

MARPA, A PRACTICAL GENERAL PARSER: THE RECOGNIZER

JEFFREY KEGLER

ABSTRACT. The Marpa recognizer is described. Marpa is a practical and fully implemented algorithm for the recognition, parsing and evaluation of context-free grammars. The Marpa recognizer is the first to unite the improvements to Earley’s algorithm found in Joop Leo’s 1991 paper to those in Ayrcock and Horspool’s 2002 paper. Marpa tracks the full state of the parse, at it proceeds, in a form convenient for the application. This greatly improves error detection and enables event-driven parsing. One such technique is “Ruby Slippers” parsing, in which the input is altered in response to the parser’s expectations.

1. INTRODUCTION

Despite the promise of general context-free parsing, and the strong academic literature behind it, it has never been incorporated into a highly available tool like those that exist for LALR[6] or regular expressions. The Marpa project was intended to take the best results from the literature on Earley parsing off the pages of the journals and bring them to a wider audience. Marpa::XS[9], a stable version of this tool, was uploaded to the CPAN Perl archive on Solstice Day in 2011. This paper describes the algorithm of Marpa::R2[8], a later version.

As implemented, Marpa parses, without exception, all context-free grammars. Time bounds are the best of Leo[10] and Earley[3]. The Leo bound, $\mathcal{O}(n)$ for LR-regular grammars, is especially relevant to Marpa’s goal of being a practical parser: If a grammar is in a class of grammar currently in practical use, Marpa parses it in linear time.

Error-detection properties, extremely important, have been overlooked in the past. Marpa breaks new ground in this respect. Marpa has the immediate error detection property, but goes well beyond that:

Date: June 19, 2013.

Copyright © 2013 Jeffrey Kegler.

This document is licensed under a Creative Commons Attribution-NoDerivs 3.0 United States License.

it is fully aware of the state of the parse, and can report this to the user while tokens are being scanned.

Marpa allows the lexer to check its list of acceptable tokens before a token is scanned. Because rejection of tokens is easily and efficiently recoverable, the lexer is also free to take an event-driven approach. Error detection is no longer an act of desperation, but a parsing technique in its own right. If a token is rejected, the lexer is free to create a new token in the light of the parser's expectations. This approach can be described as making the parser's "wishes" come true, and we have called this "Ruby Slippers Parsing".

One use of the Ruby Slippers technique is to parse with a clean but oversimplified grammar, programming the lexical analyzer to make up for the grammar's short-comings on the fly. The author has implemented an HTML parser[7], based on a grammar that assumes that all start and end tags are present. Such an HTML grammar is too simple even to describe perfectly standard-conformant HTML, but the lexical analyzer is programmed to supply start and end tags as requested by the parser. The result is a very simply and cleanly designed parser that parses very liberal HTML and accepts all input files, in the worst case treating them as highly defective HTML.

Section 2 describes the notation and conventions of this paper. Section 3 deals with Marpa's grammar rewrites. Sections 4 and 5 introduce Earley's algorithm. Section 6 describes Leo's modification to Earley's algorithm. Section 7 describes the modifications proposed by Aycock and Horspool. Section 8 presents the pseudocode for Marpa's recognizer. Section 9 describes notation and deals with other preliminaries to the theoretical results. Section 10 contain a proof of Marpa's correctness, while Section 11 contains its complexity results. Finally, section 12 generalizes Marpa's input model.

2. PRELIMINARIES

We assume familiarity with the theory of parsing, as well as Earley's algorithm. This paper will use subscripts to indicate commonly occurring types.

X_T	The variable X of type T
$\text{set-one}_{\{T\}}$	The variable set-one of type set of T
SYM	The type for a symbol
a_{SYM}	The variable a of type SYM
$\text{set-two}_{\{SYM\}}$	The variable set-two , a set of symbols

Subscripts may be omitted when the type is obvious from the context. The notation for constants is the same as that for variables. Multi-character variable names will be common, and operations will never be implicit.

Multiplication $a \times b$
 Concatenation $a . b$
 Subtraction $\text{symbol-count} - \text{terminal-count}$

Type names are often used in the text as a convenient way to refer to their type.

Where $\text{vocab}_{\{SYM\}}$ is non-empty set of symbols, let vocab^* be the set of all strings (type STR) formed from those symbols. Where s_{STR} is a string, let $|s_{STR}|$ be its length, counted in symbols. Let vocab^+ be

$$\{x_{STR} \mid x_{STR} \in \text{vocab}^* \wedge |x_{STR}| > 0\}.$$

In this paper we use, without loss of generality, the grammar g , where g is the 3-tuple

$$(\text{vocab}_{\{SYM\}}, \text{rules}, \text{accept}_{SYM}).$$

Here $\text{accept}_{SYM} \in \text{vocab}$. Call the language of g , $L(g)$, where $L(g) \subseteq \text{vocab}^*$.

$\text{rules}_{\{RULE\}}$ is a set of rules (type $RULE$), where a rule is a duple of the form $[\text{lhs}_{SYM} \rightarrow \text{rhs}_{STR}]$, such that

$$\text{lhs}_{SYM} \in \text{vocab} \quad \text{and} \quad \text{rhs}_{STR} \in \text{vocab}^+.$$

lhs_{SYM} is referred to as the left hand side (LHS) of r_{RULE} . rhs_{STR} is referred to as the right hand side (RHS) of r_{RULE} . The LHS and RHS of r_{RULE} may also be referred to as $\text{LHS}(r_{RULE})$ and $\text{RHS}(r_{RULE})$, respectively. This definition follows [2], which departs from tradition by disallowing an empty RHS.

The rules imply the traditional rewriting system, in which $x_{STR} \Rightarrow y_{STR}$ states that x_{STR} derives y_{STR} in exactly one step; $x_{STR} \xrightarrow{+} y_{STR}$ states that x_{STR} derives y_{STR} in one or more steps; and $x_{STR} \xrightarrow{*} y_{STR}$ states that x_{STR} derives y_{STR} in zero or more steps.

We say that symbol x_{SYM} is **nullable** if and only if $x_{SYM} \xrightarrow{*} \epsilon$. x_{SYM} is **nonnull** if and only if it is not nullable. Following Aycock and Horspool[2], all nullable symbols in grammar g are nulling – every symbol which can derive the null string always derives the null string. It is shown in [2] how to do this without losing generality or the ability to efficiently evaluate a semantics that is defined in terms of an original grammar that includes symbols which are both nullable and non-nulling, empty rules, etc.

Also without loss of generality, it is assumed that there is a dedicated acceptance rule, accept_{RULE} and a dedicated acceptance symbol, $\text{accept}_{SYM} = \text{LHS}(\text{accept}_{RULE})$, such that for all \mathbf{x}_{RULE} ,

$$\begin{aligned} & \text{accept}_{SYM} \notin \text{RHS}(\mathbf{x}_{RULE}) \\ \wedge & (\text{accept}_{SYM} = \text{LHS}(\mathbf{x}_{RULE}) \implies \text{accept}_{RULE} = \mathbf{x}_{RULE}). \end{aligned}$$

We define rightmost non-null symbol of a string as

$$\begin{aligned} \text{Right-NN}(\mathbf{x}_{STR}) & \stackrel{\text{def}}{=} \text{rnn}_{SYM} \text{ such that } \exists \text{pre}_{STR}, \text{post}_{STR} \mid \\ & \mathbf{x}_{STR} = \text{pre}_{STR} \cdot \text{rnn}_{SYM} \cdot \text{post}_{STR} \\ & \wedge \text{post}_{STR} \xrightarrow{*} \epsilon \\ & \wedge \neg(\text{rnn}_{SYM} \xrightarrow{*} \epsilon). \end{aligned}$$

We define the rightmost non-null symbol of a rule as

$$\text{Right-NN}([\text{lhs}_{SYM} \rightarrow \text{rhs}_{STR}]) \stackrel{\text{def}}{=} \text{Right-NN}(\text{rhs}_{STR}).$$

A rule \mathbf{x}_{RULE} is **directly right-recursive** if and only if

$$\text{LHS}(\mathbf{x}_{RULE}) = \text{Right-NN}(\mathbf{x}_{RULE}).$$

\mathbf{x}_{RULE} is **indirectly right-recursive** if and only if

$$\exists \mathbf{y}_{STR} \mid \text{Right-NN}(\mathbf{x}_{RULE}) \xrightarrow{+} \mathbf{y}_{STR} \wedge \text{Right-NN}(\mathbf{y}_{STR}) = \text{LHS}(\mathbf{x}_{RULE}).$$

\mathbf{x}_{RULE} is **right recursive**, $\text{Right-Recursive}(\mathbf{x}_{RULE})$, if and only if is either directly or indirectly right-recursive.

The definition of \mathbf{g} does not sharply distinguish terminals from non-terminals. MARPA's implementations allow terminals to be the LHS of rules, and every symbol except accept_{SYM} can be a terminal. The implementations have options that allow the user to reinstate the traditional restrictions, in part or in whole. Note that, as a result of these definitions, sentential forms will be of type STR .

Let the input to the parse be \mathbf{w} such that $\mathbf{w} \in \text{vocab}^+$. Locations in the input will be of type LOC . Let $|\mathbf{w}|$ be the length of the input, counted in symbols. When we state our complexity results later, they will often be in terms of \mathbf{n} , where $\mathbf{n} = |\mathbf{w}|$. Let $\mathbf{w}[\mathbf{i}]_{SYM}$ be character \mathbf{i} of the input, $0 \leq \mathbf{i}_{LOC} < |\mathbf{w}|$.

The alert reader may have noticed that the previous definition of \mathbf{w} did not allow zero-length inputs. To simplify the mathematics, we exclude null parses and trivial grammars from consideration. In its implementations, the Marpa parser deals with null parses and trivial grammars as special cases. (Trivial grammars are those that recognize only the null string.)

In this paper, `EARLEY` will refer to the Earley’s original recognizer[3]. `LEO` will refer to Leo’s revision of `EARLEY` as described in [10]. `AH` will refer to the Aycock and Horspool’s revision of `EARLEY` as described in [2]. `MARPA` will refer to the parser described in this paper. Where `RECCE` is a recognizer, $L(\text{RECCE}, g)$ will be the language accepted by `RECCE` when parsing g .

3. REWRITING THE GRAMMAR

We have already noted that no rules of g have a zero-length RHS, and that all symbols must be either nulling or non-nullable. These restrictions follow Aycock and Horspool[2]. The elimination of empty rules and proper nullables is done by rewriting the grammar. [2] shows how to do this without loss of generality.

Because Marpa claims to be a practical parser, it is important to emphasize that all grammar rewrites in this paper are done in such a way that the semantics of the original grammar can be reconstructed simply and efficiently at evaluation time. As one example, when a rewrite involves the introduction of new rule, semantics for the new rule can be defined to pass its operands up to a parent rule as a list. Where needed, the original semantics of a pre-existing parent rule can be “wrapped” to reassemble these lists into operands that are properly formed for that original semantics.

As implemented, the Marpa parser allows users to associate semantics with an original grammar that has none of the restrictions imposed on grammars in this paper. The user of a Marpa parser may specify any context-free grammar, including one with properly nullable symbols, empty rules, etc. The user specifies his semantics in terms of this original, “free-form”, grammar. Marpa implements the rewrites, and performs evaluation, in such a way as to keep them invisible to the user. From the user’s point of view, the “free-form” of his grammar is the one being used for the parse, and the one to which his semantics are applied.

4. EARLEY’S ALGORITHM

Let $r_{RULE} \in \text{rules}$ be a rule, and $|r|$ the length of its RHS. A dotted rule (type *DR*) is a duple, $[r_{RULE}, \text{pos}]$, where $0 \leq \text{pos} \leq |r_{RULE}|$. The position, `pos`, indicates the extent to which the rule has been recognized, and is represented with a large raised dot, so that if

$$[A_{SYM} \rightarrow X_{SYM} \cdot Y_{SYM} \cdot Z_{SYM}]$$

is a rule,

$$[A_{SYM} \rightarrow X.Y \bullet Z]$$

is the dotted rule with the dot at $\text{pos} = 2$, between Y_{SYM} and Z_{SYM} .

If we let x_{DR} be a dotted rule, such that

$$x_{DR} = [[A_{SYM} \rightarrow \text{pre}_{STR} \cdot \text{next}_{SYM} \cdot \text{post}_{STR}], \text{pos}],$$

then

$$\begin{aligned} \text{LHS}(x_{DR}) &\stackrel{\text{def}}{=} A_{SYM} \\ \text{Postdot}(x_{DR}) &\stackrel{\text{def}}{=} \begin{cases} \text{next}_{SYM}, & \text{if } x = [A \rightarrow \text{pre} \bullet \text{next} \cdot \text{post}] \\ \Lambda, & \text{if } x = [A \rightarrow \text{pre} \cdot \text{next} \cdot \text{post} \bullet] \end{cases} \\ \text{Next}(x_{DR}) &\stackrel{\text{def}}{=} \begin{cases} [A \rightarrow \text{pre} \cdot \text{next} \bullet \text{post}], & \\ \text{if } \text{Postdot}(x_{DR}) = \text{next} & \\ \Lambda, & \text{otherwise} \end{cases} \\ \text{Penult}(x_{DR}) &\stackrel{\text{def}}{=} \begin{cases} \text{next}_{SYM}, & \text{if} \\ \text{Postdot}(x_{DR}) = \text{next} & \\ \wedge \text{post}_{STR} \stackrel{*}{\Rightarrow} \epsilon & \\ \wedge \neg(\text{next}_{SYM} \stackrel{*}{\Rightarrow} \epsilon) & \\ \Lambda, & \text{otherwise} \end{cases} \end{aligned}$$

A **penult** is a dotted rule d_{DR} such that $\text{Penult}(d) \neq \Lambda$. Note that $\text{Penult}(x_{DR})$ is never a nullable symbol. The **initial dotted rule** is

$$\text{initial}_{DR} = [\text{accept}_{SYM} \rightarrow \bullet \text{start}_{SYM}].$$

A **predicted dotted rule** is a dotted rule, other than the initial dotted rule, with a dot position of zero, for example,

$$\text{predicted}_{DR} = [A_{SYM} \rightarrow \bullet \text{alpha}_{STR}].$$

A **confirmed dotted rule** is the initial dotted rule, or a dotted rule with a dot position greater than zero. A **completed dotted rule** is a dotted rule with its dot position after the end of its RHS, for example,

$$\text{completed}_{DR} = [A_{SYM} \rightarrow \text{alpha}_{STR} \bullet].$$

Predicted, confirmed and completed dotted rules are also called, respectively, **predictions**, **confirmations** and **completions**.

A traditional Earley item (type *EIMT*) is a duple

$$[\text{dotted-rule}_{DR}, x_{ORIG}]$$

of dotted rule and origin. (The origin is the location where recognition of the rule started. It is sometimes called the “parent”.) For convenience, the type *ORIG* will be a synonym for *LOC*, indicating that the variable designates the origin element of an Earley item.

An Earley parser builds a table of Earley sets,

$$\mathbf{table}[\mathbf{EARLEY}, i], \quad \text{where } 0 \leq i_{LOC} \leq |w|.$$

Earley sets are of type *ES*. Earley sets are often named by their location, so that i_{ES} means the Earley set at i_{LOC} . The type designator *ES* is often omitted to avoid clutter, especially in cases where the Earley set is not named by location.

At points, we will need to compare the Earley sets produced by the different recognizers. $\mathbf{table}[\mathbf{RECCE}, i]$ will be the Earley set at i_{LOC} in the table of Earley sets of the RECCE recognizer. For example, $\mathbf{table}[\mathbf{MARPA}, j]$ will be Earley set j_{LOC} in MARPA’s table of Earley sets. In contexts where it is clear which recognizer is intended, $\mathbf{table}[k]$, or k_{ES} , will symbolize Earley set k_{LOC} in that recognizer’s table of Earley sets. If $\mathbf{working}_{ES}$ is an Earley set, $|\mathbf{working}_{ES}|$ is the number of Earley items in $\mathbf{working}_{ES}$.

$|\mathbf{table}[\mathbf{RECCE}]|$ is the total number of Earley items in all Earley sets for RECCE,

$$|\mathbf{table}[\mathbf{RECCE}]| = \sum_{i_{LOC}=0}^{|w|} |\mathbf{table}[\mathbf{RECCE}, i]|.$$

For example, $|\mathbf{table}[\mathbf{MARPA}]|$ is the total number of Earley items in all the Earley sets of a MARPA parse.

Recall that there was a unique acceptance symbol, \mathbf{accept}_{SYM} , in *g*. The input *w* is accepted if and only if, for some \mathbf{rhs}_{STR} ,

$$[[\mathbf{accept}_{SYM} \rightarrow \mathbf{rhs}_{STR} \bullet], 0] \in \mathbf{table}[|w|]$$

5. OPERATIONS OF THE EARLEY ALGORITHM

In this section, each Earley operation is shown in the form of an inference rule, the conclusion of which is the set of the Earley items that is that operation’s **result**. The Earley sets correspond to parse locations, and for any Earley operation there is a current parse location, $\mathbf{current}_{LOC}$, and a current Earley set, $\mathbf{current}_{ES}$.

Each location starts with an empty Earley set. For the purposes of this description of EARLEY, the order of the Earley operations when building an Earley set is non-deterministic. After each Earley operation is performed, its result is unioned with the current Earley set. When

no more Earley items can be added, the Earley set is complete. The Earley sets are built in order from 0 to $|w|$.

5.1. Initialization.

$$\frac{\text{current}_{LOC} = 0}{\{[[\text{accept}_{SYM} \rightarrow \bullet \text{start}_{SYM}], 0]\}}$$

Earley **initialization** has no operands and only takes place in Earley set 0.

5.2. Scanning.

$$\frac{\begin{array}{l} \text{current}_{LOC} > 0 \\ \text{previous}_{LOC} = \text{current}_{LOC} - 1 \\ \text{token}_{SYM} = w[\text{previous}_{LOC}] \\ \text{predecessor}_{EIMT} = [\text{before}_{DR}, \text{predecessor}_{ORIG}] \\ \text{predecessor}_{EIMT} \in \text{previous}_{ES} \\ \text{Postdot}(\text{before}_{DR}) = \text{token}_{SYM} \end{array}}{\{[\text{Next}(\text{before}_{DR}), \text{predecessor}_{ORIG}]\}}$$

$\text{predecessor}_{EIMT}$ and token_{SYM} are the operands of an Earley scan. $\text{predecessor}_{EIMT}$ is called the predecessor of the scanning operation. The token, token_{SYM} , is the transition symbol of the scanning operation.

5.3. Reduction.

$$\frac{\begin{array}{l} \text{component}_{EIMT} = [[\text{lhs}_{SYM} \rightarrow \text{rhs}_{STR} \bullet], \text{component-orig}_{LOC}] \\ \text{component}_{EIMT} \in \text{current}_{ES} \\ \text{predecessor}_{EIMT} = [\text{before}_{DR}, \text{predecessor}_{ORIG}] \\ \text{predecessor}_{EIMT} \in \text{component-orig}_{ES} \\ \text{Postdot}(\text{before}_{DR}) = \text{lhs}_{SYM} \end{array}}{\{[\text{Next}(\text{before}_{DR}), \text{predecessor}_{ORIG}]\}}$$

component_{EIMT} is called the component of the reduction operation.¹ $\text{predecessor}_{EIMT}$ is the predecessor of the reduction operation. lhs_{SYM} is the transition symbol of the reduction operation.

$\text{predecessor}_{EIMT}$ and component_{EIMT} are the operands of the reduction operation. In some contexts, it is convenient to treat the transition symbol as an operand, so that the operands are $\text{predecessor}_{EIMT}$ and lhs_{SYM} .

¹The term “component” comes from Irons [5].

5.4. Prediction.

$$\frac{\begin{array}{c} \text{predecessor}_{EIMT} = [\text{predecessor}_{DR}, \text{predecessor}_{ORIG}] \\ \text{predecessor}_{EIMT} \in \text{current}_{ES} \end{array}}{\left\{ \begin{array}{l} \left[[L_{SYM} \rightarrow \bullet rh_{STR}], \text{current}_{LOC} \right] \text{ such that} \\ [L_{SYM} \rightarrow rh_{STR}] \in \text{rules} \\ \wedge (\exists z_{STR} \mid \text{Postdot}(\text{predecessor}_{DR}) \xrightarrow{*} L_{SYM} \cdot z_{STR}) \end{array} \right\}}$$

A prediction operation can add several Earley items to Earley set current_{LOC} . $\text{predecessor}_{EIMT}$ is called the predecessor of the prediction operation, and is its only operand.

5.5. Causation. An operand of an operation is also called a **cause** of that operation, and the set of operands for an Earley operation is its **causation**.

6. THE LEO ALGORITHM

In [10], Joop Leo presented a method for dealing with right recursion in $\mathcal{O}(n)$ time. Leo shows that, with his modification, Earley’s algorithm is $\mathcal{O}(n)$ for all LR-regular grammars. (LR-regular is LR where look-ahead is infinite length, but restricted to distinguishing between regular expressions.)

Summarizing Leo’s method, it consists of spotting potential right recursions and memoizing them. Leo restricts the memoization to situations where the right recursion is unambiguous. Potential right recursions are memoized by Earley set, using what Leo called “transitive items”. In this paper Leo’s “transitive items” will be called Leo items. Leo items in the form that Marpa uses will be type *LIM*. “Traditional” Leo items, that is, those of the form used in Leo’s paper[10], will be type *LIMT*.

In each Earley set, there is at most one Leo item per symbol. A traditional Leo item (LIMT) is the triple

$$[\text{top}_{DR}, \text{transition}_{SYM}, \text{top}_{ORIG}]$$

where transition_{SYM} is the transition symbol, and

$$\text{top}_{EIMT} = [\text{top}_{DR}, \text{top}_{ORIG}]$$

is the Earley item to be added on reductions over transition_{SYM} .

Leo items memoize what would otherwise be sequences of Earley items. Leo items only memoize unambiguous (or deterministic) sequences, so that the top of the sequence can represent the entire sequence – the only role the other EIMT’s in the sequence play in the

parse is to derive the top EIMT. We will call these memoized sequences, Leo sequences.

To guarantee that a Leo sequence is deterministic, LEO enforced **Leo uniqueness**. Define containment of a dotted rule in a Earley set of EIMT's as

$$\text{Contains}(\mathbf{i}_{ES}, \mathbf{d}_{DR}) \stackrel{\text{def}}{=} \exists \mathbf{b}_{EIMT}, \mathbf{j}_{ORIG} \mid \\ \mathbf{b}_{EIMT} = [\mathbf{d}_{DR}, \mathbf{j}_{ORIG}] \wedge \mathbf{b}_{EIMT} \in \mathbf{i}_{ES}.$$

A dotted rule \mathbf{d}_{DR} is **Leo unique** in the Earley set at \mathbf{i}_{ES} if and only if

$$\text{Penult}(\mathbf{d}_{DR}) \neq \Lambda \\ \wedge \forall \mathbf{d}2_{DR} (\text{Contains}(\mathbf{i}_{ES}, \mathbf{d}2_{DR}) \implies \\ \text{Postdot}(\mathbf{d}_{DR}) = \text{Postdot}(\mathbf{d}2_{DR}) \implies \mathbf{d}_{DR} = \mathbf{d}2_{DR}).$$

If \mathbf{d}_{DR} is Leo unique, then the symbol $\text{Postdot}(\mathbf{d}_{DR})$ is also said to be **Leo unique**. In cases where a symbol transition_{SYM} is Leo unique in \mathbf{i}_{ES} , we can speak of the dotted rule for transition_{SYM} , and the rule for transition_{SYM} , because there can be only one of each. In the previous definitions, it is important to emphasize that $\mathbf{d}2_{DR}$ ranges over all the dotted rules of Earley set \mathbf{i}_{ES} , even those which are ineligible for Leo memoization.

Let \mathbf{n} be the length of a Leo sequence. In EARLEY, each such sequence would be expanded in every Earley set that is the origin of an EIMT included in the sequence, and the total number of EIMT's would be $\mathcal{O}(\mathbf{n}^2)$.

With Leo memoization, a single EIMT stands in for the sequence. There are $\mathcal{O}(1)$ Leo items per Earley set, so the cost of the sequence is $\mathcal{O}(1)$ per Earley set, or $\mathcal{O}(\mathbf{n})$ for the entire sequence. If, at evaluation time, it is desirable to expand the Leo sequence, only those items actually involved in the parse need to be expanded. All the EIMT's of a potential right-recursion will be in one Earley set and the number of EIMT's will be $\mathcal{O}(\mathbf{n})$, so that even including expansion of the Leo sequence for evaluation, the time and space complexity of the sequence remains $\mathcal{O}(\mathbf{n})$.

Recall that we call a dotted rule \mathbf{d}_{DR} a **penult** if $\text{Penult}(\mathbf{d}) \neq \Lambda$. In Leo's original algorithm, any penult was treated as a potential right-recursion. MARPA applies the Leo memoizations in more restricted circumstances. For MARPA to consider a dotted rule

$$\text{candidate}_{DR} = [\text{candidate}_{RULE}, \mathbf{i}]$$

for Leo memoization, candidate_{DR} must be a penult and candidate_{RULE} must be right-recursive.

By restricting Leo memoization to right-recursive rules, MARPA incurs the cost of Leo memoization only in cases where Leo sequences could be infinitely long. This more careful targeting of the memoization is for efficiency reasons. If all penults are memoized, many memoizations will be performed where the longest potential Leo sequence is short, and the payoff is therefore very limited. One future extension might be to identify non-right-recursive rules which generate Leo sequences long enough to justify inclusion in the Leo memoizations. Such cases are unusual, but may occur.

Omission of a memoization does not affect correctness, so MARPA's restriction of Leo memoization preserves the correctness as shown in Leo[10]. Later in this paper we will show that this change also leaves the complexity results of Leo[10] intact.

Implementing the Leo logic requires adding Leo reduction as a new basic operation, adding a new premise to the Earley reduction operation, and extending the Earley sets to memoize Earley items as LIMT's.

6.1. Leo reduction.

$$\begin{array}{c} \text{component}_{EIMT} = [[\text{lhs}_{SYM} \rightarrow \text{rhs}_{STR} \bullet], \text{component-orig}_{LOC}] \\ \text{component}_{EIMT} \in \text{current}_{ES} \\ \text{predecessor}_{LIMT} = [\text{top}_{DR}, \text{lhs}_{SYM}, \text{top}_{ORIG}] \\ \text{predecessor}_{LIMT} \in \text{component-orig}_{ES} \\ \hline \{[\text{top}_{DR}, \text{top}_{ORIG}]\} \end{array}$$

The new Leo reduction operation resembles the Earley reduction operation, except that it looks for an LIMT, instead of a predecessor EIMT. $\text{predecessor}_{LIMT}$ and component_{EIMT} are the operands of the Leo reduction operation. lhs_{SYM} is the transition symbol of the Leo reduction. As with Earley reduction, it may be convenient to treat the transition symbol as an operand, so that the operands are $\text{predecessor}_{LIMT}$ and lhs_{SYM} .

6.2. Changes to Earley reduction. Earley reduction still applies, with an additional premise:

$$\neg \exists \mathbf{x}_{LIMT} \mid \mathbf{x}_{LIMT} \in \text{component-orig}_{ES} \\ \wedge \mathbf{x}_{LIMT} = [\mathbf{x}_{DR}, \text{lhs}_{SYM}, \mathbf{x}_{ORIG}]$$

The additional premise prevents Earley reduction from being applied where there is an LIMT with lhs_{SYM} as its transition symbol. This reflects the fact that Leo reduction replaces Earley reduction if and only if there is a Leo memoization.

6.3. Leo memoization. We define uniqueness of a penult in an Earley set as

$$\begin{aligned} \text{Penult-Unique}(\text{penult}_{SYM}, i_{ES}) &\stackrel{\text{def}}{=} \\ &\forall y_{DR} (\text{Contains}(\text{current}_{ES}, y_{DR}) \wedge \text{penult}_{SYM} = \text{Penult}(y_{DR})) \\ &\implies x_{DR} = y_{DR}. \end{aligned}$$

We define Leo uniqueness as

$$\begin{aligned} \text{Leo-Unique}(x_{DR}, \text{current}_{LOC}) &\stackrel{\text{def}}{=} \text{Contains}(\text{current}_{ES}, x_{DR}) \\ &\wedge \text{Penult}(x_{DR}) \neq \Lambda \\ &\wedge \text{Penult-Unique}(\text{Penult}(x_{DR}), \text{current}_{ES}) \end{aligned}$$

and Leo eligibility as

$$\begin{aligned} \text{Leo-Eligible}(x_{DR}, \text{current}_{LOC}) &\stackrel{\text{def}}{=} \\ &\exists x_{RULE}, x_{ORIG} \mid x_{DR} = [x_{RULE}, i_{ORIG}] \\ &\wedge \text{Right-Recursive}(x_{RULE}) \\ &\wedge \text{Leo-Unique}(\text{current}_{ES}, x_{DR}). \end{aligned}$$

For convenience, we define a relation that is true if pred_{LIMT} is the LIMT predecessor of an EIMT, and false otherwise:

$$\begin{aligned} \text{LIMT-Predecessor}(\text{pred}_{LIMT}, \text{bottom}_{EIMT}) &\stackrel{\text{def}}{=} \\ &\exists \text{bottom-orig}_{ES}, \text{bottom}_{DR}, \text{pred}_{DR}, \\ &\quad \text{pred-orig}_{LOC}, \text{bottom-orig}_{LOC} \text{ such that} \\ &\quad \text{bottom}_{EIMT} = [\text{bottom}_{DR}, \text{bottom-orig}_{LOC}] \\ &\quad \wedge \text{pred}_{LIMT} = [\text{pred}_{DR}, \text{LHS}(\text{bottom}_{DR}), \text{pred-orig}_{LOC}] \\ &\quad \wedge \text{pred}_{LIMT} \in \text{bottom-orig}_{ES} \end{aligned}$$

We are now ready to define an inference rule which holds if a LIMT predecessor can be found for an EIMT bottom_{EIMT}

$$\frac{\begin{aligned} &\text{LIMT-Predecessor}(\text{pred}_{LIMT}, \text{bottom}_{EIMT}) \\ &\text{pred}_{LIMT} = [\text{pred}_{DR}, \text{LHS}(\text{bottom}_{DR}), \text{pred}_{ORIG}] \\ &\text{bottom}_{EIMT} = [\text{bottom}_{DR}, \text{bottom}_{LOC}] \\ &\text{bottom}_{EIMT} \in \text{current}_{ES} \\ &\text{Leo-Eligible}(\text{bottom}_{DR}) \end{aligned}}{\{[\text{pred}_{DR}, \text{Penult}(\text{bottom}_{DR}), \text{pred}_{ORIG}]\}}$$

and another, which holds if bottom_{EIMT} has no predecessor LIMT,

$$\frac{\begin{array}{l} \neg \text{LIMT-Predecessor}(\text{pred}_{LIMT}, \text{bottom}_{EIMT}) \\ \text{bottom}_{EIMT} = [\text{bottom}_{DR}, \text{bottom}_{ORIG}] \\ \text{bottom}_{EIMT} \in \text{current}_{ES} \\ \text{Leo-Eligible}(\text{bottom}_{DR}) \end{array}}{\{[\text{Next}(\text{bottom}_{DR}), \text{Penult}(\text{bottom}_{DR}), \text{bottom}_{ORIG}]\}}$$

7. THE AYCOCK-HORSPPOOL FINITE AUTOMATON

In this paper a “split LR(0) ϵ -DFA” as described by Aycock and Horspool[2], will be called an Aycock-Horspool Finite Automaton, or AHFA. This section will summarize the ideas from [2] that are central to Marpa.

Aycock and Horspool based their AHFA’s on a few observations.

- In practice, Earley items sharing the same origin, but having different dotted rules, often appear together in the same Earley set.
- There is in the literature a method for associating groups of dotted rules that often appear together when parsing. This method is the LR(0) DFA used in the much-studied LALR and LR parsers.
- The LR(0) items that are the components of LR(0) states are, exactly, dotted rules.
- By taking into account symbols that derive the null string, the LR(0) DFA could be turned into an LR(0) ϵ -DFA, which would be even more effective at grouping dotted rules that often occur together into a single DFA state.

AHFA states are, in effect, a shorthand for groups of dotted rules that occur together frequently. Aycock and Horspool realized that, by changing Earley items to track AHFA states instead of individual dotted rules, the size of Earley sets could be reduced, and Earley’s algorithm made faster in practice.

As a reminder, the original Earley items (EIMT’s) were duples, $[\mathbf{x}_{DR}, \mathbf{x}_{ORIG}]$, where \mathbf{x}_{DR} is a dotted rule. An Aycock-Horspool Earley item is a duple

$$[\mathbf{y}_{AH}, \mathbf{y}_{ORIG}],$$

where \mathbf{y}_{AH} is an AHFA state.

MARPA uses Earley items of the form created by Aycock and Horspool. A Marpa Earley item has type *EIM*, and a Marpa Earley item is often referred to as an EIM.

Aycock and Horspool did not consider Leo’s modifications, but MARPA incorporates them, and MARPA also changes its Leo items to use AHFA states. Marpa’s Leo items (LIM’s) are triples of the form

$$[\text{top}_{AH}, \text{transition}_{SYM}, \text{top}_{ORIG}],$$

where transition_{SYM} and top_{ORIG} are as in the traditional Leo items, and top_{AH} is an AHFA state. A Marpa Leo item has type *LIM*.

[2] also defines a partial transition function for pairs of AHFA state and symbol,

$$\text{GOTO} : \text{fa}, (\epsilon \cup \text{vocab}) \mapsto \text{fa}.$$

$\text{GOTO}(\text{from}_{AH}, \epsilon)$ is a **null transition**. (AHFA’s are not fully deterministic.) If predicted_{AH} is the result of a null transition, it is called a **predicted** AHFA state. If an AHFA state is not a **predicted** AHFA state, it is called a **confirmed** AHFA state. The initial AHFA state is a confirmed AHFA state. (In [2] confirmed states are called “kernel states”, and predicted states are called “non-kernel states”.)

The states of an AHFA are not a partition of the dotted rules – a single dotted rule can occur in more than one AHFA state. In combining the improvements of Leo [10] and Aycock and Horspool[2], the following theorem is crucial.

Theorem 7.1. *If a Marpa Earley item (EIM) is the result of a Leo reduction, then its AHFA state contains only one dotted rule.*

Proof. Let the EIM that is the result of the Leo reduction be

$$\text{result}_{EIM} = [\text{result}_{AH}, \text{result}_{ORIG}]$$

Let the Earley set that contains result_{EIM} be i_{ES} . Since result_{EIM} is the result of a Leo reduction we know, from the definition of a Leo reduction, that

$$\text{complete}_{DR} \in \text{result}_{AH}$$

where complete_{DR} is a completed rule. Let c_{RULE} be the rule of complete_{DR} , and cp its dot position,

$$\text{complete}_{DR} = [c_{RULE}, cp].$$

$cp > 0$ because, in MARPA grammars, completions are never predictions.

Suppose, for a reduction to absurdity, that the AHFA state contains another dotted rule, other_{DR} , that is, that

$$\text{other}_{DR} \in \text{result}_{AH},$$

where $\mathbf{complete}_{DR} \neq \mathbf{other}_{DR}$. Let \mathbf{o}_{RULE} be the rule of \mathbf{other}_{DR} , and \mathbf{op} its dot position,

$$\mathbf{other}_{DR} = [\mathbf{o}_{RULE}, \mathbf{op}].$$

AHFA construction never places a prediction in the same AHFA state as a completion, so \mathbf{other}_{DR} is not a prediction. Therefore, $\mathbf{op} > 0$. To show this outer reduction to absurdity, we first prove by a first inner reductio that $\mathbf{c}_{RULE} \neq \mathbf{o}_{RULE}$, then by a second inner reductio that $\mathbf{c}_{RULE} = \mathbf{o}_{RULE}$.

Assume, for the first inner reductio, that $\mathbf{c}_{RULE} = \mathbf{o}_{RULE}$. By the construction of an AHFA state, both $\mathbf{complete}_{DR}$ and \mathbf{other}_{DR} resulted from the same series of transitions. But the same series of transitions over the same rule would result in the same dot position, $\mathbf{cp} = \mathbf{op}$, so that if $\mathbf{c}_{RULE} = \mathbf{o}_{RULE}$, $\mathbf{complete}_{DR} = \mathbf{other}_{DR}$, which is contrary to the assumption for the outer reductio. This shows the first inner reductio.

Next, we assume for the second inner reductio that $\mathbf{c}_{RULE} \neq \mathbf{o}_{RULE}$. Since both $\mathbf{complete}_{DR}$ and \mathbf{other}_{DR} are in the same EIM and neither is a prediction, both must result from transitions, and their transitions must have been from the same Earley set. Since they are in the same AHFA state, by the AHFA construction, that transition must have been over the same transition symbol, call it $\mathbf{transition}_{SYM}$. But Leo uniqueness applies to $\mathbf{complete}_{DR}$, and requires that the transition over $\mathbf{transition}_{SYM}$ be unique in \mathbf{i}_{ES} .

But if $\mathbf{c}_{RULE} \neq \mathbf{o}_{RULE}$, $\mathbf{transition}_{SYM}$ was the transition symbol of two different dotted rules, and the Leo uniqueness requirement does not hold. The conclusion that the Leo uniqueness requirement both does and does not hold is a contradiction, which shows the second inner reductio. Since the assumption for the second inner reductio was that $\mathbf{c}_{RULE} \neq \mathbf{o}_{RULE}$, we conclude that $\mathbf{c}_{RULE} = \mathbf{o}_{RULE}$.

By the two inner reductio's, we have both $\mathbf{c}_{RULE} \neq \mathbf{o}_{RULE}$ and $\mathbf{c}_{RULE} = \mathbf{o}_{RULE}$, which completes the outer reduction to absurdity. For the outer reductio, we assumed that \mathbf{other}_{DR} was a second dotted rule in \mathbf{result}_{AH} , such that $\mathbf{other}_{DR} \neq \mathbf{complete}_{DR}$. We can therefore conclude that

$$\mathbf{other}_{DR} \in \mathbf{result}_{AH} \implies \mathbf{other}_{DR} = \mathbf{complete}_{DR}.$$

If $\mathbf{complete}_{DR}$ is a dotted rule in the AHFA state of a Leo reduction EIM, then it must be the only dotted rule in that AHFA state. \square

8. THE MARPA RECOGNIZER

8.1. **Complexity.** Alongside the pseudocode of this section are observations about its space and time complexity. In what follows, we will charge all time and space resources to Earley items, or to attempts to add Earley items. We will show that, to each Earley item actually added, or to each attempt to add a duplicate Earley item, we can charge amortized $\mathcal{O}(1)$ time and space.

At points, it will not be immediately convenient to speak of charging a resource to an Earley item or to an attempt to add a duplicate Earley item. In those circumstances, we speak of charging time and space

- to the parse; or
- to the Earley set; or
- to the current procedure’s caller.

We can charge time and space to the parse itself, as long as the total time and space charged is $\mathcal{O}(1)$. Afterwards, this resource can be re-charged to the initial Earley item, which is present in all parses. Soft and hard failures of the recognizer use worst-case $\mathcal{O}(1)$ resource, and are charged to the parse.

We can charge resources to the Earley set, as long as the time or space is $\mathcal{O}(1)$. Afterwards, the resource charged to the Earley set can be re-charged to an arbitrary member of the Earley set, for example, the first. If an Earley set is empty, the parse must fail, and the time can be charged to the parse.

In a procedure, resource can be “caller-included”. Caller-included resource is not accounted for in the current procedure, but passed upward to the procedure’s caller, to be accounted for there. A procedure to which caller-included resource is passed will sometimes pass the resource upward to its own caller, although of course the top-level procedure does not do this.

For each procedure, we will state whether the time and space we are charging is inclusive or exclusive. The exclusive time or space of a procedure is that which it uses directly, ignoring resource charges passed up from called procedures. Inclusive time or space includes resource passed upward to the current procedure from called procedures.

Earley sets may be represented by i_{ES} , where i is the Earley set’s location i_{LOC} . The two notations should be regarded as interchangeable. The actual implementation of either should be the equivalent of a pointer to a data structure containing, at a minimum, the Earley items, a memoization of the Earley set’s location as an integer, and a per-set-list. Per-set-lists will be described in Section 8.12.

Algorithm 1 Marpa Top-level

```

1: procedure MAIN
2:   INITIAL
3:   for  $i, 0 \leq i \leq |w|$  do
4:      $\triangleright$  At this point,  $x_{ES}$  is complete, for  $0 \leq x < i$ 
5:     SCAN PASS( $i, w[i - 1]$ )
6:     if  $|i_{ES}| = 0$  then
7:       reject  $w$  and return
8:     end if
9:     REDUCTION PASS( $i$ )
10:  end for
11:  for every  $[x_{AH}, 0] \in \text{table}[|w|]$  do
12:    if  $\text{accept}_{DR} \in x_{AH}$  then
13:      accept  $w$  and return
14:    end if
15:  end for
16:  reject  $w$ 
17: end procedure

```

8.2. **Top-level code.** Exclusive time and space for the loop over the Earley sets is charged to the Earley sets. Inclusive time and space for the final loop to check for accept_{DR} is charged to the Earley items at location $|w|$. Overhead is charged to the parse. All these resource charges are obviously $\mathcal{O}(1)$.

8.3. **Ruby Slippers parsing.** This top-level code represents a significant change from AH. SCAN PASS and REDUCTION PASS are separated. As a result, when the scanning of tokens that start at location i_{LOC} begins, the Earley sets for all locations prior to i_{LOC} are complete. This means that the scanning operation has available, in the Earley sets, full information about the current state of the parse, including which tokens are acceptable during the scanning phase.

Algorithm 2 Initialization

```

1: procedure INITIAL
2:   ADD EIM PAIR( $0_{ES}, \text{start}_{AH}, 0$ )
3: end procedure

```

8.4. **Initialization.** Inclusive time and space is $\mathcal{O}(1)$ and is charged to the parse.

Algorithm 3 Marpa Scan pass

```

1: procedure SCAN PASS( $i_{LOC}$ ,  $a_{SYM}$ )
2:   Note: Each pass through this loop is an EIM attempt
3:   for each predecessor $_{EIM} \in \text{transitions}((i - 1), a)$  do
4:     [ $\text{from}_{AH}$ ,  $\text{origin}_{LOC}$ ]  $\leftarrow$  predecessor $_{EIM}$ 
5:      $\text{to}_{AH} \leftarrow \text{GOTO}(\text{from}_{AH}, a_{SYM})$ 
6:     ADD EIM PAIR( $i_{ES}$ ,  $\text{to}_{AH}$ ,  $\text{origin}_{LOC}$ )
7:   end for
8: end procedure

```

8.5. **Scan pass.** `transitions` is a set of tables, one per Earley set. The tables in the set are indexed by symbol. Symbol indexing is $\mathcal{O}(1)$, since the number of symbols is a constant, but since the number of Earley sets grows with the length of the parse, it cannot be assumed that Earley sets can be indexed by location in $\mathcal{O}(1)$ time. For the operation `transitions(l_{LOC} , s_{SYM})` to be in $\mathcal{O}(1)$ time, l_{LOC} must represent a link directly to the Earley set. In the case of scanning, the lookup is always in the previous Earley set, which can easily be tracked in $\mathcal{O}(1)$ space and retrieved in $\mathcal{O}(1)$ time. Inclusive time and space can be charged to the predecessor $_{EIM}$. Overhead is charged to the Earley set at i_{LOC} .

Algorithm 4 Reduction pass

```

1: procedure REDUCTION PASS( $i_{LOC}$ )
2:   Note: table[i] may include EIM's added by
3:         by REDUCE ONE LHS and
4:         the loop must traverse these
5:   for each Earley item  $\text{work}_{EIM} \in \text{table}[i]$  do
6:     [ $\text{work}_{AH}$ ,  $\text{origin}_{LOC}$ ]  $\leftarrow$   $\text{work}_{EIM}$ 
7:      $\text{lhs\_sides}_{\{SYM\}} \leftarrow$  a set containing the LHS
8:       of every completed rule in  $\text{work}_{EIM}$ 
9:     for each  $\text{lhs}_{SYM} \in \text{lhs\_sides}_{\{SYM\}}$  do
10:      REDUCE ONE LHS( $i_{LOC}$ ,  $\text{origin}_{LOC}$ ,  $\text{lhs}_{SYM}$ )
11:    end for
12:  end for
13:  MEMOIZE TRANSITIONS( $i_{LOC}$ )
14: end procedure

```

8.6. **Reduction pass.** The loop over `table[i]` must also include any items added by REDUCE ONE LHS. This can be done by implementing `table[i]` as an ordered set and adding new items at the end.

Exclusive time is clearly $\mathcal{O}(1)$ per work_{EIM} , and is charged to the work_{EIM} . Additionally, some of the time required by REDUCE ONE LHS is caller-included, and therefore charged to this procedure. Inclusive time from REDUCE ONE LHS is $\mathcal{O}(1)$ per call, as will be seen in section 8.8, and is charged to the work_{EIM} that is current during that call to REDUCE ONE LHS. Overhead may be charged to the Earley set at i_{LOC} .

Algorithm 5 Memoize transitions

```

1: procedure MEMOIZE TRANSITIONS( $i_{LOC}$ )
2:   for every  $\text{postdot}_{SYM}$ , a postdot symbol of  $i_{ES}$  do
3:     Note:  $\text{postdot}_{SYM}$  is “Leo eligible” if it is
4:       Leo unique and its rule is right recursive
5:     if  $\text{postdot}_{SYM}$  is Leo eligible then
6:       Set  $\text{transitions}(i_{LOC}, \text{postdot}_{SYM})$ 
7:         to a LIM
8:     else
9:       Set  $\text{transitions}(i_{LOC}, \text{postdot}_{SYM})$ 
10:        to the set of EIM’s that have
11:         $\text{postdot}_{SYM}$  as their postdot symbol
12:     end if
13:   end for
14: end procedure

```

8.7. Memoize transitions. The transitions table for i_{ES} is built once all EIMs have been added to i_{ES} . We first look at the resource, excluding the processing of Leo items. The non-Leo processing can be done in a single pass over i_{ES} , in $\mathcal{O}(1)$ time per EIM. Inclusive time and space are charged to the Earley items being examined. Overhead is charged to i_{ES} .

We now look at the resource used in the Leo processing. A transition symbol transition_{SYM} is Leo eligible if it is Leo unique and its rule is right recursive. (If transition_{SYM} is Leo unique in i_{ES} , it will be the postdot symbol of only one rule in i_{ES} .) All but one of the determinations needed to decide if transition_{SYM} is Leo eligible can be precomputed from the grammar, and the resource to do this is charged to the parse. The precomputation, for example, for every rule, determines if it is right recursive.

One part of the test for Leo eligibility cannot be done as a precomputation. This is the determination whether there is only one dotted rule in i_{ES} whose postdot symbol is transition_{SYM} . This can be done in

a single pass over the EIM's of i_{ES} that notes the postdot symbols as they are encountered and whether any is encountered twice. The time and space, including that for the creation of a LIM if necessary, will be $\mathcal{O}(1)$ time per EIM examined, and can be charged to EIM being examined.

Algorithm 6 Reduce one LHS symbol

```

1: procedure REDUCE ONE LHS( $i_{LOC}$ ,  $origin_{LOC}$ ,  $lhs_{SYM}$ )
2:   Note: Each pass through this loop is an EIM attempt
3:   for each  $pim \in transitions(origin_{LOC}, lhs_{SYM})$  do
4:      $\triangleright pim$  is a “postdot item”, either a LIM or an EIM
5:     if  $pim$  is a LIM,  $pim_{LIM}$  then
6:       Perform a LEO REDUCTION OPERATION
7:       for operands  $i_{LOC}$ ,  $pim_{LIM}$ 
8:     else
9:       Perform a EARLEY REDUCTION OPERATION
10:      for operands  $i_{LOC}$ ,  $pim_{EIM}$ ,  $lhs_{SYM}$ 
11:    end if
12:  end for
13: end procedure

```

8.8. **Reduce one LHS.** To show that

$$transitions(origin_{LOC}, lhs_{SYM})$$

can be traversed in $\mathcal{O}(1)$ time, we note that the number of symbols is a constant and assume that $origin_{LOC}$ is implemented as a link back to the Earley set, rather than as an integer index. This requires that $work_{EIM}$ in REDUCTION PASS carry a link back to its origin. As implemented, Marpa's Earley items have such links.

Inclusive time for the loop over the EIM attempts is charged to each EIM attempt. Overhead is $\mathcal{O}(1)$ and caller-included.

Algorithm 7 Earley reduction operation

```

1: procedure EARLEY REDUCTION OPERATION( $i_{LOC}$ ,  $from_{EIM}$ ,
    $trans_{SYM}$ )
2:   [ $from_{AH}$ ,  $origin_{LOC}$ ]  $\leftarrow from_{EIM}$ 
3:    $to_{AH} \leftarrow GOTO(from_{AH}, trans_{SYM})$ 
4:   ADD EIM PAIR( $i_{ES}$ ,  $to_{AH}$ ,  $origin_{LOC}$ )
5: end procedure

```

8.9. **Earley Reduction operation.** Exclusive time and space is clearly $\mathcal{O}(1)$. EARLEY REDUCTION OPERATION is always called as part of an EIM attempt, and inclusive time and space is charged to the EIM attempt.

Algorithm 8 Leo reduction operation

```

1: procedure LEO REDUCTION OPERATION( $i_{LOC}$ ,  $from_{LIM}$ )
2:   [ $from_{AH}$ ,  $trans_{SYM}$ ,  $origin_{LOC}$ ]  $\leftarrow$   $from_{LIM}$ 
3:    $to_{AH} \leftarrow$  GOTO( $from_{AH}$ ,  $trans_{SYM}$ )
4:   ADD EIM PAIR( $i_{ES}$ ,  $to_{AH}$ ,  $origin_{LOC}$ )
5: end procedure

```

8.10. **Leo reduction operation.** Exclusive time and space is clearly $\mathcal{O}(1)$. LEO REDUCTION OPERATION is always called as part of an EIM attempt, and inclusive time and space is charged to the EIM attempt.

Algorithm 9 Add EIM pair

```

1: procedure ADD EIM PAIR( $i_{ES}$ ,  $confirmed_{AH}$ ,  $origin_{LOC}$ )
2:    $confirmed_{EIM} \leftarrow$  [ $confirmed_{AH}$ ,  $origin_{LOC}$ ]
3:    $predicted_{AH} \leftarrow$  GOTO( $confirmed_{AH}$ ,  $\epsilon$ )
4:   if  $confirmed_{EIM}$  is new then
5:     Add  $confirmed_{EIM}$  to table[i]
6:   end if
7:   if  $predicted_{AH} \neq \Lambda$  then
8:      $predicted_{EIM} \leftarrow$  [ $predicted_{AH}$ ,  $i_{LOC}$ ]
9:     if  $predicted_{EIM}$  is new then
10:      Add  $predicted_{EIM}$  to table[i]
11:    end if
12:   end if
13: end procedure

```

8.11. **Adding a pair of Earley items.** This operation adds a confirmed EIM item and, if it exists, the EIM for its null-transition. Inclusive time and space is charged to the calling procedure. Trivially, the space is $\mathcal{O}(1)$ per call.

We show that time is also $\mathcal{O}(1)$ by singling out the two non-trivial cases: checking that an Earley item is new, and adding it to the Earley set. MARPA checks whether an Earley item is new in $\mathcal{O}(1)$ time by using a data structure called a PSL. PSL's are the subject of Section 8.12. An Earley item can be added to the current set in $\mathcal{O}(1)$ time if

Earley set is seen as a linked list, to the head of which the new Earley item is added.

The resource used by ADD EIM PAIR is always caller-included. No time or space is ever charged to a predicted Earley item. At most one attempt to add a `predictedEIM` will be made per attempt to add a `confirmedEIM`, so that the total resource charged remains $\mathcal{O}(1)$.

8.12. Per-set lists. In the general case, where \mathbf{x} is an arbitrary datum, it is not possible to use duple $[\mathbf{i}_{ES}, x]$ as a search key and expect the search to use $\mathcal{O}(1)$ time. Within MARPA, however, there are specific cases where it is desirable to do exactly that. This is accomplished by taking advantage of special properties of the search.

If it can be arranged that there is a link direct to the Earley set \mathbf{i}_{ES} , and that $0 \leq \mathbf{x} < \mathbf{c}$, where \mathbf{c} is a constant of reasonable size, then a search can be made in $\mathcal{O}(1)$ time, using a data structure called a PSL. Data structures identical to or very similar to PSL's are briefly outlined in both [3, p. 97] and [1, Vol. 1, pages 326-327]. But neither source gives them a name. The term PSL ("per-Earley set list") is new with this paper.

A PSL is a fixed-length array of integers, indexed by an integer, and kept as part of each Earley set. While MARPA is building a new Earley set, \mathbf{j}_{ES} , the PSL for every previous Earley set, \mathbf{i}_{LOC} , tracks the Earley items in \mathbf{j}_{ES} that have \mathbf{i}_{LOC} as their origin. The maximum number of Earley items that must be tracked in each PSL is the number of AHFA states, $|\mathbf{fa}|$, which is a constant of reasonable size that depends on \mathbf{g} .

It would take more than $\mathcal{O}(1)$ time to clear and rebuild the PSL's each time that a new Earley set is started. This overhead is avoided by "time-stamping" each PSL entry with the Earley set that was current when that PSL entry was last updated.

As before, where \mathbf{i}_{ES} is an Earley set, let \mathbf{i}_{LOC} be its location, and vice versa. \mathbf{i}_{LOC} is an integer which is assigned as Earley sets are created. Let $\text{ID}(\mathbf{x}_{AH})$ be the integer ID of an AHFA state. Numbering the AHFA states from 0 on up as they are created is an easy way to create $\text{ID}(\mathbf{x}_{AH})$. Let $\text{PSL}[\mathbf{x}_{ES}][y]$ be the entry for integer y in the PSL in the Earley set at \mathbf{x}_{LOC} .

Consider the case where Marpa is building \mathbf{j}_{ES} and wants to check whether Earley item \mathbf{x}_{EIM} is new, where $\mathbf{x}_{EIM} = [\mathbf{x}_{AH}, \mathbf{x}_{ORIG}]$. To check if \mathbf{x}_{EIM} is new, Marpa checks

$$\text{time-stamp} = \text{PSL}[\mathbf{x}_{ES}][\text{ID}(\mathbf{x}_{AH})]$$

If the entry has never been used, we assume that $\mathbf{time-stamp} = \Lambda$. If $\mathbf{time-stamp} \neq \Lambda \wedge \mathbf{time-stamp} = j_{LOC}$, then \mathbf{x}_{EIM} is not new, and will not be added to the Earley set.

If $\mathbf{p}_{LOC} = \Lambda \vee \mathbf{time-stamp} \neq j_{LOC}$, then \mathbf{x}_{EIM} is new. \mathbf{x}_{EIM} is added to the Earley set, and a new time-stamp is set, as follow:

$$\text{PSL}[\mathbf{x}_{ES}][\text{ID}(\mathbf{x}_{AH})] \leftarrow j_{LOC}.$$

8.13. Complexity summary. For convenience, we collect and summarize here some of the observations of this section.

Observation 8.1. The time and space charged to an Earley item which is actually added to the Earley sets is $\mathcal{O}(1)$.

Observation 8.2. The time charged to an attempt to add a duplicate Earley item to the Earley sets is $\mathcal{O}(1)$.

For evaluation purposes, MARPA adds a link to each EIM that records each attempt to add that EIM, whether originally or as a duplicate. Traditionally, complexity results treat parsers as recognizers, and such costs are ignored. This will be an issue when the space complexity for unambiguous grammars is considered.

Observation 8.3. The space charged to an attempt to add a duplicate Earley item to the Earley sets is $\mathcal{O}(1)$ if links are included, zero otherwise.

As noted in Section 8.11, the time and space used by predicted Earley items and attempts to add them is charged elsewhere.

Observation 8.4. No space or time is charged to predicted Earley items, or to attempts to add predicted Earley items.

9. PRELIMINARIES TO THE THEORETICAL RESULTS

9.1. Nulling symbols. Recall that Marpa grammars, without loss of generality, contain neither empty rules or properly nullable symbols. This corresponds directly to a grammar rewrite in the MARPA implementation, and its reversal during MARPA's evaluation phase. For the correctness and complexity proofs in this paper, we assume an additional rewrite, this time to eliminate nulling symbols.

Elimination of nulling symbols is also without loss of generality, as can be seen if we assume that a history of the rewrite is kept, and that the rewrite is reversed after the parse. Clearly, whether a grammar g accepts an input w will not depend on the nulling symbols in its rules.

In its implementation, MARPA does not directly rewrite the grammar to eliminate nulling symbols. But nulling symbols are ignored in

creating the AHFA states, and must be restored during MARPA's evaluation phase, so that the implementation and this simplification for theory purposes track each other closely.

9.2. Comparing Earley items.

Definition. A Marpa Earley item **corresponds** to a traditional Earley item $\mathbf{x}_{EIMT} = [\mathbf{x}_{DR}, \mathbf{x}_{ORIG}]$ if and only if the Marpa Earley item is a $\mathbf{y}_{EIM} = [\mathbf{y}_{AH}, \mathbf{x}_{ORIG}]$ such that $\mathbf{x}_{DR} \in \mathbf{y}_{AH}$. A traditional Earley item, \mathbf{x}_{EIMT} , corresponds to a Marpa Earley item, \mathbf{y}_{EIM} , if and only if \mathbf{y}_{EIM} corresponds to \mathbf{x}_{EIMT} .

Definition. A set of EIM's is **consistent** with respect to a set of EIMT's, if and only if each of the EIM's in the first set corresponds to at least one of the EIMT's in the second set. A Marpa Earley set $\mathbf{table}[\text{MARPA}, i]$ is **consistent** if and only if all of its EIM's correspond to EIMT's in $\mathbf{table}[\text{LEO}, i]$.

Definition. A set of EIM's is **complete** with respect to a set of EIMT's, if and only if for every EIMT in the second set, there is a corresponding EIM in the first set. A Marpa Earley set $\mathbf{table}[\text{MARPA}, i]$ is **complete** if and only if for every traditional Earley item in $\mathbf{table}[\text{LEO}, i]$ there is a corresponding Earley item in $\mathbf{table}[\text{MARPA}, i]$.

Definition. A Marpa Earley set is **correct** if and only that Marpa Earley set is complete and consistent.

9.3. **About AHFA states.** Several facts from [2] will be heavily used in the following proofs. For convenience, they are restated here.

Observation 9.1. Every dotted rule is an element of one or more AHFA states, that is,

$$\forall \mathbf{x}_{DR} \exists \mathbf{y}_{AH} \mid \mathbf{x}_{DR} \in \mathbf{y}_{AH}.$$

Observation 9.2. AHFA confirmation is consistent with respect to the dotted rules. That is, for all \mathbf{from}_{AH} , \mathbf{t}_{SYM} , \mathbf{to}_{AH} , \mathbf{to}_{DR} such that

$$\begin{aligned} & \mathbf{GOTO}(\mathbf{from}_{AH}, \mathbf{t}_{SYM}) = \mathbf{to}_{AH} \\ \wedge & \quad \mathbf{to}_{DR} \in \mathbf{to}_{AH}, \end{aligned}$$

there exists \mathbf{from}_{DR} such that

$$\begin{aligned} & \mathbf{from}_{DR} \in \mathbf{from}_{AH} \\ \wedge & \quad \mathbf{t}_{SYM} = \mathbf{Postdot}(\mathbf{from}_{DR}) \\ \wedge & \quad \mathbf{Next}(\mathbf{from}_{DR}) = \mathbf{to}_{DR}. \end{aligned}$$

Observation 9.3. AHFA confirmation is complete with respect to the dotted rules. That is, for all from_{AH} , t_{SYM} , from_{DR} , to_{DR} if

$$\begin{aligned} & \text{from}_{DR} \in \text{from}_{AH} \\ \wedge & \text{Postdot}(\text{from}_{DR}) = \text{t}_{SYM}, \\ \wedge & \text{Next}(\text{from}_{DR}) = \text{to}_{DR} \end{aligned}$$

then there exists to_{AH} such that

$$\begin{aligned} & \text{GOTO}(\text{from}_{AH}, \text{t}_{SYM}) = \text{to}_{AH} \\ \wedge & \text{to}_{DR} \in \text{to}_{AH}. \end{aligned}$$

Observation 9.4. AHFA prediction is consistent with respect to the dotted rules. That is, for all from_{AH} , to_{AH} , to_{DR} such that

$$\text{GOTO}(\text{from}_{AH}, \epsilon) = \text{to}_{AH} \wedge \text{to}_{DR} \in \text{to}_{AH},$$

there exists from_{DR} such that

$$\text{from}_{DR} \in \text{from}_{AH} \wedge \text{to}_{DR} \in \text{Predict}(\text{from}_{DR}).$$

Observation 9.5. AHFA prediction is complete with respect to the dotted rules. That is, for all from_{AH} , from_{DR} , to_{DR} , if

$$\text{from}_{DR} \in \text{from}_{AH} \wedge \text{to}_{DR} \in \text{Predict}(\text{from}_{DR}),$$

then there exists to_{AH} such that

$$\text{to}_{DR} \in \text{to}_{AH} \wedge \text{GOTO}(\text{from}_{AH}, \epsilon) = \text{to}_{AH}$$

10. MARPA IS CORRECT

10.1. Marpa's Earley sets grow at worst linearly.

Theorem 10.1. *For a context-free grammar, and a parse location i_{LOC} ,*

$$|\text{table}[\text{MARPA}, i]| = \mathcal{O}(i).$$

Proof. EIM's have the form $[\mathbf{x}_{AH}, \mathbf{x}_{ORIG}]$. \mathbf{x}_{ORIG} is the origin of the EIM, which in Marpa cannot be after the current Earley set at i_{LOC} , so that

$$0 \leq \mathbf{x}_{ORIG} \leq i_{LOC}.$$

The possibilities for \mathbf{x}_{AH} are finite, since the number of AHFA states is a constant, $|\mathbf{fa}|$, which depends on \mathbf{g} . Since duplicate EIM's are never added to an Earley set, the maximum size of Earley set i_{LOC} is therefore

$$i_{LOC} \times |\mathbf{fa}| = \mathcal{O}(i_{LOC}). \quad \square$$

10.2. Marpa's Earley sets are correct.

Theorem 10.2. *Marpa's Earley sets are correct.*

The proof is by triple induction, that is, induction with a depth down to 3 levels. We number the levels of induction 0, 1 and 2, starting with the outermost. The level 0 induction is usually called the outer induction. The level 1 induction is usually called the inner induction. Level 2 induction is referred to by number.

The outer induction is on the Earley sets. The outer induction hypothesis is that all Earley sets $\mathbf{table}[\mathbf{MARPA}, i]$, $0 \leq i_{LOC} \leq n_{LOC}$, are complete and consistent, and therefore correct. We leave it as an exercise to show, as the basis of the induction, that $\mathbf{table}[\mathbf{MARPA}, 0]$ is complete and consistent.

To show the outer induction step, we show first consistency, then completeness. We show consistency by an inner induction on the Marpa operations. The inner induction hypothesis is that $\mathbf{table}[\mathbf{MARPA}, i]$, as so far built, is consistent with respect to $\mathbf{table}[\mathbf{LEO}, i]$.

As the basis of the inner induction, an empty Marpa Earley set is consistent, trivially. We show the step of the inner induction by cases:

- MARPA scanning operations;
- MARPA reductions when there are no Leo reductions; and
- MARPA's Leo reductions

10.2.1. *Marpa scanning is consistent.* For Marpa's scanning operation, we know that the predecessor EIM is correct by the outer induction hypothesis, and that the token is correct by the definitions in the preliminaries. We know, from Section 8.5, that at most two EIM's will be added. We now examine them in detail.

Let

$$\mathbf{confirmed}_{AH} = \mathbf{GOTO}(\mathbf{predecessor}_{AH}, \mathbf{token}_{SYM})$$

If $\mathbf{confirmed}_{AH} = \Lambda$, the pseudocode of Section 8.5 shows that we do nothing. If we do nothing, since $\mathbf{table}[\mathbf{MARPA}, i]$ is consistent by the inner induction hypothesis, it remains consistent, trivially.

Otherwise, let $\mathbf{confirmed}_{EIM} = [\mathbf{confirmed}_{AH}, i_{LOC}]$. We see that $\mathbf{confirmed}_{EIM}$ is consistent with respect to $\mathbf{table}[\mathbf{LEO}, i]$, by the definition of Earley scanning (Section 5.2) and Observation 9.2. Consistency is invariant under union, and since $\mathbf{table}[\mathbf{MARPA}, i]$ is consistent by the inner induction, $\mathbf{table}[\mathbf{MARPA}, i]$ remains consistent after $\mathbf{confirmed}_{EIM}$ is added.

For predictions, if $\mathbf{confirmed}_{AH} \neq \Lambda$, let

$$\mathbf{predicted}_{AH} = \mathbf{GOTO}(\mathbf{confirmed}_{AH}, \epsilon)$$

If $\text{predicted}_{AH} = \Lambda$, the pseudocode of Section 8.11 shows that we do nothing. If we do nothing, since $\text{table}[\text{MARPA}, i]$ is consistent by the inner induction hypothesis, it remains consistent, trivially. Otherwise, let

$$\text{predicted}_{EIM} = [\text{predicted}_{AH}, i_{LOC}].$$

We see that predicted_{EIM} is consistent with respect to $\text{table}[\text{LEO}, i]$, by the definition of Earley prediction (Section 5.4) and Observation 9.4. Consistency is invariant under union and, since $\text{table}[\text{MARPA}, i]$ is consistent by the inner induction, $\text{table}[\text{MARPA}, i]$ remains consistent after predicted_{EIM} is added.

10.2.2. *Earley reduction is consistent.* Next, we show that MARPA's reduction operation is consistent, in the case where there is no Leo reduction. There will be two cause EIM's, predecessor_{EIM} and component_{EIM} . predecessor_{EIM} will be correct by the outer induction hypothesis and component_{EIM} will be consistent by the inner induction hypothesis. From component_{EIM} , we will find zero or more transition symbols, lhs_{SYM} . From this point, the argument is very similar to that for the case of the scanning operation.

Let

$$\text{confirmed}_{AH} = \text{GOTO}(\text{predecessor}_{AH}, \text{lhs}_{SYM})$$

If $\text{confirmed}_{AH} = \Lambda$, we do nothing, and $\text{table}[\text{MARPA}, i]$ remains consistent, trivially. Otherwise, let

$$\text{confirmed}_{EIM} = [\text{confirmed}_{AH}, i_{LOC}].$$

We see that confirmed_{EIM} is consistent with respect to $\text{table}[\text{LEO}, i]$ by the definition of Earley reduction (Section 5.3), and Observation 9.2. By the invariance of consistency under union, $\text{table}[\text{MARPA}, i]$ remains consistent after confirmed_{EIM} is added.

For predictions, the argument exactly repeats that of Section 10.2.1. $\text{table}[\text{MARPA}, i]$ remains consistent, whether or not a predicted_{EIM} is added.

10.2.3. *Leo reduction is consistent.* We now show consistency for MARPA's reduction operation, in the case where there is a Leo reduction. If there is a Leo reduction, it is signaled by the presence of predecessor_{LIM} ,

$$\text{predecessor}_{LIM} = [\text{top}_{AH}, \text{lhs}_{SYM}, \text{top}_{ORIG}]$$

in the Earley set where we would look for the predecessor_{EIM} . We treat the logic to create predecessor_{LIM} as a matter of memoization of the previous Earley sets, and its correctness follows from the outer induction hypothesis.

As the result of a Leo reduction, LEO will add $[\mathbf{top}_{DR}, \mathbf{top}_{ORIG}]$ to $\mathbf{table}[LEO, j]$. Because the MARPA LIM is correct, using Observations 9.2 and 9.3 and Theorem 7.1, we see that \mathbf{top}_{AH} is the singleton set $\{\mathbf{top}_{DR}\}$. From Section 8.10, we see that, as the result of the Leo reduction, MARPA will add

$$\mathbf{leo}_{EIM} = [\mathbf{top}_{AH}, \mathbf{top}_{ORIG}]$$

to $\mathbf{table}[MARPA, j]$. The consistency of \mathbf{leo}_{EIM} follows from the definition of EIM consistency. The consistency of $\mathbf{table}[MARPA, i]$, once \mathbf{leo}_{EIM} is added, follows by the invariance of consistency under union.

10.2.4. *Marpa's Earley sets are consistent.* Sections 10.2.1, 10.2.2 and 10.2.3 show the cases for the step of the inner induction, which shows the induction. It was the purpose of the inner induction to show that consistency of $\mathbf{table}[MARPA, i]$ is invariant under Marpa's operations.

10.2.5. *The inner induction for completeness.* It remains to show that, when Marpa's operations are run as described in the pseudocode of Section 8, that $\mathbf{table}[MARPA, i]$ is complete. To do this, we show that at least one EIM in $\mathbf{table}[MARPA, i]$ corresponds to every EIMT in $\mathbf{table}[LEO, i]$. We will proceed by cases, where the cases are LEO operations. For every operation that LEO would perform, we show that MARPA performs an operation that produces a corresponding Earley item. Our cases for the operations of LEO are Earley scanning operations; Earley reductions; Leo reductions; and Earley predictions.

10.2.6. *Scanning is complete.* For scanning, the Marpa pseudocode shows that a scan is attempted for every pair

$$[\mathbf{predecessor}_{EIM}, \mathbf{token}_{SYM}],$$

where $\mathbf{predecessor}_{EIM}$ is an EIM in the previous Earley set, and \mathbf{token}_{SYM} is the token scanned at i_{LOC} . (The pseudocode actually finds $\mathbf{predecessor}_{EIM}$ in a set returned by $\mathbf{transitions}()$. This is a memoization for efficiency and we will ignore it.)

By the preliminary definitions, we know that \mathbf{token}_{SYM} is the same in both EARLEY and LEO. By the outer induction hypothesis we know that, for every traditional Earley item in the previous Earley set, there is at least one corresponding Marpa Earley item. Therefore, MARPA performs its scan operation on a complete set of correct operands.

Comparing the Marpa pseudocode (section 8.5), with the Earley scanning operation (section 5.2) and using Observations 9.3 and 9.5, we see that a Earley item will be added to $\mathbf{table}[MARPA, i]$ corresponding to every scanned Earley item of $\mathbf{table}[LEO, i]$. We also see, from the pseudocode of Section 8.11, that the MARPA scanning operation

will add to `table[MARPA, i]` an Earley item for every prediction that results from a scanned Earley item in `table[LEO, i]`.

10.2.7. *Earley reduction is complete.* We now examine Earley reduction, under the assumption that there is no Leo transition. The Marpa pseudocode shows that the Earley items in `table[MARPA, i]` are traversed in a single pass for reduction.

To show that we traverse a complete and consistent series of component Earley items, we stipulate that the Earley set is an ordered set, and that new Earley items are added at the end. From Theorem 10.1, we know that the number of Earley items is finite, so a traversal of them must terminate.

Consider, for the purposes of the level 2 induction, the reductions of LEO to occur in generations. Let the scanned Earley items be generation 0. An EIMT produced by a reduction is generation $n + 1$ if its component Earley item was generation n . Predicted Earley items do not need to be assigned generations. In Marpa grammars they can never contain completions, and therefore can never act as the component of a reduction.

The induction hypothesis for the level 2 induction is that for some n , the Earley items of `table[MARPA, i]` for generations 0 through n are complete and consistent. From Section 10.2.4, we know that all Earley items in Marpa's sets are consistent. In Section 10.2.6, we showed that generation 0 is complete – it contains Earley items corresponding to all of the generation 0 EIMT's of LEO. This is the basis of the level 2 induction.

Since we stipulated that MARPA adds Earley items at the end of each set, we know that they occur in generation order. Therefore MARPA, when creating Earley items of generation $n + 1$ while traversing `table[MARPA, i]`, can rely on the level 2 induction hypothesis for the completeness of Earley items in generation n .

Let `workingEIM ∈ iES` be the Earley item currently being considered as a potential component for an Earley reduction operation. From the pseudocode, we see that reductions are attempted for every pair `predecessorEIM, workingEIM`. (Again, `transitions()` is ignored as a memoization.) By the outer induction hypothesis we know that, for every traditional Earley item in the previous Earley set, there is at least one corresponding Marpa Earley item. We see from the pseudocode, therefore, that for each `workingEIM` that MARPA performs its reduction operation on a complete set of correct predecessors. Therefore MARPA performs its reduction operations on a complete set of operand pairs.

Comparing the Marpa pseudocode (Section 8.9) with the Earley reduction operation (Section 5.3) and using Observations 9.3 and 9.5, we see that a Earley reduction result of generation $n + 1$ will be added to $\text{table}[\text{MARPA}, i]$ corresponding to every Earley reduction result in generation $n + 1$ of $\text{table}[\text{LEO}, i]$, as well as one corresponding to every prediction that results from an Earley reduction result of generation $n + 1$ in $\text{table}[\text{LEO}, i]$. This shows the level 2 induction and the case of reduction completeness.

10.2.8. *Leo reduction is complete.* We now show completeness for MARPA's reduction operation, in the case where there is a Leo reduction. In Section 10.2.3, we found that where LEO would create the EIMT $[\text{top}_{DR}, \text{top}_{ORIG}]$, Marpa adds $[\text{top}_{AH}, \text{top}_{ORIG}]$ such that $\text{top}_{DR} \in \text{top}_{AH}$. Since top_{DR} is a completed rule, there are no predictions. This shows the case immediately, by the definition of completeness.

10.2.9. *Prediction is complete.* Predictions result only from items in the same Earley set. In Sections 10.2.6, 10.2.7 and 10.2.8, we showed that, for every prediction that would result from an item added to $\text{table}[\text{LEO}, i]$, a corresponding prediction was added to $\text{table}[\text{MARPA}, i]$.

10.2.10. *Finishing the proof.* Having shown the cases in Sections 10.2.6, 10.2.7, 10.2.8 and 10.2.9, we know that Earley set $\text{table}[\text{MARPA}, i]$ is complete. In section 10.2.4 we showed that $\text{table}[\text{MARPA}, i]$ is consistent. It follows that $\text{table}[\text{MARPA}, i]$ is correct, which is the step of the outer induction. Having shown its step, we have the outer induction, and the theorem. \square

10.3. **Marpa is correct.** We are now in a position to show that Marpa is correct.

Theorem 10.3. $L(\text{MARPA}, g) = L(g)$

Proof. From Theorem 10.2, we know that

$$[\text{accept}_{DR}, 0] \in \text{table}[\text{LEO}, |w|]$$

if and only there is a

$$[\text{accept}_{AH}, 0] \in \text{table}[\text{MARPA}, |w|]$$

such that $\text{accept}_{DR} \in \text{accept}_{AH}$. From the acceptance criteria in the LEO definitions and the MARPA pseudocode, it follows that

$$L(\text{MARPA}, g) = L(\text{LEO}, g).$$

By Theorem 4.1 in [10], we know that

$$L(\text{LEO}, \mathbf{g}) = L(\mathbf{g}).$$

The theorem follows from the previous two equalities. \square

11. MARPA RECOGNIZER COMPLEXITY

11.1. Complexity of each Earley item. For the complexity proofs, we consider only Marpa grammars without nulling symbols. We showed that this rewrite is without loss of generality in Section 9.1, when we examined correctness. For complexity we must also show that the rewrite and its reversal can be done in amortized $\mathcal{O}(1)$ time and space per Earley item.

Lemma 11.1. *All time and space required to rewrite the grammar to eliminate nulling symbols, and to restore those rules afterwards in the Earley sets, can be allocated to the Earley items in such a way that each Earley item requires $\mathcal{O}(1)$ time and space.*

Proof. The time and space used in the rewrite is a constant that depends on the grammar, and is charged to the parse. The reversal of the rewrite can be done in a loop over the Earley items, which will have time and space costs per Earley item, plus a fixed overhead. The fixed overhead is $\mathcal{O}(1)$ and is charged to the parse. The time and space per Earley item is $\mathcal{O}(1)$ because the number of rules into which another rule must be rewritten, and therefore the number of Earley items into which another Earley item must be rewritten, is a constant that depends on the grammar. \square

Theorem 11.2. *All time in MARPA can be allocated to the Earley items, in such a way that each Earley item, and each attempt to add a duplicate Earley item, requires $\mathcal{O}(1)$ time.*

Theorem 11.3. *All space in MARPA can be allocated to the Earley items, in such a way that each Earley item requires $\mathcal{O}(1)$ space and, if links are not considered, each attempt to add a duplicate Earley item adds no additional space.*

Theorem 11.4. *If links are considered, all space in MARPA can be allocated to the Earley items in such a way that each Earley item and each attempt to add a duplicate Earley item requires $\mathcal{O}(1)$ space.*

Proof of Theorems 11.2, 11.3, and 11.4. These theorems follows from the observations in Section 8 and from Lemma 11.1. \square

11.2. Duplicate dotted rules. The same complexity results apply to MARPA as to LEO, and the proofs are very similar. LEO's complexity results[10] are based on charging resource to Earley items, as were the results in Earley's paper[3]. But both assume that there is one dotted rule per Earley item, which is not the case with MARPA.

MARPA's Earley items group dotted rules into AHFA states, but this is not a partitioning in the strict sense – dotted rules can fall into more than one AHFA state. This is an optimization, in that it allows dotted rules, if they often occur together, to be grouped together aggressively. But it opens up the possibility that, in cases where EARLEY and LEO disposed of a dotted rule once and for all, MARPA might have to deal with it multiple times.

From an efficiency perspective, MARPA's duplicate rules are by all the evidence, a plus. And they do not change the complexity results, although the price of showing this is the theoretical apparatus of this section.

Theorem 11.5.

$$|\mathbf{table}[\mathbf{MARPA}]| < c \times |\mathbf{table}[\mathbf{LEO}]|,$$

where c is a constant that depends on the grammar.

Proof. We know from Theorem 10.2 that every Marpa Earley item corresponds to one of LEO's traditional Earley items. If an EIM corresponds to an EIMT, the AHFA state of the EIM contains the EIMT's dotted rule, while their origins are identical. Even in the worst case, a dotted rule cannot appear in every AHFA state, so that the number of Marpa items corresponding to a single traditional Earley item must be less than $|\mathbf{fa}|$. Therefore,

$$|\mathbf{table}[\mathbf{MARPA}]| < |\mathbf{fa}| \times |\mathbf{table}[\mathbf{LEO}]| \quad \square$$

Earley[3] shows that, for unambiguous grammars, every attempt to add an Earley item will actually add one. In other words, there will be no attempts to add duplicate Earley items. Earley's proof shows that for each attempt to add a duplicate, the causation must be different – that the EIMT's causing the attempt differ in either their dotted rules or their origin. Multiple causations for an Earley item would mean multiple derivations for the sentential form that it represents. That in turn would mean that the grammar is ambiguous, contrary to assumption.

In MARPA, there is a slight complication. A dotted rule can occur in more than one AHFA state. Because of that, it is possible that two

of MARPA's operations to add an EIM will represent identical Earley causations, and therefore will be consistent with an unambiguous grammar. Dealing with this complication requires us to prove a result that is weaker than that of [3], but that is still sufficient to produce the same complexity results.

Theorem 11.6. *For an unambiguous grammar, the number of attempts to add Earley items will be less than or equal to*

$$c \times |\text{table}[\text{MARPA}]|,$$

where c is a constant that depends on the grammar.

Proof. Let **initial-tries** be the number of attempts to add the initial item to the Earley sets. For Earley set 0, it is clear from the pseudocode that there will be no attempts to add duplicate EIM's:

$$\text{initial-tries} = |\text{table}[0]|$$

Let **leo-tries** be the number of attempted Leo reductions in Earley set j_{LOC} . For Leo reduction, we note that by its definition, duplicate attempts at Leo reduction cannot occur. Let **max-AHFA** be the maximum number of dotted rules in any AHFA state. From the pseudo-code of Sections 8.8 and 8.10, we know there will be at most one Leo reduction for each each dotted rule in the current Earley set, j_{LOC} .

$$\text{leo-tries} \leq \text{max-AHFA} \times |\text{table}[j]|$$

Let **scan-tries** be the number of attempted scan operations in Earley set j_{LOC} . Marpa attempts a scan operation, in the worst case, once for every EIM in the Earley set at $j_{LOC} - 1$. Therefore, the number of attempts to add scans must be less than equal to $|\text{table}[j - 1]|$, the number of actual Earley items at $j_{LOC} - 1$.

$$\text{scan-tries} \leq |\text{table}[j - 1]|$$

Let **predict-tries** be the number of attempted predictions in Earley set j_{LOC} . MARPA includes prediction in its scan and reduction operations, and the number of attempts to add duplicate predicted EIM's must be less than or equal to the number of attempts to add duplicate confirmed EIM's in the scan and reduction operations.

$$\text{predict-tries} \leq \text{reduction-tries} + \text{scan-tries}$$

The final and most complicated case is Earley reduction. Recall that j_{ES} is the current Earley set. Consider the number of reductions attempted. MARPA attempts to add an Earley reduction result once for every triple

$$[\text{predecessor}_{EIM}, \text{transition}_{SYM}, \text{component}_{EIM}].$$

where

$$\begin{aligned} & \text{component}_{EIM} = [\text{component}_{AH}, \text{component-origin}_{LOC}] \\ \wedge & \text{component}_{DR} \in \text{component}_{AH} \\ \wedge & \text{transition}_{SYM} = \text{LHS}(\text{component}_{DR}). \end{aligned}$$

We now put an upper bound on number of possible values of this triple. The number of possibilities for transition_{SYM} is clearly at most $|\text{symbols}|$, the number of symbols in g . We have $\text{component}_{EIM} \in j_{ES}$, and therefore there are at most $|j_{ES}|$ choices for component_{EIM} .

We can show that the number of possible choices of predecessor_{EIM} is at most the number of AHFA states, $|\text{fa}|$, by a reductio. Suppose, for the reduction, there were more than $|\text{fa}|$ possible choices of predecessor_{EIM} . Then there are two possible choices of predecessor_{EIM} with the same AHFA state. Call these choice1_{EIM} and choice2_{EIM} . We know, by the definition of Earley reduction, that $\text{predecessor}_{EIM} \in j_{ES}$, and therefore we have $\text{choice1}_{EIM} \in j_{ES}$ and $\text{choice2}_{EIM} \in j_{ES}$. Since all EIM's in an Earley set must differ, and choice1_{EIM} and choice2_{EIM} both have the same AHFA state, they must differ in their origin. But two different origins would produce two different derivations for the reduction, which would mean that the parse was ambiguous. This is contrary to the assumption for the theorem that the grammar is unambiguous. This shows the reductio and that the number of choices for predecessor_{EIM} , compatible with component_{ORIG} , is at most $|\text{fa}|$.

Collecting the results we see that the possibilities for each component_{EIM} are

$$\begin{array}{ll} |\text{fa}| & \text{choices of } \text{predecessor}_{EIM} \\ \times |\text{symbols}| & \text{choices of } \text{transition}_{SYM} \\ \times |j_{ES}| & \text{choices of } \text{component}_{EIM} \end{array}$$

The number of reduction attempts will therefore be at most

$$\text{reduction-tries} \leq |\text{fa}| \times |\text{symbols}| \times |j_{ES}|.$$

Summing

$$\begin{aligned} \text{tries} = & \text{scan-tries} + \text{leo-tries} + \\ & \text{predict-tries} + \text{reduction-tries} + \text{initial-tries}, \end{aligned}$$

we have, where $n = |\mathbf{w}|$, the size of the input,

$$\begin{aligned}
& |\mathbf{table}[0]| && \text{initial EIM's} \\
& + \sum_{i=0}^n \mathbf{max-AHFA} \times |\mathbf{table}[j]| && \text{LIM's} \\
& + 2 \times \sum_{i=1}^n |\mathbf{table}[j - 1]| && \text{scanned EIM's} \\
& + 2 \times \sum_{i=0}^n |\mathbf{fa}| \times |\mathbf{symbols}| \times |\mathbf{j}_{ES}| && \text{reduction EIM's.}
\end{aligned}$$

In this summation, `prediction-tries` was accounted for by counting the scanned and predicted EIM attempts twice. Since `max-AHFA` and `symbols` are both constants that depend only on \mathbf{g} , if we collect the terms of the summation, we will find a constant c such that

$$\mathbf{tries} \leq c \times \sum_{i=0}^n |\mathbf{table}[j]|,$$

and

$$\mathbf{tries} \leq c \times |\mathbf{table}[\mathbf{MARPA}]|,$$

where c is a constant that depends on \mathbf{g} . \square

As a reminder, we follow tradition by stating complexity results in terms of n , setting $n = |\mathbf{w}|$, the length of the input.

Theorem 11.7. *For a context-free grammar,*

$$|\mathbf{table}[\mathbf{MARPA}]| = \mathcal{O}(n^2).$$

Proof. By Theorem 10.1, the size of the Earley set at i_{LOC} is $\mathcal{O}(i)$. Summing over the length of the input, $|\mathbf{w}| = n$, the number of EIM's in all of MARPA's Earley sets is

$$\sum_{i_{LOC}=0}^n \mathcal{O}(i) = \mathcal{O}(n^2). \quad \square$$

Theorem 11.8. *For a context-free grammar, the number of attempts to add Earley items is $\mathcal{O}(n^3)$.*

Proof. Reexamining the proof of Theorem 11.6, we see that the only bound that required the assumption that \mathbf{g} was unambiguous was `reduction-tries`, the count of the number of attempts to add Earley reductions. Let `other-tries` be attempts to add EIM's other than as the result of Earley reductions. By Theorem 11.7,

$$|\mathbf{table}[\mathbf{MARPA}]| = \mathcal{O}(n^2),$$

and by Theorem 11.6,

$$\text{other-tries} \leq c \times |\text{table}[\text{MARPA}]|,$$

so that $\text{other-tries} = \mathcal{O}(n^2)$.

Looking again at reduction-tries for the case of ambiguous grammars, we need to look again at the triple

$$[\text{predecessor}_{EIM}, \text{transition}_{SYM}, \text{component}_{EIM}].$$

We did not use the fact that the grammar was unambiguous in counting the possibilities for transition_{SYM} or component_{EIM} , but we did make use of it in determining the count of possibilities for predecessor_{EIM} . We still know that

$$\text{predecessor}_{EIM} \in \text{component-origin}_{ES},$$

where $\text{component-origin}_{LOC}$ is the origin of component_{EIM} . Worst case, every EIM in $\text{component-origin}_{ES}$ is a possible match, so that the number of possibilities for predecessor_{EIM} now grows to $|\text{component-origin}_{ES}|$, and

$$\text{reduction-tries} = |\text{component-origin}_{ES}| \times |\text{symbols}| \times |j_{ES}|.$$

In the worst case $\text{component-origin} \simeq j$, so that by Theorem 10.1,

$$|\text{component-origin}_{ES}| \times |j_{ES}| = \mathcal{O}(j^2).$$

Adding other-tries and summing over the Earley sets, we have

$$\mathcal{O}(n^2) + \sum_{j_{LOC}=0}^n \mathcal{O}(j^2) = \mathcal{O}(n^3). \quad \square$$

Theorem 11.9. *Either a right derivation has a step that uses a right recursive rule, or it has length is at most c , where c is a constant which depends on the grammar.*

Proof. Let the constant c be the number of symbols. Assume, for a reductio, that a right derivation expands to a Leo sequence of length $c + 1$, but that none of its steps uses a right recursive rule.

Because it is of length $c + 1$, the same symbol must appear twice as the rightmost symbol of a derivation step. (Since for the purposes of these complexity results we ignore nulling symbols, the rightmost symbol of a string will also be its rightmost non-nulling symbol.) So part of the rightmost derivation must take the form

$$\text{earlier-prefix}_{STR} \cdot A_{SYM} \stackrel{\pm}{\Rightarrow} \text{later-prefix}_{STR} \cdot A_{SYM}.$$

But the first step of this derivation sequence must use a rule of the form

$$A_{SYM} \rightarrow \text{rhs-prefix}_{STR} . \text{rightmost}_{SYM},$$

where $\text{rightmost}_{SYM} \stackrel{+}{\Rightarrow} A_{SYM}$. Such a rule is right recursive by definition. This is contrary to the assumption for the reductio. We therefore conclude that the length of a right derivation must be less than or equal to c , unless at least one step of that derivation uses a right recursive rule. \square

11.3. The complexity results. We are now in a position to show specific time and space complexity results.

Theorem 11.10. *For every LR-regular grammar, MARPA runs in $\mathcal{O}(n)$ time and space.*

Proof. By Theorem 4.6 in [10, p. 173], the number of traditional Earley items produced by LEO when parsing input w with an LR-regular grammar g is

$$\mathcal{O}(|w|) = \mathcal{O}(n).$$

MARPA may produce more Earley items than LEO for two reasons: First, MARPA does not apply Leo memoization to Leo sequences which do not contain right recursion. Second, MARPA's Earley items group dotted rules into states and this has the potential to increase the number of Earley items.

By theorem 7.1, the definition of an EIMT, and the construction of a Leo sequence, it can be seen that a Leo sequence corresponds step-for-step with a right derivation. It can therefore be seen that the number of EIMT's in the Leo sequence and the number of right derivation steps in its corresponding right derivation will be the same.

Consider one EIMT that is memoized in LEO. By theorem 7.1 it corresponds to a single dotted rule, and therefore a single rule. If not memoized because it is not a right recursion, this EIMT will be expanded to a sequence of EIMT's. How long will this sequence of non-memoized EIMT's be, if we still continue to memoize EIMT's which correspond to right recursive rules? The EIMT sequence, which was formerly a memoized Leo sequence, will correspond to a right derivation that does not include any steps that use right recursive rules. By Theorem 11.9, such a right derivation can be of length at most c_1 , where c_1 is a constant that depends on g . As noted, this right derivation has the same length as its corresponding EIMT sequence, so that each EIMT not memoized in MARPA will expand to at most c_1 EIMT's.

By Theorem 11.5, when EIMT's are replaced with EIM's, the number of EIM's MARPA requires is at worst, c_2 times the number of EIMT's,

where c_2 is a constant that depends on g . Therefore the number of EIM's per Earley set for an LR-regular grammar in a MARPA parse is less than

$$c_1 \times c_2 \times \mathcal{O}(n) = \mathcal{O}(n).$$

LR-regular grammar are unambiguous, so that by Theorem 11.6, the number of attempts that MARPA will make to add EIM's is less than or equal to c_3 times the number of EIM's, where c_3 is a constant that depends on g . Therefore, by Theorems 11.2 and 11.4, the time and space complexity of MARPA for LR-regular grammars is

$$c_3 \times \mathcal{O}(n) = \mathcal{O}(n). \quad \square$$

Theorem 11.11. *For every unambiguous grammar, MARPA runs in $\mathcal{O}(n^2)$ time and space.*

Proof. By assumption, g is unambiguous, so that by Theorem 11.6, and Theorem 11.7, the number of attempts that MARPA will make to add EIM's is

$$c \times \mathcal{O}(n^2),$$

where c is a constant that depends on g . Therefore, by Theorems 11.2 and 11.4, the time and space complexity of MARPA for unambiguous grammars is $\mathcal{O}(n^2)$. \square

Theorem 11.12. *For every context-free grammar, MARPA runs in $\mathcal{O}(n^3)$ time.*

Proof. By Theorem 11.2, and Theorem 11.8. \square

Theorem 11.13. *For every context-free grammar, MARPA runs in $\mathcal{O}(n^2)$ space, if it does not track links.*

Proof. By Theorem 11.3 and Theorem 11.7. \square

Traditionally only the space result stated for a parsing algorithm is that without links, as in 11.13. This is sufficiently relevant if the parser is only used as a recognizer. In practice, however, algorithms like MARPA are typically used in anticipation of an evaluation phase, for which links are necessary.

Theorem 11.14. *For every context-free grammar, MARPA runs in $\mathcal{O}(n^3)$ space, including the space for tracking links.*

Proof. By Theorem 11.4, and Theorem 11.8. \square

12. THE MARPA INPUT MODEL

In this paper, up to this point, the traditional input stream model has been assumed. As implemented, Marpa generalizes the idea of input streams beyond the traditional model.

Marpa’s generalized input model replaces the input w with a set of tokens, \mathbf{tokens} , whose elements are triples of symbol, start location and length:

$$[t_{SYM}, \mathbf{start}_{LOC}, \mathbf{length}]$$

such that $\mathbf{length} \geq 1$ and $\mathbf{start}_{LOC} \geq 0$. The size of the input, $|w|$, is the maximum over \mathbf{tokens} of $\mathbf{start}_{LOC} + \mathbf{length}$.

Multiple tokens can start at a single location. (This is how MARPA supports ambiguous tokens.) The variable-length, ambiguous and overlapping tokens of MARPA bend the conceptual framework of “parse location” beyond its breaking point, and a new term for parse location is needed. Start and end of tokens are described in terms of **earleme** locations, or simply **earlemes**. Token length is also measured in earlemes.

Like standard parse locations, earlemes start at 0, and run up to $|w|$. Unlike standard parse locations, there is not necessarily a token “at” any particular earleme. (A token is considered to be “at an earleme” if it ends there, so that there is never a token “at” earleme 0.) In fact, there may be earlemes at which no token either starts or ends, although for the parse to succeed, such an earleme would have to be properly inside at least one token. Here “properly inside” means after the token’s start earleme and before the token’s end earleme.

In the Marpa input stream, tokens may interweave and overlap freely, but gaps are not allowed. That is, for all i_{LOC} such that $0 \leq i_{LOC} < |w|$, there must exist

$$\mathbf{token} = [t_{SYM}, \mathbf{start}_{LOC}, \mathbf{length}]$$

such that

$$\begin{aligned} &\mathbf{token} \in \mathbf{tokens} \quad \text{and} \\ &\mathbf{start}_{LOC} \leq i_{LOC} < \mathbf{start}_{LOC} + \mathbf{length}. \end{aligned}$$

The intent of Marpa’s generalized input model is to allow users to define alternative input models for special applications. An example that arises in current practice is natural language, features of which are most naturally expressed with ambiguous tokens. The traditional input stream can be seen as the special case of the Marpa input model where for all $x_{SYM}, y_{SYM}, x_{LOC}, y_{LOC}, x_{length}, y_{length}$, if we have

both of

$$\begin{aligned} [\mathbf{x}_{SYM}, \mathbf{x}_{LOC}, \mathbf{x}_{length}] &\in \text{tokens} \quad \text{and} \\ [\mathbf{y}_{SYM}, \mathbf{y}_{LOC}, \mathbf{y}_{length}] &\in \text{tokens}, \end{aligned}$$

then we have both of

$$\begin{aligned} \mathbf{x}_{length} &= \mathbf{y}_{length} = 1 \quad \text{and} \\ \mathbf{x}_{LOC} = \mathbf{y}_{LOC} &\implies \mathbf{x}_{SYM} = \mathbf{y}_{SYM}. \end{aligned}$$

The correctness results hold for Marpa input streams, but to preserve the time complexity bounds, restrictions must be imposed. In stating them, let it be understood that

$$[\mathbf{x}_{SYM}, \mathbf{x}_{LOC}, \mathbf{length}]_{TOK} \in \text{tokens}$$

We require that, for some constant c , possibly dependent on the grammar \mathbf{g} , that every token length be less than c ,

$$(1) \quad \forall [\mathbf{x}_{SYM}, \mathbf{x}_{LOC}, \mathbf{length}]_{TOK}, \mathbf{length} < c,$$

and that the cardinality of the set of tokens starting at any one location be less than c ,

$$(2) \quad \forall i_{LOC}, \left| \{ [\mathbf{x}_{SYM}, \mathbf{x}_{LOC}, \mathbf{length}]_{TOK} \mid \mathbf{x}_{LOC} = i_{LOC} \} \right| < c$$

Restrictions 1 and 2 impose little or no obstacle to the practical use of Marpa's generalized input model. And with them, the complexity results for MARPA stand.

REFERENCES

- [1] Alfred H. Aho and Jeffrey D. Ullman. *The Theory of Parsing, Translation, and Computing* Prentice-Hall, Englewood Cliff, N.J., 1972.
- [2] John Aycock and R. Nigel Horspool. Practical Earley Parsing *The Computer Journal*, Vol. 45, No. 6, 2002, pp. 620-630.
- [3] J. Earley. An efficient context-free parsing algorithm. *Communications of the Association for Computing Machinery*, 13(2):94-102, 1970.
- [4] Dirk Grune and Criel J.H Jacobs. *Parsing Techniques: A Practical Guide*. Springer, Amsterdam, 2008.
- [5] Edgar T. Irons. A syntax-directed compiler for ALGOL 60. *Communications of the Association for Computing Machinery*, 4(1):51-55, Jan. 1961
- [6] Stephen C. Johnson. Yacc: Yet another compiler-compiler. In *Unix Programmer's Manual Supplementary Documents 1*. 1986.
- [7] Jeffrey Kegler, 2011: Marpa-HTML. <http://search.cpan.org/dist/Marpa-HTML/>.
- [8] Jeffrey Kegler, 2013: Marpa-R2. <http://search.cpan.org/dist/Marpa-R2/>.
- [9] Jeffrey Kegler, 2011: Marpa-XS-1.002000. <http://search.cpan.org/dist/Marpa-XS/>.

- [10] J. M. I. M. Leo. A general context-free parsing algorithm running in linear time on every LR(k) grammar without using lookahead. *Theoretical Computer Science*, 82:165–176, 1991.

CONTENTS

1. Introduction	1
2. Preliminaries	2
3. Rewriting the grammar	5
4. Earley's algorithm	5
5. Operations of the Earley algorithm	7
5.1. Initialization	8
5.2. Scanning	8
5.3. Reduction	8
5.4. Prediction	9
5.5. Causation	9
6. The Leo algorithm	9
6.1. Leo reduction	11
6.2. Changes to Earley reduction	11
6.3. Leo memoization	12
7. The Aycock-Horspool finite automaton	13
8. The Marpa Recognizer	16
8.1. Complexity	16
8.2. Top-level code	17
8.3. Ruby Slippers parsing	17
8.4. Initialization	17
8.5. Scan pass	18
8.6. Reduction pass	18
8.7. Memoize transitions	19
8.8. Reduce one LHS	20
8.9. Earley Reduction operation	21
8.10. Leo reduction operation	21
8.11. Adding a pair of Earley items	21
8.12. Per-set lists	22
8.13. Complexity summary	23
9. Preliminaries to the theoretical results	23
9.1. Nulling symbols	23
9.2. Comparing Earley items	24
9.3. About AHFA states	24
10. Marpa is correct	25
10.1. Marpa's Earley sets grow at worst linearly	25
10.2. Marpa's Earley sets are correct	26
10.3. Marpa is correct	30
11. Marpa recognizer complexity	31
11.1. Complexity of each Earley item	31
11.2. Duplicate dotted rules	32

MARPA, A PRACTICAL GENERAL PARSER: THE RECOGNIZER	43
11.3. The complexity results	37
12. The Marpa Input Model	39
References	40