h \bowtie gh'$ for some graph $h' \in \mathcal{G}_{\mathcal{T}}$, i.e., g consists of input graph edges. For instance, the memo store after step 16 in Fig. 3 consists of the

following pairs, produced by the reduce actions in one of the previous steps:

$$\{\langle D^{bhj}, 245 \rangle, \langle D^{bd^c}, 2 \rangle, \langle D^{cjl}, 367 \rangle, \langle D^{c^fg}, 3 \rangle, \\ \langle D^{abc}, 1 \rangle, \langle D^{d^hi}, 4 \rangle, \langle D^{eij}, 5 \rangle, \langle D^{fj^e}, 6 \rangle, \langle D^{g^el}, 7 \rangle\},$$

Edges in the second components of pairs are again represented by their numbers.

The lookup operation is controlled by the nodes bound to parameters of the current state, by the (nonterminal) label of the transition, and by the set R of edges that are marked as read in the current state. The lookup operation may return *valid* pairs only. These are pairs $\langle A, g \rangle$ whose graph g does not contain any edge that is also a member of R . Otherwise, edges in g and in R would be read twice.⁷

As an example, let us now consider state 3_{1245}^{bhjac} after step 16. The CFA (Fig. 2) has three outgoing transitions with nonterminal triggers D^{hxy} , D^{cix} , and D^{cxy} when replacing parameters by nodes bound to them. x and y may be bound only to nodes that have not yet been read in state 3_{1245}^{bhjac} . Unread nodes are determined by the set 1245 of read edges, i.e., f, g, e, l are unread in this state. The memo store, therefore, does not contain a pair for D^{hxy} , but it contains $\langle D^{cjl}, 367 \rangle$ for D^{cix} and $\langle D^{c^fg}, 3 \rangle$ for D^{cxy} . Note that $\langle D^{cjl}, 367 \rangle$ does not fit D^{cxy} because node j has been read already. The lookup operation, therefore, has in fact found two valid pairs, and the parser could perform goto actions with both of them. Moreover, it could ignore them both and continue in the regular way, i.e., shift edge $3 = t^{c^fg}$ (see step 17 in Fig. 3). Because the GPSR parser, by design, does not rule out any choice, it will consider all of the three choices here. That way, memoization does not affect the correctness of the parser; if reusing of nonterminals does not lead to acceptance of a valid input graph, regular GPSR will do. But the parser needs a criterion which of the choices to try first. The obvious criterion is to prioritize the nonterminal edge that represents the largest subgraph; the corresponding goto has the potential to skip the longest sequence of parsing steps. In our example, this is $\langle D^{cjl}, 367 \rangle$, i.e., just the case described at the beginning of this section.

The GRAPPA¹ parser distiller has been extended to generate parsers that maintain a memo store in hash tables and that look up all valid pairs when the parser reaches a state with outgoing nonterminal edges. Looked up pairs are ordered by the size of their represented subgraph and tried in that sequence. And it tries the regular GPSR actions if none of these choices leads to acceptance.

6 Evaluation

We now report on running time experiments when parsing Sierpinski graphs. We generated three different parsers: a CYK parser using DIAGEN⁸ [13], a GPSR parser using the depth-first strategy described in Sect. 4, and finally a GPSR parser using the depth-first strategy and memoization as described in the previous section. The GPSR parsers have been generated using GRAPPA, and they stop as soon they can accept the input graph. The CYK parser was in fact optimized in two ways: the parser creates nonterminal edges by dynamic programming, and each of these edges can be derived to a certain subgraph of the input graph. The optimized parser makes sure that it does not create two or more indistinguishable nonterminals for the same subgraph, even if the nonterminals represent different derivation trees (which does not occur here.) And it stops as soon as it finds a derivation of the entire input graph.

Running time of the three parsers has been measured for Sierpinski graphs T_n for different values of n . Each T_n consists of $2n + 1$ triangles. T_0 is just a single triangle, and T_n (for $n > 0$) is made of T_k ,

⁷We assume that there are no parallel edges with the same label. Otherwise, each edge must have a unique name and the lookup operation must make sure that it does not return an edge with a name that is also a member of R .

⁸Homepage: www.unibw.de/inf2/diagen




Figure 6: Recursive definition of Sierpinski graphs T_n for $n > 0$ (left) and running time (in ms) of different parsers analyzing T_n for varying values of n (right).

T_m , and $T_{n-k-m-1}$ as shown in Fig. 6 where $k = \lfloor (n-1)/3 \rfloor$ and $m = \lfloor (n-k-1)/2 \rfloor$, i.e., the $2n+1$ triangles of T_n are as equally distributed to T_k , T_m , and $T_{n-k-m-1}$ as possible.

Fig. 6 shows the running time of the different parsers applied to T_n with varying value n , measured on an iMac 2017, 4.2 GHz Intel Core i7, OpenJDK 12.0.1 with standard configuration, and is shown in milliseconds on the y-axis while n is shown on the x-axis. Note the substantial speed-up when using memoization (called “Memo” in the legend) compared to the plain GPRS parser (called “GPSR”). In fact, the GPSR parser using memoization allows to parse Sierpinski graphs which cannot be parsed in practice by the other two parsers. Moreover, maintaining the memoization store is insignificant with respect to memory consumption: The memoization store grows to $7n+2$ pairs when parsing T_n , i.e., memoization adds only linear space requirements.

Moreover, we reconsider the examples of series-parallel graphs and structured flowcharts, which we have used in [9]:

The following rules generate series-parallel graphs [8, p. 99]:

$$Z^E \xrightarrow{0} G^{xy} \quad G^{xy} \xrightarrow{1} e^{xy} \quad G^{xy} \xrightarrow{2} G^{xy} G^{xy} \quad G^{xy} \xrightarrow{3} G^{xz} G^{zy}$$

Structured flowcharts are flowcharts that do not allow arbitrary jumps, but represent structured programs with conditional statements and while loops. They consist of rectangles containing instructions, diamonds that indicate conditions, and ovals indicating begin and end of the program. Arrows indicate control flow; see Fig. 7 for an example (text within the blocks has been omitted). Flowcharts are easily represented by graphs as also shown in Fig. 7. The following rules generate all graphs representing structured flowcharts:

$$\begin{aligned} Z^E &\rightarrow \text{begin}^x P^{xy} \text{end}^y \\ P^{xy} &\rightarrow S^{xy} \mid P^{xz} S^{zy} \\ S^{xy} &\rightarrow \text{instr}^{xy} \mid \text{cond}^{xuv} P^{uy} P^{vy} \mid \text{cond}^{xy} P^{ux} \end{aligned}$$

None of these grammars is PSR because their CFAs have conflicts. We used these examples in [9] to compare GPSR parsers with CYK parsers. We extend these experiments here and additionally compare these parsers with a GPSR parser using memoization.

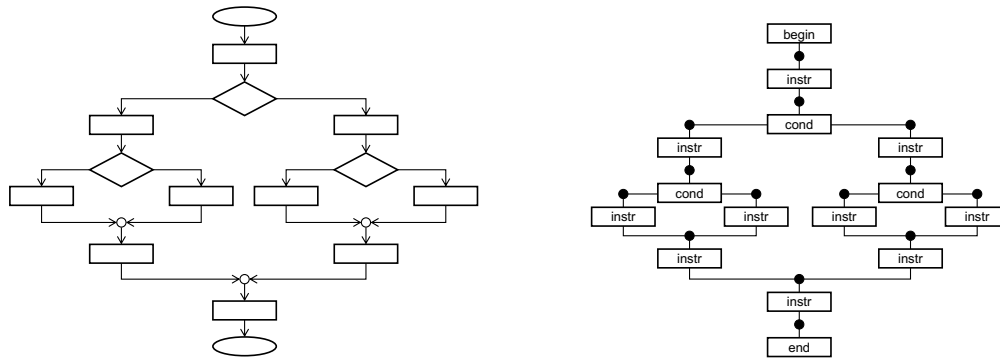


Figure 7: A structured flowchart and its graph representation.




Figure 8: Definition of series-parallel graphs S_n .




Figure 9: Definition of flowchart graphs F_n .

As in [9], we employ GPSR parsers with two different strategies for series-parallel graphs and for structured flowcharts. GPSR 1 employs a breadth-first strategy whereas GPSR 2 applies a more sophisticated strategy. It requires grammar rules to be annotated with either first or second priority. The GPSR 2 parser for series-parallel graphs gives rule 3 (series) precedence over rule 2 (parallel) whereas the GPSR 2 parser for structured flowcharts gives sequences priority over conditional statements.

Running time of the parsers has been measured for series-parallel graphs S_n as shown in Fig. 8 and for flowcharts F_n defined in Fig. 9. Each F_n consists of n conditions and $3n + 1$ instructions. The flowchart in Fig. 7 is in fact F_3 . F_n has a subgraph D_n , which, for $n > 0$, contains subgraphs D_m and $D_{m'}$ with $n = m + m' + 1$. Note that the conditions in F_n form a binary tree with n nodes when we ignore instructions. We always choose m and m' such that it is a complete binary tree. These shapes S_n and F_n turned out to be typical for series-parallel graphs and flowcharts. Other shapes that linearly grow with a parameter n show comparable results and could have been used instead.

Fig. 10 shows the running time of the different parsers applied to S_n and F_n with varying value n on the same platform as for Sierpinski graphs. The experiments again show that the GPSR parser with memoization is substantially faster than the CYK parser and even more faster than most of the GPSR parsers. Only GPSR 2 for structured flowcharts is a bit faster than the memoization parser because it need not maintain the memo store. But note that realizing the hand-tailored strategy for the GPSR 2 parser required additional programming work, whereas the memoization parser has been generated by the GRAPPA distiller without any further manual work.

Maintaining the memoization store when parsing series-parallel graphs and structured flowcharts is insignificant with respect to memory consumption, just as for Sierpinski graphs: The memoization store grows to $38n - 39$ (for $n > 2$) and $18n + 4$ pairs when parsing S_n and F_n , respectively, i.e., memoization

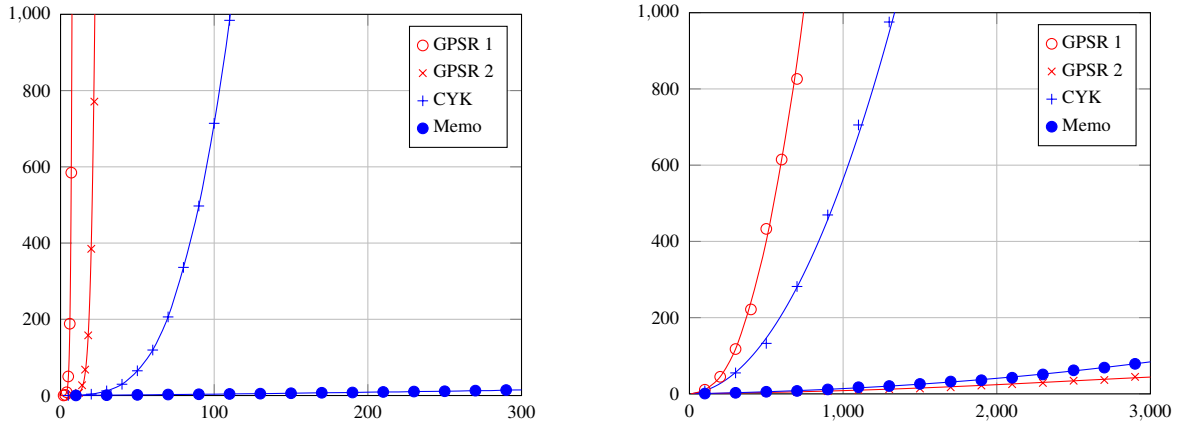


Figure 10: Running time (in ms) of different parsers analyzing series-parallel graphs S_n (left) and structured flowcharts F_n (right) with varying value n .

adds only linear space requirements.

7 Conclusions

We have proposed to use memoization to make GPSR parsing faster by memorizing nonterminal edges that have been created in the search process and that are discarded by plain GPSR parsing although this information could be reused later. Our experiments with three example languages (Sierpinski graphs, series-parallel graphs, and structured flowcharts) have shown that GPSR parsing with memoization is in fact substantially faster and does not increase memory consumption significantly for these examples. However, memoization is not a silver bullet. It cannot speed up GPSR parsing when analyzing invalid input graphs. In these cases, they must completely traverse the entire search space, essentially falling back to plain GPSR parsing. The same applies if one is not only interested in one successful parse, but in all parses if the input graph is ambiguous.

Memoization techniques have also been used to speed up GLR parsers for strings; J.R. Kipps improved the original GLR algorithm from $O(n^{k+1})$ where k is the length of the longest rule to $O(n^3)$ using memoization [10]. And this speed-up is independent of the input string being valid or invalid.⁹ But memoization for GLR parsing differs entirely from memoization for GPSR parsers proposed here: A GLR parser searches for all parses of the input graph in parallel, and all these “parsing processes” are synchronized by reading one input string token after the other. Memoization helps to speed up reduce steps in the graph-structured stack. A GPSR parser, instead, must try different “reading sequences” of the input graph, and memoization helps to reuse information that has been found earlier in a different reading sequence. Kipps’s memoization approach in fact resulted in a parsing algorithm with the same running time complexity as the newer BRNGLR parsers [15], which do not need memoization at all. But parsing for HR grammars is in general NP-hard. So there cannot be a general fast parsing algorithm for HR grammars, and one must depend on techniques like memoization, as suggested in this paper, to obtain efficient parsers.

⁹As a matter of fact, in his thesis, G.R. Economopoulos [6, p. 184] questions the correctness of this claim because Kipps’s parsing algorithm does not terminate on grammars with hidden-left recursion, similar to Tomita’s original GLR parsing algorithm [17].

In future work, we will apply GPSR parsing with memoization to examples from natural language processing, in particular for parsing *Abstract Meaning Representations* (AMR) [1]. PSR parsing cannot be applied there because almost all grammars are ambiguous in this field. In particular, we would like to compare our parser with the state of the art in this field, i.e., the Bolinas parser [2] by D. Chiang, K. Knight *et al.* that implements the polynomial algorithm for HR grammars devised in [12] and the s-graph parser [7] by A. Koller *et al.*

References

- [1] Laura Banarescu, Claire Bonial, Shu Cai, Madalina Georgescu, Kira Griffitt, Ulf Hermjakob, Kevin Knight, Philipp Koehn, Martha Palmer & Nathan Schneider (2013): *Abstract Meaning Representation for Sembanking*. In: *Proc. 7th Linguistic Annotation Workshop at ACL 2013 Workshop*, pp. 178–186. Available at <https://www.aclweb.org/anthology/W13-2322>.
- [2] David Chiang, Jacob Andreas, Daniel Bauer, Karl Moritz Hermann, Bevan Jones & Kevin Knight (2013): *Parsing Graphs with Hyperedge Replacement Grammars*. In: *Proc. 51st Ann. Meeting of the Assoc. for Computational Linguistic (Vol. 1: Long Papers)*, pp. 924–932. Available at <https://www.aclweb.org/anthology/P13-1091>.
- [3] Frank Drewes, Annegret Habel & Hans-Jörg Kreowski (1997): *Hyperedge Replacement Graph Grammars*. In Grzegorz Rozenberg, editor: *Handbook of Graph Grammars and Computing by Graph Transformation. Vol. I: Foundations*, chapter 2, World Scientific, Singapore, pp. 95–162, doi:10.1142/9789812384720_0002.
- [4] Frank Drewes, Berthold Hoffmann & Mark Minas (2016): *Approximating Parikh Images for Generating Deterministic Graph Parsers*. In: *Software Technologies: Applications and Foundations - STAF 2016 Collocated Workshops, Revised Selected Papers, LNCS 9946*, pp. 112–128, doi:10.1007/978-3-319-50230-4_9.
- [5] Frank Drewes, Berthold Hoffmann & Mark Minas (2019): *Formalization and Correctness of Predictive Shift-Reduce Parsers for Graph Grammars based on Hyperedge Replacement*. *Journal of Logical and Algebraic Methods in Programming* 104, pp. 303–341, doi:10.1016/j.jlamp.2018.12.006. Preprint available also at <https://arxiv.org/abs/1812.11927>.
- [6] Giorgios Robert Economopoulos (2006): *Generalised LR parsing algorithms*. Ph.D. thesis, Department of Computer Science, Royal Holloway, University of London.
- [7] Jonas Groschwitz, Alexander Koller & Christoph Teichmann (2015): *Graph parsing with s-graph grammars*. In: *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics, ACL 2015, Volume 1: Long Papers*, pp. 1481–1490, doi:10.3115/v1/P15-1143.
- [8] Annegret Habel (1992): *Hyperedge Replacement: Grammars and Languages*. LNCS 643, Springer, doi:10.1007/BFb0013875.
- [9] Berthold Hoffmann & Mark Minas (2019): *Generalized Predictive Shift-Reduce Parsing for Hyperedge Replacement Graph Grammars*. In Carlos Martín-Vide, Alexander Okhotin & Dana Shapira, editors: *Language and Automata Theory and Applications (LATA 2019), Proceedings, Lecture Notes in Computer Science 11417*, Springer, pp. 233–245, doi:10.1007/978-3-030-13435-8_17.
- [10] James R. Kipps (1991): *GLR Parsing in Time $O(n^3)$* . In Masaru Tomita, editor: *Generalized LR Parsing*, Springer US, Boston, MA, pp. 43–59, doi:10.1007/978-1-4615-4034-2_4.
- [11] Donald E. Knuth (1965): *On the translation of languages from left to right*. *Information and Control* 8(6), pp. 607 – 639, doi:10.1016/S0019-9958(65)90426-2.
- [12] Clemens Lautemann (1990): *The complexity of graph languages generated by hyperedge replacement*. *Acta Informatica* 27, pp. 399–421, doi:10.1007/BF00289017.
- [13] Mark Minas (2002): *Concepts and Realization of a Diagram Editor Generator Based on Hypergraph Transformation*. *Science of Computer Programming* 44(2), pp. 157–180, doi:10.1016/S0167-6423(02)00037-0.

- [14] Elizabeth Scott & Adrian Johnstone (2006): *Right nulled GLR parsers*. *ACM Transactions on Programming Languages and Systems* 28(4), pp. 577–618, doi:10.1145/1146809.1146810.
- [15] Elizabeth Scott, Adrian Johnstone & Giorgios Economopoulos (2007): *BRNGLR: a cubic Tomita-style GLR parsing algorithm*. *Acta Informatica* 44(6), pp. 427–461, doi:10.1007/s00236-007-0054-z.
- [16] Gabriele Taentzer & Arend Rensink (2005): *Ensuring Structural Constraints in Graph-Based Models with Type Inheritance*. In Maura Cerioli, editor: *Fundamental Approaches to Software Engineering*, Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 64–79, doi:10.1007/978-3-540-31984-9_6.
- [17] Masaru Tomita (1985): *An Efficient Context-Free Parsing Algorithm for Natural Languages*. In: *Proceedings of the 9th International Joint Conference on Artificial Intelligence*, pp. 756–764. Available at <http://dl.acm.org/citation.cfm?id=1623611.1623625>.