

1 Comparing Bottom-up with Top-down Parsing 2 Architectures for the Syntax Definition Formalism 3 from a Disambiguation standpoint

4 **Jurgen J. Vinju**   

5 NWO-I Centrum Wiskunde & Informatica

6 TU Eindhoven, The Netherlands

7 **Abstract**

8 Context-free general parsing and disambiguation algorithms are threaded throughout the research and
9 engineering career of Eelco Visser. Both our Ph.D. theses featured the study of “disambiguation.” Disambiguation
10 is the declarative definition of choices among different parse trees, derived using the same context-free grammar,
11 for the same input sentence.

12 This essay highlights the differences between syntactic disambiguation for context-free general parsing in
13 a top-down architecture and a bottom-up architecture. The differences between top-down and bottom-up are
14 mainly observed as practical aspects of the software architecture and software implementation. Eventually, the
15 concept of data-dependent context-free grammar brings all engineering perspectives of disambiguation back
16 into a conceptual (declarative) framework independent of the parsing architecture. The novelty in this essay
17 is the juxtaposition of three general parsing architectures from a disambiguation point of view: SGLR, SGLL,
18 and DDGLL. It also motivates design decisions in the parsing architectures for $SDF\{1,2\}$ and Rascal with
19 previously unpublished detail. The essay falls short of a literature review and a tool evaluation since it does not
20 investigate the disambiguation methods of the many other parser generator tools that exist. The fact that only the
21 implementation algorithms are different between the compared parsing architectures, while the syntax definition
22 formalisms have practically the same formal semantics for historical reasons, nicely “isolates the variable” of
23 interest.

24 We hope this essay lives up to the enormous enthusiasm, curiosity, and drive for perfection in syntax
25 definition and parsing that Eelco always radiated. We dearly miss him.

26 **2012 ACM Subject Classification** Software and its engineering → Syntax

27 **Keywords and phrases** parser generation, context-free grammars, GLR, GLL, algorithms, disambiguation

28 **Digital Object Identifier** [10.4230/OASIScs.EVCS.2023.12](https://doi.org/10.4230/OASIScs.EVCS.2023.12)

29 **Acknowledgements** This essay would not have been possible without the past teamwork with Eelco Visser,
30 Paul Klint, Jan Heering, Jeroen Scheerder, Mark van den Brand, Chris Verhoef, Alex Sellink, Pieter Olivier,
31 Hayco de Jong, Georgios (Rob) Economopoulos, Martin Bravenboer, Tijs van der Storm, Joost Visser, Merijn de
32 Jonge, Jørgen Iversen, Arnold Lankamp, Ali Afroozeh, Anastasia Izmaylova, Bas Basten, Martin Bravenboer,
33 Adrian Johnstone and Elizabeth Scott on the topic of context-free general parsing and disambiguation and
34 supporting technologies. However, any error or inconsistency in the following is all mine.

35 **1 Introduction**

36 This essay focuses on qualitative differences in the design and implementation of Syntax Definition,
37 Parser Generation, and Disambiguation between two classes of Parsing algorithms: GLR and GLL.
38 Disambiguation is a function of an entire Parsing Architecture with the goal of reducing the set of
39 Parse Trees (the Parse Forest) to exactly one using declarative definitions. Extensions of the Syntax
40 Definition Formalisms (languages for context-free grammars) allow expressing preferences between
41 different Parse Trees as produced by the Grammar. It is helpful to see Disambiguation as orthogonal
42 to Parsing, where the latter produces Parse Trees and the former removes them again. However, in
43 actual Parsing Architectures, this distinction is virtually invisible due to efficiency considerations. In
44 this essay, we emphasize declarative and correct parsing over efficiency.



© Jurgen J. Vinju, NWO-I Centrum Wiskunde & Informatica;
licensed under Creative Commons License CC-BY 4.0

Eelco Visser Commemorative Symposium (EVCS 2023).

Editors: Ralf Lämmel, Peter D. Mosses, and Friedrich Steimann; Article No. 12; pp. 12:1–12:15



OpenAccess Series in Informatics

OASIS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

45 **1.1 History of Disambiguation with the SDF**

46 A brief and selective history of disambiguating context-free grammars written in the “Syntax Defini-
 47 tion Formalism” (SDF) is due first. We focus on the parsing architectures of the Syntax Definition
 48 Formalism (SDF) [18, 15, 19] and its later incarnations SDF2 [26, 45] (a part of the “new” ASF+SDF
 49 Meta-Environment [7, 12] and of StrategoXT [13]), and its further parallel offspring in Rascal [24, 25]
 50 and Spoofox [22] (SDF3 [4]). This history motivates the comparison of parsing architectures later
 51 and establishes their origins and dependencies for the sake of full disclosure. We are not comparing
 52 independently designed artifacts.

53 Already in the first SDF from the early 1980’s its users wrote context-free grammars in a BNF-like
 54 format [18], plus regular extensions such as lists and optionals (a.k.a. “EBNF”), plus disambiguation
 55 declarations. The non-terminal notation of SDF was taken from the meta notation used in Paul
 56 Klint’s PhD thesis [23]. These definitions were then used to generate parsers and other useful
 57 language tooling such as unparsers, syntax highlighters, structure editors, etc. SDF used general
 58 parsing algorithms from the very start. Initially, it was based on Jay Earley’s algorithm [16], but
 59 SDF switched to Tomita’s GLR parsing algorithm [37] quickly. Earley’s and GLR would allow any
 60 kind of grammars —not just LL(1) or LR(1) or LALR, but *any*. This was a unique and stimulating
 61 feature for a parser generator, since putting together rules from different modules would also work
 62 (i.e. parsing would happen) and often it would do the right thing. SDF with GLR as its underlying
 63 execution mechanism offered the powers of modularity, language embeddings, and compositionality
 64 of grammars that were eminently useful and also deemed elegant.

65 However, ambiguity and its cure disambiguation are neither modular nor compositional. By this,
 66 we mean that two arbitrary unambiguous modules when composed can easily become ambiguous
 67 and that two arbitrary modules with disambiguation constructs without parsing errors could when
 68 composed, easily generate spurious parsing errors. Ambiguity of context-free grammars is generally
 69 undecidable (with and without disambiguation constructs) [36], and this conundrum is the main
 70 motivation for all our research into the “diagnostics and treatment” of ambiguity in the SDF world [43].
 71 Ambiguity made the SDF sometimes hard to use, despite its elegance and compositionality, or perhaps
 72 because of it. On top of this, the non-determinism of SDF grammars (ambiguous or not) is also a
 73 source of inefficiency of Tomita’s GLR. The same disambiguation constructs that were proposed
 74 would also have a positive effect on efficiency as well.

75 Rewinding, this history starts with the implementation of SDF which was documented in Jan
 76 Rekers’ thesis on incremental general parser generation in 1992. SDF existed before this and
 77 was documented in the technical reports of the ESPRIT project “GIPE - Generating Interactive
 78 Programming Environments” and GIPE II in ESPRIT 2 ([18] not digitized). SDF and the initial
 79 implementations of SDF2 were implemented using ASF+SDF itself, which was implemented in
 80 “LeLisp” [15]. The scanner in SDF was non-deterministic as well: scanning would produce all
 81 possible tokens at a given input position, from which the non-deterministic parsers could choose. This
 82 was necessary to achieve the desired modularity and compositionality of real programming languages
 83 like PL/I, Pascal, and COBOL. Disambiguation was featured in SDF by the associativity and priority
 84 declarations between rules, to help declare the binding strength of unary, binary, and n-ary expression
 85 operators without having to factor a grammar and without introducing any helper non-terminals.
 86 Writing and maintaining SDF grammars was attractive because due to these two disambiguation
 87 constructs the number of rules needed to define the syntax of a language was kept close to the number
 88 of actual constructs in the language. The interactions between the scanner and parser were sometimes
 89 unpredictable due to complex feature interactions with language composition and rules like longest
 90 match and keyword reservation. Incremental parsing was important at that time for feasibility of
 91 running complex and lively updated IDEs on small and slow machines. Wilco Koorn and Jan Rekers
 92 extended the GLR algorithm for substring parsing [32] to that end, and also to improve error recovery

93 and auto-completion IDE features. This pushed the interaction between the parsing and the scanning
94 algorithms to the limit.

95 Eelco's Ph.D. thesis on SDF2 in 1997 followed [45], and we completed an implementation of its
96 parsing and disambiguation mechanisms (SGLR and PGEN) in C and ASF+SDF in 2000 together
97 with Jeroen Scheerder and Mark van den Brand [7]. A main driving force at that time was the
98 COBOL grammar by Chris Verhoef and Ralf Lämmel [27], as well as the "Island Grammars" by Leon
99 Moonen [29] that kept pushing the boundaries of what was possible with declarative disambiguation.
100 Peter Mosses and the Action Notation [9] (design and implementation) we considered important
101 to satisfy as our "customer on-site." Eelco had been inspired by Solomon and Cormack [33] and
102 introduced "scannerless parsing" to SDF by removing the entire scanner from the architecture,
103 thereby introducing lexical ambiguity to the playground of SDF2. This design decision removed
104 the aforementioned hard-to-predict interactions between a non-deterministic scanner and a non-
105 deterministic parser.

106 Eelco's thesis contains the mitigations necessary to make scannerless parsing workable. These
107 are two new disambiguation constructs: *follow restrictions* for longest match and *reject rules* for
108 keyword reservation [39]. He also introduced a simplified representation of parse trees, with only
109 three kinds of nodes: applications of grammar rules, unordered ambiguity clusters, and terminal
110 characters. The grammar rules were represented in the abstract syntax of SDF2 as an algebraic data
111 type in the ATerm format (Pieter Olivier, Hayco de Jong, Mark van den Brand, Paul Klint) [11, 10].
112 This algebraic format of constructors for parse trees, called AsFix2, was derived from earlier parse
113 tree export formats (AsFix1) that were designed by Mark van den Brand to bootstrap ASF+SDF off
114 of the LeLisp implementation [38] and as the intended exchange format between the various tools of
115 the ASF+SDF Meta-Environment [7] that was under development at that time.

116 Since SDF2 was partly implemented in ASF+SDF itself, a bootstrap was required. Mark van
117 den Brand and Pieter Olivier completed the LeLisp-independent ASF+SDF compiler in ASF+SDF
118 in 2001, which included a compilation of the implementation of SDF2 and a re-implementation
119 of the back-end table generator in C. My own thesis work included disambiguation for SDF2 [42].
120 That was in collaboration with Eelco, Jeroen Scheerder, and Mark van den Brand [39], on how to
121 implement the filtering ideas in Eelco's thesis [26, 45]. The work with Rob Economopoulos was
122 about integrating RNGLR [34] into SGLR in 2013 [17], which could simplify some of the filters but
123 complicated others. Diagnosing ambiguity with Bas Basten was a parallel track [43, 5] (2011).

124 SDF2 as-is was used for many years by the ASF+SDF Meta-Environment, ELAN4 environ-
125 ment [12, 41], Action Environment [9] and StrategoXT [13] communities. In this essay we focus on
126 the differences between the bottom-up parsing architecture of SDF2 as it was in the early 2010's and
127 the Rascal top-down architecture in the same period.

128 SDF2's Scannerless Generalized LR parsing algorithm (SGLR) is by Eelco Visser [45]. What
129 makes it different from its predecessors, namely Rekers' fixed version [31] of Tomita's GLR [37], is
130 the semantics or implementation of disambiguation filters specific to scannerless parsing and specific
131 to declarative definitions of operator precedence and associativity in expression grammars. SGLR
132 is the parsing architecture made ready to implement SDF2, the syntax definition formalism from
133 the thesis of Eelco Visser. The predecessor of SDF2, SDF had similar disambiguation constructs
134 for operator precedence, but not for disambiguating lexical syntax. Hence Eelco named "SGLR":
135 "Scannerless GLR", and the efficient and declarative disambiguation of lexical syntax is its core
136 contribution.

137 Lexical ambiguity aside, what we never really solved (at that time) was the context-free ambiguity
138 problem in general. Context-free general parsing produces multiple parse trees, sometimes, and it
139 is hard to predict when. And, SGLR offers specific disambiguation constructs for specific kinds of
140 ambiguity [5], but it can not solve arbitrary ambiguity. So, we experimented (from the very start)

12:4 Comparing Bottom-up with Top-down Parsing from a Disambiguation standpoint

141 with far more general disambiguation constructs, such as the “multiset filter” algorithm [26, 45]. The
142 multiset filter lifts the strict partial order of priority rules to entire parse trees by considering them
143 “sets of rules” and applying a strict partial order of sets of partially ordered elements on entire trees.
144 The more advanced the filters became, the less declarative and the more heuristic they became. Also,
145 new disambiguation filters tended to be hard to implement correctly (more on this later) or efficiently
146 (they all became back-end tree filters). Every new disambiguation concept added to SGLR required
147 almost the effort of a PhD thesis. At the end of the 2010’s, we found a resting point, eventually, to be
148 able to filter parse trees using general purposes tree manipulators, such as Stratego, ASF+SDF, and
149 Rascal—a.k.a. semantics directed disambiguation [8].

150 The second bootstrap stage of Rascal in 2011 provided us with a choice again. We had the
151 opportunity to start from scratch in terms of parsing architecture. Having wrestled with the GLR
152 algorithm for decades, we decided to flip the perspective and go top-down to Scott and Johnstone’s
153 GLL [35]. The one and only motivation was the simplicity and elegance of the new architecture,
154 promising also simple and elegant disambiguation filters. A scannerless version of GLL [35] with
155 disambiguations based on all the filters of SDF2 was indeed produced for Rascal 0.4.x. After this
156 experimenting with new filters became really easy, intuitive, and fast. It is the goal of the current
157 essay to substantiate this story.

158 The PhD thesis of Izmaylova and Afroozeh contains the idea of mapping the semantics of
159 disambiguation filters to (lexical) constraints in data-dependent context-free grammars [1]. The
160 team of Jim, Mandelbaum, and Walker had shown with their Jakker parser generator that data-
161 dependent grammars are an elegant formalism for expressing the unambiguous syntax of programming
162 languages [20, 21]. Moreover, such data-dependent grammars with lexical constraints seem to
163 effectively model almost every hack we have seen in hand-written top-down parsers as well. Even the
164 offside rule in Haskell, which is described as a hack of introducing an extra token in the token stream
165 in an error state, can be simulated using a DDCFG in Iguana [2], and also symbol tables such as used
166 in the scanning of C programs are expressible in data-dependent context-free grammars [1].

167 **2** Comparing Syntax Definition Formalisms

168 These are the three main disambiguation constructs in SDF2 and Rascal:

- 169 ■ **Priorities and associativity** for binding strength of operators in expression languages;
- 170 ■ **Reject rules** for keyword reservation
- 171 ■ **Follow restrictions** for longest and first match

172 **2.1 Associativity and Priority Disambiguation**

173 In [Figure 1](#) the use of priority and associativity declarations is shown for the same language “Exp”
174 in both SDF2 and Rascal. Associativity can be applied to a single rule, e.g. *, or a group of rules
175 (– and +). We see a binary ordering $>$ which is to be interpreted as a strict partial order (transitive,
176 irreflexive, asymmetric) between rules.

177 In both formalisms, each priority and associativity declaration defines a set of disallowed deriva-
178 tion steps in the respective grammar. Namely for “left” associative binary recursive rules a rule in the
179 same group must not be derived on the right-hand side non-terminal of the rule. Vice versa for “right”
180 associativity. If a rule is not binary recursive (i.e. it is of the form $X ::= X \dots X$ then the filter has *no*
181 *effect*.

182 Also, both formalisms share similar semantics for the priority relation. If production A has a
183 higher priority than another production B, then never shall B be a direct child of A. No non-terminal
184 of A will be recognized using the application of rule B. SDF2 priority rules have no effect unless
185 the two rules in question are shaped like so: $X ::= \epsilon X \dots \mid \dots X \epsilon$ or so: $X ::= \dots X \epsilon \mid \epsilon X \dots$

<pre> 1 module ExpInSDF2 2 context-free syntax 3 Id -> Exp 4 "(" Exp ")" -> Exp {bracket} 5 6 context-free priorities 7 Exp "[" Exp "]" -> Exp <0> > 8 Exp "*" Exp -> Exp {left} > 9 { left: 10 Exp "+" Exp -> Exp 11 Exp "-" Exp -> Exp } </pre>	<pre> 1 module ExpInRascal 2 syntax Exp 3 = Id 4 bracket "(" Exp ")" 5 Exp "[" Exp "]" 6 > left Exp "*" Exp 7 > left (Exp "+" Exp 8 Exp "-" Exp 9); 10 _ 11 _ </pre>
---	--

■ **Figure 1** Comparing expression language definition between SDF2 and Rascal; minor meta-syntax differences but conceptually the same.

<pre> 1 module RejectInSDF2 2 context-free syntax 3 Id -> Exp 4 Keyword -> Id {reject} 5 "begin" -> Keyword 6 "end" -> Keyword </pre>	<pre> 1 module RejectInRascal 2 syntax Exp = Id \ Keywords; 3 keyword Keywords 4 = "begin" 5 "end" 6 ; </pre>
--	---

■ **Figure 2** Comparing reject rules between SDF2 and Rascal.

186 The recursive positions must overlap at a left-most or right-most position (or their pre/postfixes must
 187 derive the empty sequence ϵ).

188 For SDF2, if the user makes a mistake and defines a priority relation that is reflexive or symmetric
 189 ($A > B$ and $B > A$), then the filter may remove too many derivations from the generated parser with
 190 parse errors as a result. Also, SDF2 removes *all* nested derivations of B under A if $A > B$, including
 191 the position between brackets in `Exp "[" Exp "]"`. This would make it impossible to write `a[1+2]`. And
 192 so SDF2 users write `< 0 >` before that rule to limit the filter to the first argument position and ignore it
 193 for the second. For Rascal the set of generated filter positions is filtered itself; only the positions that
 194 are guaranteed to generate ambiguity are filtered and the other positions are ignored. If the priority
 195 relation is not a partial order, the user is provided with an error message.

196 In short: the semantics of associativity and priority disambiguation is described in terms of
 197 derivation step filters on the grammar level. We refer to Eelco's thesis [45], this paper on ambiguity
 198 diagnostics [5], and this on disambiguating expression grammars [3], which explain the above
 199 formally.

200 The correctness of the semantics of these constructs relies on the guaranteed ambiguity of binary
 201 expression operators that are left and/or right-recursive, such that the filter does not remove the
 202 last derivation from a Parse Tree [3]. So, their implementation should prevent the effects of such
 203 derivations or remove them, somewhere in the respective parsing architecture in order to implement
 204 this filtering behavior.

205 2.2 Reject Rules

206 **Figure 2** depicts the two styles of reserving keywords as a disambiguation mechanism. The `reject`
 207 tag was introduced by Eelco Visser in 1997. The semantics is that any subsentence recognized by *any*
 208 derivation for `Id` which can also be recognized as `Keyword`, at that same input position and having the
 209 same length, must be filtered. As described, the mechanism extends context-free grammars to include
 210 the *intersection* of context-free non-terminals and it should be possible to generate parsers for the
 211 famous non-context-free example $a^n b^n c^n$. Later we read why this was not accomplished with SDF2.

12:6 Comparing Bottom-up with Top-down Parsing from a Disambiguation standpoint

```
1 module RestrictionsInSDF2           1 module RestrictionsInRascal
2 lexical syntax                     2 lexical Id
3 [A-Za-z][A-Za-z0-9]* -> Id         3 = [A-Za-z][A-Za-z0-9]*
4 lexical restrictions               4 !>> [A-Za-z0-9];
5 Id -/- [A-Za-z0-9]                5 _
```

■ **Figure 3** Comparing follow restriction between SDF2 and Rascal.

212 The Rascal reflection of this design is the `\` operator that removes the language of `keywords` from
213 *that specific* use of `Id` in `Exp`. The difference is thus that Rascal defined a derivation step filter for
214 a specific position of `Id` in a specific rule, while SDF2 defined a filter for all uses of `Id` in any rule.
215 Nevertheless, the expressive power would be the same, if the Rascal designers had not limited the
216 keyword non-terminals to generate non-empty and finite languages only.

217 2.3 Follow restrictions

218 Here in **Figure 3** we even more clearly step out of the realm of context-free grammars. Follow
219 restrictions in Rascal and SDF2 express, literally, constraints on what comes *after* the character yield
220 of a recognized non-terminal. In SDF2 we define a filter that removes all derivations of `Id` anywhere
221 if the single character that would follow it in the input is a member of the character class `[A-Za-z0-9]`.
222 Although this is not a local property of a Parse Tree, it is a local input of most parsing algorithms that
223 move from left to right through the input. The next character or token in the stream is usually referred
224 to as the “lookahead token.” SDF2 also has multi-character lookahead tokens for restrictions, which
225 lead to the same semantics (but an arguably much more complex implementation).

226 In Rascal the semantics is similar, but we define the derivation filter not for all instances of the
227 non-terminal, but only for the position in the rule that the restriction is applied. Next to this Rascal
228 features the complement: a follow requirement declares that a non-terminal *must* be followed by a
229 certain character class, and their analogous duals: precede restrictions and requirements. Rascal, like
230 SDF2, also features multiple look-ahead characters.

231 While explaining the semantics of the three disambiguation constructs we used only ideas such
232 as Grammar, Derivation steps, Parse Trees, and Input sentences. There is no distinction between
233 top-down and bottom-up because we have yet to dive into the implementations of these filters.

234 3 Comparing Parsing and Disambiguation Algorithms


235 We are comparing the *parsing and disambiguation* algorithms SGLR and SGLL as they were part
236 of the parsing architectures for SDF2 and Rascal as they were in 2010. We will show details of the
237 SGLR implementation in C and ASF+SDF and in Java and Rascal of the Rascal implementation.

238 **Figure 4** shows the parsing and disambiguation architecture of SDF2 with SGLR in it. As you can
239 see a disambiguation filter may end up filtering a rule from a grammar completely, modify the SLR
240 parse table to prevent certain derivations to occur at all, or inject itself in the parser run-time to prevent
241 parser driver operations that have the same effect. Eventually, every filter could be implemented by a
242 post-parse tree transformation.


243 **Figure 5** shows a different architecture because there is no intermediate stage for parse table
244 generation. However, Rascal does have a grammar rule merging step while loading a new grammar
245 that captures some of the partial evaluation that is done while building parse tables as well.

246 The two premises of implementing SDF2’s disambiguation constructs are [26, 39]

- 247 ■ efficiency: implement filters as early as possible in the parsing architecture; preventing non-
248 determinism is better than fixing ambiguity later. For Rascal, we adopted a similar but less



■ **Figure 4** Bottom-up SDF2 architecture based on SGLR



■ **Figure 5** Top-down Rascal architecture based on SGLL

- 249 far-fetching dogma: better filter while predicting than filter while accepting a derivation.
- 250 ■ grammar neutrality: do not change the shape of the grammar rules, such that also the shape of
- 251 the Parse Trees is unaffected. This adds to the predictability of the shape of the forests as well as
- 252 efficiency since tree structure does not need to be reconstructed. For Rascal, this is also a core
- 253 design constraint.

254 3.1 Implementing priority and associativity

255 In a bottom-up parser, it is possible to *completely prevent* the creation of derivations that do not satisfy

256 the constraints generated from associativity and priority declarations¹ The major vehicle for this is

257 the parse table generator.

258 SDF2 uses DeRemer's SLR(1) [14] table construction. An SLR(1) parse table is something most

259 students must be able to generate from a grammar by hand in their Compiler Construction course.

260 The table represents a state machine for a pushdown automaton driver that will recognize a string or

261 not, using the information in the table. One could look at the table as a partially evaluated parsing

262 algorithm (typically Earley's algorithm [16]) where the grammar is interpreted but the input sentence

263 is left open. Eelco's idea is that any **reduce** P_1 actions going out of a state can be completely removed

264 if said state witnesses that P_2 is its parent at the wrong position according to $P_2 > P_1$.

265 However, SLR(1)'s follow sets are defined on non-terminals and not production rules, so they do

266 not represent rules that clearly. Every goto action out of a state may represent a union of rules for

267 different non-terminals and different positions in different rules of the same non-terminal. Seeing that

268 we use a non-deterministic parsing driver, Eelco observed that it was possible to redefine SLR follow

269 sets per production instead of per non-terminal; and this enabled a full implementation of the priority

270 and associativity semantics. The overhead is that different rules would exit a state on sometimes the

271 same follow set to the same state, but at least never could an illegal transition be made anymore. The

¹ Later Peter Mosses found out that there are cases where SDF2 is theoretically incomplete for expressing binding strength using single derivation step filters and that required the development of a new theory and new implementations. This is out of the scope of the current paper.

272 breakthrough here was that this solution, on top of Rekers' version of Tomita's GLR, could deal with
 273 *any* context-free grammar and not only the LR(k) class.

274 The actual code that implements this filter in SDF2 is written in ASF+SDF and C (against ApiGen-
 275 generated ATerm interfaces). The expensive part is the transitive closures and cartesian products part,
 276 which was ported to C after the declarative version in ASF+SDF proved to be a bottleneck. This
 277 implementation derives triples (P_1, pos, P_2) that explain for every rule P_1 at which position pos the
 278 other rule P_2 should be disallowed. The worst-case size of the set of these triples is quadratic in the
 279 number of original production rules. Later the relatively simple SLR table generator takes this set as
 280 an additional argument and surgically does not add reductions if they occur in the set (Figure 6).

281 The Rascal implementation of the same filter uses a comparable triplet computation but is written
 282 in (higher-level) Rascal. The Rascal code also first proves ambiguity for left-most and right-most
 283 recursive positions, instead of applying the filter to all arguments of every production. With SGLL,
 284 the set of triplets is not used at parser generation time but during the prediction stages of the top-down
 285 parsing algorithm (Figure 7). Rascal's SGLL (designed and implemented by Arnold Lankamp) filters
 286 rules the moment they are applied by looking at their parent node. At this moment of rule reduction,
 287 the triplet set is queried. This has the exact same effect as the SDF2 implementation, at the cost of a
 288 hash-table lookup. In fact, the SGLL implementation numbers every production with a low integer
 289 and uses a small array to look up all the conflicting rules at a certain position. In a correct SGLL or
 290 SGLR implementation, all these solutions for priority and associativity filtering have the same effect
 291 of removing reductions without breaking the algorithm, at the cost of extra bookkeeping.

292 The SGLL algorithm prevents certain recursive steps dynamically while the SDF2 algorithm
 293 filters certain derivations. From a slight distance, both algorithms simulate a grammar transformation
 294 that would introduce non-terminals for every production rule and disallow certain rules based on
 295 the same constraint triplets. A factored grammar, such as we see when using LALR parsers, would
 296 probably not be much different. One could say that by requiring not to change the grammar, we
 297 have to simulate those grammar changes on a *lower level of abstraction* to achieve the same effect.
 298 Accidentally, a similar "set of integer production rules representation" used by Arnold Lankamp in
 299 SGLL, Eelco had envisioned earlier with "character-class grammars" [44]. There each non-terminal
 300 would also be represented by a set of active production rules for that level in the grammar and
 301 removing a production would entail implementing a filter on the grammar level.

302 The SDF2/SGLR solution requires theory: does the implementation satisfy all the constraints
 303 derived from the semantics of the formalism? Well, only if we change the concept of what an SLR(1)
 304 table is a bit, such that it fits. The solution does not generalize to other standard table formats, like
 305 LALR(1). The Rascal/SGLL solution sits very tightly on the concept of top-down parsing where
 306 recursion on the way down models prediction and coming back up models acceptance/reduction; it
 307 can be added to any (G)LL(1) algorithm implementation technique.

308 3.2 Implementing reject

309 The way we write SDF2 **reject** rules already leaks something about the implementation strategy of
 310 keyword reservation. We simply schedule the rejected rule along with the rest of the grammar. Then,
 311 when *all alternatives* for the `Id` non-terminal have finished, we check if one of them was accidentally
 312 a **reject** rule and if so we remove all those derivations from the computation.

313 With this elegant approach, Eelco had found a near-optimal solution [45]. We do not have to start
 314 a whole other parser with its own stack. Instead, we surf on the non-deterministic graph-structured
 315 stack of Tomita's algorithm and pay only a low overhead of recognizing an additional alternative.

316 However, this implementation had to be revisited and revised many times between the years 2000
 317 and 2005, and eventually, we had to admit general non-terminals could not be "rejected" by this
 318 algorithm. The problem with **reject** is very much akin to the original bug in Tomita's algorithm, which


```

1  ATermList shift_prod(ItemSet items, int prodNr) {
2      Item item, newitem;
3      PT_Symbol symbol;
4      ATermList newvertex = AEmpty;
5      ItemSetIterator iter;
6      PT_Production prod = PGEN_getProductionOfProdNumber(prodNr);
7
8      symbol = PT_getProductionRhs(prod);
9
10     ITS_iteratorPerDotSym(items, symbol, &iter);
11     while (ITS_hasNext(&iter)) {
12         item = ITS_next(&iter);
13         assert(PT_isEqualSymbol(symbol, IT_getDotSymbol(item)));
14         newitem = IT_shiftDot(item);
15         if (newitem != NO_ITEM
16             && !PGEN_isPriorityConflict(item, prodNr)) {
17             newvertex = Ainsert(newvertex, IT_ItemToTerm(newitem));
18         } }
19
20     return newvertex; }

```

■ **Figure 6** C code snippet in the SDF2 parser generator with a single surgical addition predicate on line 16 to implement associativity/priority filtering. Small intervention: big impact.

319 occurred with hidden left recursion. Sometimes the graph-structured stack would miss reductions
320 and Farshi fixed that [30] with an additional stage to search for the missing reductions. The **reject**
321 implementation very much depends on the algorithm identifying a moment where all rules for the
322 restricted non-terminal `Id` have reduced, but the original algorithm for it did not achieve this. The
323 reject “aspect” of SGLR is scattered in different places making it hard to theorize what its effect
324 really is. And so sometimes rules would continue even though they should have been rejected with
325 spurious ambiguity as a result. After several experiments, the diagnosis was left-nullable rules for
326 the restricting could lead to “escaped” reductions. Further complicating the algorithm with yet more
327 searching on top of Farshi’s fix was deemed inefficient.

328 The reject filter in Rascal is either an implementation of an `ICompletionFilter` or an `IEnterFilter`.
329 The latter prevents going into a production when it is predicted while the former removes the effect
330 of recognizing a production when it is completed. This is the general filtering scheme, which gives
331 access to the input character array, the start, and end index of the currently recognized input, etc.
332 Other (static) information, like which non-terminal or production is captured by the object that
333 implements the filter. The reject filter in SGLL simply compares the input subsentence with a given
334 list of keywords that are not allowed and fails on a match. A more complex implementation could
335 parse the substring using another parser (or recognizer). Setting up a nested parser is not as complex
336 or expensive in Java as it was in SGLR’s C version.

337 The parameters of a completion filter in SGLL (see **Figure 8**) make explicit what can be safely
338 used as information to filter without breaking the algorithm’s assumptions. The dynamic programming
339 techniques that are used to stay in polynomial time for an exponential number of parse trees, or even
340 cubic, are predicated upon the identification of reusable parse stacks for reusable subsentences. When
341 context information breaks into this equation, it breaks the underlying assumptions for sharing com-
342 putations leading to false positives (spurious derivations) as well as false negatives (spurious filtering
343 of derivations). Every new filter introduced requires new theory. The definition of `ICompletionFilter`
344 and `IEnterFilter` in Rascal’s SGLL mitigate this by allowing any filter based on the given information
345 and guaranteeing algorithmic correctness. If you need something more, it’s back to the drawing board
346 just as with SGLR.

12:10 Comparing Bottom-up with Top-down Parsing from a Disambiguation standpoint

```
1 private void handleEdgeListWithRestrictions(...) {
2     firstTimeRegistration.clear(); firstTimeReductions.clear();
3     for (int j = edgeSet.size() - 1; j >= 0; --j) {
4         AbstractStackNode<P> edge = edgeSet.get(j);
5         int resultStoreId = getResultStoreId(edge.getId());
6
7         if (!firstTimeReductions.contains(resultStoreId)) {
8             if (firstTimeRegistration.contains(resultStoreId)) continue;
9
10            firstTimeRegistration.add(resultStoreId);
11
12            if (!filteredParents.contains(edge.getId())) {
13                AbstractContainerNode<P> resultStore = null;
14                if (edgeSet.getLastVisitedLevel(resultStoreId) == loc)
15                    resultStore = edgeSet.getLastResult(resultStoreId);
16                ... /* elided error recovery code */
17                resultStore.addAlternative(production, resultLink);
18            } else {
19                AbstractContainerNode<P> resultStore = edgeSet.getLastResult(resultStoreId);
20                stacksWithNonTerminalsToReduce.push(edge, resultStore);
21            } } }
```

■ **Figure 7** Java code snippet in Rascal's SGLL parser run-time, with a single additional predicate on line 14 to filter associativity/priority violations. Again: small intervention; big impact.

```
1 public interface ICompletionFilter {
2     boolean isFiltered(int[] input, int start, int end, PositionStore positionStore); }
```

■ **Figure 8** Rascal's SGLL completion filter interface code.

347 3.3 Implementing follow restrictions

348 The final filter, however innocuous, proved to be another grand challenge for implementing in SDF2
349 and SGLR. Firstly, the basic implementation for single-character lookahead was simply to remove
350 the given characters from the lookahead sets in the SLR(1) table [45]. Secondly, multiple character
351 lookahead would be implemented by dynamically filtering reductions in the inner parser loop. With
352 the right internal administration that associates the lookaheads with every production rule, the code
353 change in the algorithm is minimal.

354 However, at the time we were using the ATerm library for representing parse trees with its *maximal*
355 *sharing* ability. Parse tree nodes that were structurally equal, would always be shared. Sharing was
356 thus based on the contents of the trees' sub-nodes, and not on the context. This design decision is
357 efficient in the context of ambiguity where lots of sub-nodes would be structurally equal (whitespace).
358 With a theory of context-free grammars this all works fine. With the theory of context information
359 with follow restrictions, maximal sharing breaks the parser. To fix the bugs, we ended up introducing
360 an additional lookup table for ambiguous parse trees where the starting position in the input sentence
361 became part of the lookup key.

362 Having learned from the experience, the Rascal SGLL implementation used its own intermediate
363 SPPF-like [37] data structure for parse forests, which is then serialized to AsFix Parse Trees after the
364 parse is done. The aforementioned ICompletionFilter can be used to implement follow restrictions in
365 a single line of code. IEnterFilter would be used for the precede variants. If you'd start from scratch
366 to implement SGLR in Java, for example, you would use the same trick. Indeed the JSGLR code by
367 Karl Trygve Kalleberg in Spoofox, uses an ambiguity hash-table that also includes the start position
368 of every tree.

Feature	Grammar	Implementation
SDF2 Priority / Associativity	$P_1 > P_2$, left , right	<ol style="list-style-type: none"> 1. Compute constraint triplets 2. SLR(1) follow-sets per rule 3. Filtering goto's in SLR(1)
Rascal Priority / Associativity	$P_1 > P_2$, left , right	<ol style="list-style-type: none"> 1. Compute constraint triplets 2. Filter forest edge creation
SDF2 Keyword reservation	$Kw \rightarrow Id$ { reject }	<ol style="list-style-type: none"> 1. Grouping reductions 2. Filter reductions
Rascal Keyword reservation	$Id \setminus Kw$	1. <code>ICompletionFilter</code>
SDF2 Longest Match	$Id \text{ -/- } [A-Z]$	<ol style="list-style-type: none"> 1. Goto Filter on follow sets 2. Reduction filter for k lookahead
Rascal Longest Match	$Id \text{ !>> } [A-Z]$	1. <code>ICompletionFilter</code>

■ **Table 1** Overview of the disambiguation implementations for either SDF2 (bottom-up) or Rascal (top-down).

369 3.4 Top-down disambiguation is easier to get right

370 **Table 1** summarizes how the three declarative implementation constructs were implemented in either
 371 a top-down (Rascal) or bottom-up (SDF2) parsing architecture. Once you have the theory straight, the
 372 implementation of a disambiguation filter seems a surgical incision in either algorithm: a conditional
 373 around the scheduling of the next step. If only this were true.

374 Implementing the Reject and Follow Restriction filters has proven to be complex for the SGLR
 375 architecture while it is straightforward in SGLL. Moreover, the priorities and associativity filter on the
 376 parse table level can never be complete, even though it is elegant. Small changes in the parse table,
 377 such as filtering a goto edge, may break the conditions under which parse stacks or parse tree nodes
 378 can be shared later in the GLR algorithm. These semantic links are platonic, in the sense that they do
 379 not lead to explicit dependencies on the source code level of SGLR, however, when not taken care of
 380 complex bugs do arise. Concretely, the addition of Follow Restrictions to the SLR table construction
 381 algorithm broke many of the underlying assumptions of the SGLR implementation and required the
 382 reconsideration of all of its internal data structures. On the other hand, for SGLL a follow restriction
 383 was a simple conditional while scheduling the next algorithmic step.

384 The reject filter for SGLR proved to be even more difficult to implement. The actual filter
 385 operation is to remove all other reductions for a non-terminal in the presence of a "rejected" one for
 386 the same sub-sentence and non-terminal. The SGLR algorithm does not schedule reductions in such a
 387 way that a clear moment arises when all possible reductions for a subsentence have been collected.
 388 Sometimes graph stack nodes are processed already for further reductions (chain rules), and that way
 389 a tree would escape that would otherwise have been filtered. Sometimes the rejected stack node itself
 390 would be processed too early, letting later nodes escape. The first problem was solved by changing the
 391 GLR algorithm to group reductions in the same starting position of the input. The latter problem was
 392 solved by disallowing more complex non-terminals to be rejected, limiting them to finite non-nullable
 393 languages. The reject filter in SGLL is a simple reduction filter.

394 We conclude that top-down is much easier to experiment with and extend. Bottom-up could be
 395 faster due to partial evaluation, but still, additional bookkeeping and less sharing are required to filter
 396 correctly. Rascal's SGLL in Java is as fast as the SDF2's SGLR in C.

397 **4** Perspective on contextual disambiguation with DDCFGs

398 Let's step back from the comparison made between top-down general and bottom-up general parsing
399 with disambiguation and zoom out to the general problem of disambiguation.

- 400 ■ All three disambiguation constructs use *context information*.
- 401 ■ Each of the three disambiguation constructs is an *ad-hoc* extension of the SDF [43, 6].
- 402 ■ Each disambiguation construct deeply impacts the parsing algorithm.

403 Never mind that they are easier to implement in GLL, but what is the best way of formulating the
404 next disambiguation construct on the SDF level? Say we want to support the offside rule [28]. How
405 to implement it in (S)GLL? Disambiguation always adds “context” to the algorithm of constructing
406 parse trees as compared to “context-free” grammars. At the LDTA conference in 2011, Trevor
407 Jim and Yithzak Mandelbaum demonstrated the utility and elegance of data-dependent context-free
408 grammars with their Yakker parser generator [21, 20]. Much earlier, Mark van den Brand in his PhD
409 thesis [40] also demonstrated that parse-time semantic predicates can be used elegantly and efficiently
410 to disambiguate (lexical and context-free) ambiguity.

411 A Data-dependent Constraint Grammar is a context-free grammar with three major extensions: (a)
412 the non-terminal on the left-hand side of any rule may receive additional data parameters, (b) every
413 symbol on the right-hand side may be conditional on said data parameters using constraint formulas,
414 and (c) data from the input sentence or of syntax trees already processed may be passed as parameters
415 to non-terminals or constraints. Typical “data” would be the character string of a sub-sentence, the
416 start and ending position of every rule, the current indentation level, etc. Typical formulas would
417 be integer arithmetic (for layout positioning), and string (in)equality, but in general, any predicate
418 *without side-effects* written in the host programming language is ok.

419 Afrozeh and Izmaylova [2, 1] mapped all of SDF's and Rascal's disambiguation mechanisms to
420 their Iguana formalism which is based on data-dependent grammars, and then immediately added
421 many more disambiguation constructs. For example, with Iguana it is possible to declaratively express
422 the offside rule and many other “two-dimensional” layout constraints for programming languages
423 such as Haskell and Python. Iguana is a top-down parsing architecture, as the reader might expect. It
424 is also possible to implement data-dependent grammars on top of Earley's algorithm [21, 16].

425 However, the semantics of Disambiguation remains the same whether you implement your
426 DDCFG parsing algorithm in a top-down or a bottom-up data-dependent context-free general frame-
427 work. The formal semantics of data-dependent context-free grammars acts as a virtual machine
428 for disambiguation constructs, making it easier to reason about correctness independent of the
429 implementation in a complex parsing algorithm [1].

430 **5** Conclusion

431 First, contextual disambiguation is a pleonasm. Second, it is arguably easier to design and implement
432 (new) disambiguation constructs with GLL than with GLR. Third, data-dependent context-free
433 grammars add the level of formality and generality that we were always searching for when inventing
434 new disambiguation schemes (as exemplified by the Jakker and Iguana Parsing Architectures). We
435 conclude that a *top-down* implementation of *data-dependent* context-free parsing is the way to go for
436 Rascal as well as SDF3.

437 ——— References ———

- 438 1 Ali Afrozeh and Anastasia Izmaylova. One Parser to Rule Them All. In *Proceedings of the 2015 ACM*
439 *International Symposium on New Ideas, New Paradigms, and Reflections on Programming & Software*,
440 *Onward!* 2015, pages 151–170. ACM, 2015. doi:10.1145/2814228.2814242.

- 441 2 Ali Afroozeh and Anastasia Izmaylova. Iguana: a practical data-dependent parsing framework. In
 442 Ayal Zaks and Manuel V. Hermenegildo, editors, *Proceedings of the 25th International Conference on*
 443 *Compiler Construction, CC 2016, Barcelona, Spain, March 12-18, 2016*, pages 267–268. ACM, 2016.
 444 doi:10.1145/2892208.2892234.
- 445 3 Ali Afroozeh, Mark van den Brand, Adrian Johnstone, Elizabeth Scott, and Jurgen J. Vinju. Safe
 446 specification of operator precedence rules. In *International Conference on Software Language Engineering*
 447 *(SLE)*, LNCS. Springer, 2013.
- 448 4 Luís Amorim and Eelco Visser. *Multi-purpose Syntax Definition with SDF3*, pages 1–23. 09 2020.
 449 doi:10.1007/978-3-030-58768-0_1.
- 450 5 Bas Basten and Jurgen Vinju. Parse forest diagnostics with dr. ambiguity. In *International Conference on*
 451 *Software Language Engineering (SLE)*, LNCS. Springer, 2011.
- 452 6 Bas Basten and Jurgen Vinju. Parse forest diagnostics with Dr. Ambiguity. In *International Conference*
 453 *on Software Language Engineering (SLE)*, LNCS. Springer, 2011.
- 454 7 Mark G.J. van den Brand, Arie van Deursen, Jan Heering, Hayco A. de Jong, Merijn de Jonge, Tobias
 455 Kuipers, Paul. Klint, Leon Moonen, Pieter .A. Olivier, Jeroen Scheerder, Jurgen J. Vinju, Eelco Visser,
 456 and Joost Visser. The ASF+SDF Meta-Environment: a Component-Based Language Development
 457 Environment. In R. Wilhelm, editor, *Compiler Construction (CC '01)*, volume 2027 of *Lecture Notes in*
 458 *Computer Science*, pages 365–370. Springer-Verlag, 2001.
- 459 8 Mark G.J. van den Brand, Steven Klusener, Leon Moonen, and Jurgen J. Vinju. Generalized Parsing and
 460 Term Rewriting - Semantics Directed Disambiguation. In Barret Bryant and João Saraiva, editors, *Third*
 461 *Workshop on Language Descriptions Tools and Applications*, Electronic Notes in Theoretical Computer
 462 Science. Elsevier, 2003.
- 463 9 Mark van den Brand, Jørgen Iversen, and Peter Mosses. An Action Environment. *Electr. Notes Theor.*
 464 *Comput. Sci.*, 110:149–168, 12 2004.
- 465 10 Mark van den Brand and Paul Klint. ATerms for manipulation and exchange of structured data: It's all
 466 about sharing. *Information & Software Technology*, 49:55–64, 01 2007.
- 467 11 Mark van den Brand, Paul Klint, Hayco de Jong, and Pieter Olivier. Efficient annotated terms. *Software—*
 468 *Practice & Experience*, 30(2), January 2000.
- 469 12 Mark van den Brand, Pierre-Etienne Moreau, and Jurgen Vinju. Environments for term rewriting engines
 470 for free! In *Proceedings of the 14th International Conference on Rewriting Techniques and Applications*,
 471 RTA'03, page 424–435, Berlin, Heidelberg, 2003. Springer-Verlag.
- 472 13 Martin Bravenboer, Karl Trygve Kalleberg, Rob Vermaas, and Eelco Visser. Stratego/XT 0.17. a language
 473 and toolset for program transformation. *Science of Computer Programming*, 72(1):52–70, 2008. Special
 474 Issue on Second issue of experimental software and toolkits (EST). doi:https://doi.org/10.1016/j.
 475 scico.2007.11.003.
- 476 14 Frank DeRemer. Simple lr(k) grammars. *Commun. ACM*, 14(7):453–460, 1971. doi:10.1145/362619.
 477 362625.
- 478 15 Arie Van Deursen, Jan Heering, and Paul Klint. *Language Prototyping: An Algebraic Specification*
 479 *Approach: Vol. V*. World Scientific Publishing Co., Inc., USA, 1996.
- 480 16 Jay Earley. An efficient context-free parsing algorithm. *Commun. ACM*, 13:94–102, February 1970.
 481 doi:http://doi.acm.org/10.1145/362007.362035.
- 482 17 Giorgios R. Economopoulos, Paul Klint, and Jurgen J. Vinju. Faster scannerless GLR parsing. In
 483 Oege de Moor and Michael I. Schwartzbach, editors, *Compiler Construction (CC)*, volume 5501 of
 484 *Lecture Notes in Computer Science*, pages 126–141. Springer, 2009. doi:http://dx.doi.org/10.1007/
 485 978-3-642-00722-4_10.
- 486 18 J. Heering and P. Klint. A syntax definition formalism. Technical report, 1986. ESPRIT'86: Results and
 487 Achievements, page 619–630.
- 488 19 Jan Heering, Paul R. H. Hendriks, Paul Klint, and Jan Rekers. The syntax definition formalism SDF —
 489 reference manual. *SIGPLAN Not.*, 24(11):43–75, nov 1989. doi:10.1145/71605.71607.
- 490 20 Trevor Jim and Yitzhak Mandelbaum. Delayed semantic actions in yakker. In Claus Brabrand and Eric Van
 491 Wyk, editors, *Language Descriptions, Tools and Applications, LDTA 2011, Saarbrücken, Germany, March*
 492 *26-27, 2011. Proceeding*, page 8. ACM, 2011. doi:10.1145/1988783.1988791.

- 493 21 Trevor Jim, Yitzhak Mandelbaum, and David Walker. Semantics and algorithms for data-dependent
494 grammars. In Manuel V. Hermenegildo and Jens Palsberg, editors, *Proceedings of the 37th ACM SIGPLAN-*
495 *SIGACT Symposium on Principles of Programming Languages, POPL 2010, Madrid, Spain, January*
496 *17-23, 2010*, pages 417–430. ACM, 2010. doi:[10.1145/1706299.1706347](https://doi.org/10.1145/1706299.1706347).
- 497 22 Lennart C.L. Kats and Eelco Visser. The Spoofox language workbench: Rules for declarative specification
498 of languages and IDEs. *SIGPLAN Not.*, 45(10):444–463, oct 2010.
- 499 23 Paul Klint. *From SPRING to SUMMER: design, definition and implementation of programming languages*
500 *for string manipulation and pattern matching*. PhD thesis, Technische Hogeschool Eindhoven, March
501 1982.
- 502 24 Paul Klint, Tijs van der Storm, and J.J. Vinju. EASY meta-programming with Rascal. In João Fernandes,
503 Ralf Lämmel, Joost Visser, and João Saraiva, editors, *Generative and Transformational Techniques in*
504 *Software Engineering III*, volume 6491 of *LNCSE*, pages 222–289. Springer Berlin / Heidelberg, 2011.
- 505 25 Paul Klint, Tijs van der Storm, and Jurgen J. Vinju. Rascal: A domain specific language for source
506 code analysis and manipulation. In *Ninth IEEE International Working Conference on Source Code*
507 *Analysis and Manipulation (SCAM)*, pages 168–177. IEEE Computer Society, 2009. doi:[http://doi.](http://doi.ieeecomputersociety.org/10.1109/SCAM.2009.28)
508 [ieeecomputersociety.org/10.1109/SCAM.2009.28](http://doi.ieeecomputersociety.org/10.1109/SCAM.2009.28).
- 509 26 Paul Klint and Eelco Visser. Using filters for the disambiguation of context-free grammars. In G. Pighizzini
510 and P. San Pietro, editors, *Proc. ASMICS Workshop on Parsing Theory*, pages 1–20, Milano, Italy, 1994.
511 Tech. Rep. 126–1994, Dipartimento di Scienze dell’Informazione, Università di Milano.
- 512 27 Ralf Lämmel and Chris Verhoef. Semi-automatic Grammar Recovery. *Software—Practice & Experience*,
513 31(15):1395–1438, December 2001.
- 514 28 P. J. Landin. The next 700 programming languages. *Commun. ACM*, 9:157–166, March 1966.
- 515 29 Leon Moonen. Generating robust parsers using island grammars. *Proceedings Eighth Working Conference*
516 *on Reverse Engineering*, pages 13–22, 2001.
- 517 30 Rohman Nozohoor-Farshi. Handling of ill-designed grammars in tomita’s parsing algorithm. In *Pro-*
518 *ceedings of the First International Workshop on Parsing Technologies*, pages 182–192, Pittsburgh,
519 Pennsylvania, USA, August 1989. Carnegy Mellon University. URL: [https://aclanthology.org/](https://aclanthology.org/W89-0219)
520 [W89-0219](https://aclanthology.org/W89-0219).
- 521 31 J. Rekers. *Parser Generation for Interactive Environments*. PhD thesis, University of Amsterdam, 1992.
- 522 32 Jan Rekers and Wilco Koorn. Substring parsing for arbitrary context-free grammars. In *Proceedings of the*
523 *Second International Workshop on Parsing Technologies*, pages 218–224, Cancun, Mexico, February 13-25
524 1991. Association for Computational Linguistics. URL: <https://aclanthology.org/1991.iwpt-1.25>.
- 525 33 D. J. Salomon and G. V. Cormack. Scannerless NSLR(1) parsing of programming languages. In *Pro-*
526 *ceedings of the ACM SIGPLAN 1989 Conference on Programming language design and implementation*,
527 *PLDI 1989*, pages 170–178. ACM, 1989. doi:<http://doi.acm.org/10.1145/73141.74833>.
- 528 34 Elizabeth Scott and Adrian Johnstone. Right nulled GLR parsers. *ACM Trans. Program. Lang. Syst.*,
529 28(4):577–618, jul 2006.
- 530 35 Elizabeth Scott and Adrian Johnstone. GLL parsing. *ENTCS*, 253(7):177 – 189, 2010. Proceedings of the
531 Ninth Workshop on Language Descriptions Tools and Applications (LDTA 2009).
- 532 36 Thomas A. Sudkamp. *Languages and Machines: An Introduction to the Theory of Computer Science*.
533 Addison-Wesley Longman Publishing Co., Inc., USA, 1997.
- 534 37 M. Tomita. *Efficient Parsing for Natural Languages. A Fast Algorithm for Practical Systems*. Kluwer
535 Academic Publishers, 1985.
- 536 38 Mark van den Brand, Jan Heering, Paul Klint, and Pieter A. Olivier. Compiling language definitions: The
537 ASF+SDF compiler. *CoRR*, cs.PL/0007008, 2000. URL: <https://arxiv.org/abs/cs/0007008>.
- 538 39 Mark van den Brand, Jeroen Scheerder, Jurgen J. Vinju, and Eelco Visser. Disambiguation filters
539 for scannerless generalized LR parsers. In R. Nigel Horspool, editor, *Compiler Construction, 11th*
540 *International Conference, CC 2002*, volume 2304 of *LNCSE*, pages 143–158. Springer, 2002.
- 541 40 Mark G. J. van den Brand. *PREGMATIC - a generator for incremental programming environments*. PhD
542 thesis, Radboud University Nijmegen, 1992.
- 543 41 Mark G. J. van den Brand, Pierre-Etienne Moreau, and Christophe Ringeissen. The ELAN Environment:
544 an Rewriting Logic Environment based on ASF+SDF Technology. In *Workshop on Language Descriptions*,

- 545 *Tools and Applications - LDTA'02*, volume 65/3 of *Electronic Notes in Theoretical Computer Science*,
546 page 7 p, Grenoble, France, April 2002. Colloque avec actes et comité de lecture. internationale. URL:
547 <https://hal.inria.fr/inria-00101028>.
- 548 42 J.J. Vinju. *Analysis and Transformation of Source Code by Parsing and Rewriting*. PhD thesis, Universiteit
549 van Amsterdam, November 2005.
- 550 43 Jurgen J. Vinju. SDF disambiguation medkit for programming languages. Technical Report SEN-1107,
551 Centrum Wiskunde & Informatica, 2011. <http://oai.cwi.nl/oai/asset/18080/18080D.pdf>.
- 552 44 Eelco Visser. From context-free grammars with priorities to character class grammars. In Mieke Brune
553 Arie van Deursen and Jan Heering, editors, *Dat Is Dus Heel Interessant, Liber Amicorum dedicated to*
554 *Paul Klint*. Centrum Wiskunde & Informatica and IVI Universiteit van Amsterdam, 1997.
- 555 45 Eelco Visser. *Syntax Definition for Language Prototyping*. PhD thesis, Universiteit van Amsterdam, 1997.