presented in Figure 1, and to thank his colleagues, Richard Conway, William Maxwell, and Robert Wagner for their helpful discussions of the material.
Received April, 1969; revised August, 1969
REFERENCES

1. Alberga, Cyril, N. String similarity and misspellings. Comm. ACM 10, 5 (May 1967), 302-313.
2. Blair, Charles R. A program for correcting spelling errors. Information and Control 3 (Mar. 1960), 60-67.
3. Conway, R. W., and Maxwell, W. L. CORC-the Cornell computing language. Comm. ACM 6 (June 1963), 317-321.
4. Conway, R. W., and Maxwell, W. L. CUPL-an approach to introductory computing instruction. Tech. Rep. No. 68-4, Dept. of Computer Science, Cornell U., Ithaca, N. Y.
5. Conway, R. W., and Worley, W. S. The Cornell-HASP system for the 360/65. Tech. Rep. No. 68-A1, Office of Computer Services, Cornell U., Ithaca, N. Y.
6. Damerau, F. A technique for computer detection and correction of spelling errors. Comm. ACM 7, 3 (Mar. 1964), 171-176.
7. Davidson, L. Retrieval of misspelled names in an airlines passenger reservation system. Comm. ACM 5, 3 (Mar. 1962), 169-171.
8. Freeman, D. N. Error correction in CORC: The Cornell Computing Language. Ph.D. Th., Cornell U., Ithaca, N. Y., Sept. 1963.
9. Glantz, H. T. On the recognition of information with a digital computer. J. ACM 4, 2 (Apr. 1957), 178-188.
10. Hamming, R. W. One Man's View of Computer Science. $J$. ACM 16, 1 (Jan. 1969), 3-12.
11. Jackson, M. Mnemonics, Datamation 13 (Apr. 1967), 26-29.

# An Efficient Context-Free Parsing Algorithm 

Jay Earley<br>University of California,* Berkeley, California


#### Abstract

A parsing algorithm which seems to be the most efficient general context-free algorithm known is described. It is similar to both Knuth's LR(k) algorithm and the familiar top-down algorithm. It has a time bound proportional to $n^{3}$ (where $n$ is the length of the string being parsed) in general; it has an $n^{2}$ bound for unambiguous grammars; and it runs in linear time on a large class of grammars, which seems to include most practical context-free programming language grammars. In an empirical comparison it appears to be superior to the top-down and bottom-up algorithms studied by Griffiths and Petrick.


KEY WORDS AND PHRASES: syntax analysis, parsing, context-free grammar, compilers, computational complexity
CR CATEGORIES: 4.12, 5.22, 5.23

## 1. Introduction

Context-free grammars (BNF grammars) have been used extensively for describing the syntax of programming languages and natural languages. Parsing algorithms for context-free grammars consequently play a large role in the implementation of compilers and interpreters for pro-

[^0]gramming languages and of programs which "understand" or translate natural languages.
Numerous parsing algorithms have been developed. Some are general, in the sense that they can handle all context-free grammars, while others can handle only subclasses of grammars. The latter, restricted algorithms tend to be much more efficient. The algorithm described here seems to be the most efficient of the general algorithms, and also it can handle a larger class of grammars in linear time than most of the restricted algorithms. We back up these claims of efficiency with both a formal investigation and an empirical comparison.
This paper is based on the author's 1968 report [1] where many of the points studied here appear in much greater detail. In Section 2 the terminology used in this paper is defined. In Section 3 the algorithm is described informally and in Section 4 it is described precisely. Section 5 is a study of the formal efficiency properties of the algorithm and may be skipped by those not interested in this aspect. Section 6 has the empirical comparison and in Section 7 the practical use of the algorithm is discussed.

## 2. Terminology

A language is a set of strings over a finite set of symbols. We call these terminal symbols and represent them by lowercase letters: a, b, c. We use a context-free grammar as a formal device for specifying which strings are in the set. This grammar uses another set of symbols, the nonterminals, which we can think of as syntactic classes. We use capitals for nonterminals: A, B, C. Strings of either terminals or nonterminals are represented by Greek letters: $\alpha, \beta, \gamma$. The empty string is $\lambda . \alpha^{k}$ represents

$$
\overbrace{\alpha \cdots \alpha}^{k \text { times }} .
$$

$|\alpha|$ is the number of symbols in $\alpha$. There is a finite set of productions or rewriting rules of the form $\mathrm{A} \rightarrow \alpha$. The nonterminal which stands for "sentence" is called the root $R$
of the grammar. The productions with a particular nonterminal D on their left sides are called the alternatives of D. Hereafter we use grammar to mean context-free grammar.

We will work with this example grammar of simple arithmetic expressions, grammar AE:

$$
\begin{aligned}
& \mathrm{E} \rightarrow \mathrm{~T} \\
& \mathrm{E} \rightarrow \mathrm{E}+\mathrm{T} \\
& \mathrm{~T} \rightarrow \mathrm{P} \\
& \mathrm{~T} \rightarrow \mathrm{~T} * \mathrm{P} \\
& \mathrm{P} \rightarrow \mathrm{a}
\end{aligned}
$$

The terminal symbols are $\{a,+, *\}$, the nonterminals are $\{\mathrm{E}, \mathrm{T}, \mathrm{P}\}$, and the root is E .

Most of the rest of the definitions are understood to be with respect to a particular grammar G. We write $\alpha \Rightarrow \beta$ if $\exists \gamma, \delta, \eta, \mathrm{A}$ such that $\alpha=\gamma \mathrm{A} \delta$ and $\beta=\gamma \eta \delta$ and $\mathrm{A} \rightarrow \eta$ is a production. We write $\alpha \stackrel{*}{\Rightarrow} \beta$ ( $\beta$ is derived from $\alpha$ ) if $\boldsymbol{\exists}$ strings $\alpha_{0}, \alpha_{1}, \cdots, \alpha_{m}(m \geq 0)$ such that

$$
\alpha=\alpha_{0} \Rightarrow \alpha_{1} \Rightarrow \cdots \Rightarrow \alpha_{m}=\beta
$$

The sequence $\alpha_{0}, \cdots, \alpha_{m}$ is called a derivation (of $\beta$ from $\alpha)$.

A sentential form is a string $\alpha$ such that the $\operatorname{root} \mathrm{R} \stackrel{*}{\Rightarrow} \alpha$. A sentence is a sentential form consisting entirely of terminal symbols. The language defined by a grammar $L(G)$ is the set of its sentences. We may represent any sentential form in at least one way as a derivation tree (or parse tree) reflecting the steps made in deriving it (though not the order of the steps). For example, in grammar AE, either derivation

$$
\mathrm{E} \Rightarrow \mathrm{E}+\mathrm{T} \Rightarrow \mathrm{~T}+\mathrm{T} \Rightarrow \mathrm{~T}+\mathrm{P} \Rightarrow \mathrm{~T} * \mathrm{P}+\mathrm{P}
$$

or

$$
\mathrm{E} \Rightarrow \mathrm{E}+\mathrm{T} \Rightarrow \mathrm{E}+\mathrm{P} \Rightarrow \mathrm{~T}+\mathrm{P} \Rightarrow \mathrm{~T} * \mathrm{P}+\mathrm{P}
$$

is represented by


The degree of ambiguity of a sentence is the number of its distinct derivation trees, A sentence is unambiguous if it has degree 1 of ambiguity. A grammar is unambiguous if each of its sentences is unambiguous. A grammar has bounded ambiguity if there is a bound $b$ on the degree of ambiguity of any sentence of the grammar. A grammar is reduced if every nonterminal appears in some derivation of some sentence.

A recognizer is an algorithm which takes as input a string and either accepts or rejects it depending on whether or not the string is a sentence of the grammar. A parser is a recognizer which also outputs the set of all legal derivation trees for the string.

## 3. Informal Explanation

The following is an informal description of the algorithm as a recognizer: It scans an input string $\mathrm{X}_{1} \cdots \mathrm{X}_{n}$ from left to right looking ahead some fixed number $k$ of symbols. As each symbol $X_{i}$ is scanned, a set of states $S_{i}$ is constructed which represents the condition of the recognition process at that point in the scan. Each state in the set represents (1) a production such that we are currently scanning a portion of the input string which is derived from its right side, (2) a point in that production which shows how much of the production's right side we have recognized so far, (3) a pointer back to the position in the input string at which we began to look for that instance of the production, and (4) a $k$-symbol string which is a syntactically allowed successor to that instance of the production. This quadruple is represented here as a production, with a dot in it, followed by an integer and a string.

For example, if we are recognizing a $*$ a with respect to grammar AE and we have scanned the first a, we would be in the state set $S_{1}$ consisting of the following states (excluding the $k$-symbol strings):

| $\mathrm{P} \rightarrow \mathrm{a}$. | 0 |
| :--- | :--- |
| $\mathrm{~T} \rightarrow \mathrm{P}$. | 0 |
| $\mathrm{~T} \rightarrow \mathrm{~T} . * \mathrm{P}$ | 0 |
| $\mathrm{E} \rightarrow \mathrm{T}$. | 0 |
| $\mathrm{E} \rightarrow \mathrm{E} .+\mathrm{T}$ | 0 |

Each state represents a possible parse for the beginning of the string, given that only the a has been seen. All the states have 0 as a pointer, since all the productions represented must have begun at the beginning of the string.

There will be one such state set for each position in the string. To aid in recognition, we place $k+1$ right terminators " - " (a symbol which does not appear elsewhere in the grammar) at the right end of the input string.

To begin the algorithm, we put the single state

$$
\phi \rightarrow . \mathrm{R} \dashv \quad \dashv^{k} \quad 0
$$

into state set $S_{0}$, where $R$ is the root of the grammar and where $\phi$ is a new nonterminal.

In general, we operate on a state set $S_{i}$ as follows: we process the states in the set in order, performing one of three operations on each one depending on the form of the state. These operations may add more states to $S_{i}$ and may also put states in a new state set $\mathrm{S}_{i+1}$. We describe these three operations by example.

In grammar AE , with $k=1, \mathrm{~S}_{0}$ starts as the single state

$$
\begin{equation*}
\phi \rightarrow . \mathrm{E}-1 \quad-1 \quad 0 \tag{1}
\end{equation*}
$$

The predictor operation is applicable to a state when there is a nonterminal to the right of the dot. It causes us to add one new state to $S_{i}$ for each alternative of that nonterminal. We put the dot at the beginning of the production in each new state, since we have not scanned any of its symbols yet. The pointer is set to $i$, since the state was created in $\mathrm{S}_{i}$. Thus the predictor adds to $\mathrm{S}_{\boldsymbol{i}}$ all productions which might generate substrings beginning at $\mathrm{X}_{i+1}$.
In our example, we add to $\mathrm{S}_{0}$

$$
\begin{align*}
& \mathrm{E} \rightarrow . \mathrm{E}+\mathrm{T} \not-0  \tag{2}\\
& \mathrm{E} \rightarrow . \mathrm{T} \not-0 \tag{3}
\end{align*}
$$

The $k$-symbol look-ahead string is $\not-$, since it is after E in the original state. We must now process these two states. The predictor is also applicable to them. Operating on (2), it produces

$$
\begin{align*}
& \mathrm{E} \rightarrow . \mathrm{E}+\mathrm{T}+0  \tag{4}\\
& \mathrm{E} \rightarrow . \mathrm{T}+0 \tag{5}
\end{align*}
$$

with a look-ahead symbol + because it appears after E in (2). Operating on (3), it produces

$$
\begin{aligned}
& \mathrm{T} \rightarrow . \mathrm{T} * \mathrm{P} \not \dashv 0 \\
& \mathrm{~T} \rightarrow . \mathrm{P} \not-0
\end{aligned}
$$

Here the look-ahead symbol is - because T is last in the production and $\dagger$ is its look-ahead symbol. Now, the predictor, operating on (4) produces (4) and (5) again, but they are already in $S_{0}$, so we do nothing. From (5) it produces

$$
\begin{aligned}
& \mathrm{T} \rightarrow . \mathrm{T} * \mathrm{P}+0 \\
& \mathrm{~T} \rightarrow . \mathrm{P}+0
\end{aligned}
$$

The rest of $\mathrm{S}_{0}$ is

$$
\begin{aligned}
& \mathrm{T} \rightarrow . \mathrm{T} * \mathrm{P} \quad * \quad 0 \\
& \mathrm{~T} \rightarrow . \mathrm{P} * \\
& \mathrm{P} \rightarrow . \mathrm{a} \\
& \mathrm{P} \\
& \mathrm{P} \rightarrow . \mathrm{a}+0 \\
& \mathrm{P} \rightarrow . \mathrm{a}
\end{aligned}+0
$$

The predictor is not applicable to any of the last three states. Instead the scanner is, because it is applicable just in case there is a terminal to the right of the dot. The scanner compares that symbol with $\mathrm{X}_{i+1}$, and if they match, it adds the state to $S_{i+1}$, with the dot moved over one in the state to indicate that that terminal symbol has been scanned.

If $X_{1}=a$, then $S_{1}$ is

$$
\begin{array}{lll}
P \rightarrow a . & \dashv & 0 \\
P \rightarrow a . & + & 0  \tag{6}\\
P \rightarrow \text { a. } & * & 0
\end{array}
$$

these states being added by the scanner.

If we finish processing $S_{i}$ and $S_{i+1}$ remains empty, an error has occurred in the input string. Otherwise, we start to process $\mathrm{S}_{i+1}$.

The third operation, the completer, is applicable to a state if its dot is at the end of its production. Thus the completer is applicable to each of these states in $S_{1}$. It compares the look-ahead string with $\mathrm{X}_{i+1} \cdots \mathrm{X}_{i+k}$. If they match, it goes back to the state set indicated by the pointer, in this case $S_{0}$, and adds all states from $S_{0}$ which have $P$ to the right of the dot. It moves the dot over $P$ in these states. Intuitively, $\mathrm{S}_{0}$ is the state set we were in when we went looking for that P. We have now found it, so we go back to all the states in $S_{0}$ which caused us to look for a P , and we move the dot over the P in these states to show that it has been successfully scanned.
If $\mathrm{X}_{2}=+$, then the completer applied to (6) causes us to add to $\mathrm{S}_{1}$

$$
\begin{array}{lll}
\mathrm{T} \rightarrow \mathrm{P} . & \dagger & 0 \\
\mathrm{~T} \rightarrow \mathrm{P} . & + & 0 \\
\mathrm{~T} \rightarrow \mathrm{P} . & * & 0
\end{array}
$$

Applying the completer to the second of these produces

$$
\begin{array}{lll}
\mathrm{E} \rightarrow \mathrm{~T} . & - & 0 \\
\mathrm{E} \rightarrow \mathrm{~T} . & + & 0 \\
\mathrm{~T} \rightarrow \mathrm{~T} . * \mathrm{P} & \dashv & 0 \\
\mathrm{~T} \rightarrow \mathrm{~T} . * \mathrm{P} & + & 0 \\
\mathrm{~T} \rightarrow \mathrm{~T} . * \mathrm{P} & * & 0
\end{array}
$$

and finally, from the second of these, we get

$$
\begin{aligned}
& \phi \rightarrow \text { E. } \dashv 千- \\
& E \rightarrow E .+T \not- \\
& E \rightarrow E .+T+0
\end{aligned}
$$

The scanner then adds to $\mathrm{S}_{2}$

$$
\begin{aligned}
& \mathrm{E} \rightarrow \mathrm{E}+. \mathrm{T} \nsucc \\
& \mathrm{E} \rightarrow \mathrm{E}+. \mathrm{T}+ \\
& 0
\end{aligned}
$$

If the algorithm ever produces an $S_{i+1}$ consisting of the single state

$$
\phi \rightarrow E \neq \quad \dagger 0
$$

then we have correctly scanned an $E$ and the $-\dashv$, so we are finished with the string, and it is a sentence of the grammar.

A complete run of the algorithm on grammar AE is given in Figure 1. In this example, we have written as one all the states in a state set which differ only in their look-ahead string. (Thus " $-+*$ " as a look-ahead string stands for three states, with " - ", "十", and "*" as their respective look-ahead strings.)

The technique of using state sets and the look-ahead are derived from Knuth's work on LR (k) grammars [2]. In fact our algorithm bears a close relationship to Knuth's algorithm on $\operatorname{LR}(k)$ grammars except for the fact that

| GRAMMAR AE |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $\begin{aligned} & \text { root: } \mathrm{E} \rightarrow \mathrm{~T} \mid \mathrm{E}+\mathrm{T} \quad \text { input string }=a+a * a \\ & \mathrm{~T} \rightarrow \mathrm{P} \mid \mathrm{T} * \mathrm{P} \\ & \mathrm{P} \rightarrow \mathrm{a} \end{aligned}$ |  |  |  |  |  |  |  |
| $k=1$ |  |  |  |  |  |  |  |
| So | $\phi \rightarrow . \mathrm{E}-1$ | -1 | 0 | $\mathrm{S}^{3}$ | $\mathrm{P} \rightarrow \mathrm{a}$. | - +* | 2 |
| ( $\mathrm{X}_{1}=\mathrm{a}$ ) | $\mathrm{E} \rightarrow . \mathrm{E}+\mathrm{T}$ | -1+ | 0 | ( $\mathrm{X}_{4}=*$ ) | $\mathrm{T} \rightarrow \mathrm{P}$. | -1+* | 2 |
|  | $\mathrm{E} \rightarrow$. T | -1+ | 0 |  | $\mathrm{E} \rightarrow \mathrm{E}+\mathrm{T}$. | -1+ | 0 |
|  | $\mathrm{T} \rightarrow . \mathrm{T} * \mathrm{P}$ | $-1+*$ | 0 |  | $\mathbf{T} \rightarrow \mathrm{T} \cdot * \mathrm{P}$ | - +* | 2 |
|  | $\mathrm{T} \rightarrow . \mathrm{P}$ | -1+* |  |  |  |  |  |
|  | $\mathrm{P} \rightarrow \mathrm{a}$ | +1* | 0 | S | $\mathbf{T} \rightarrow \mathbf{T *} . \mathbf{P}$ | -1+* | 2 |
|  |  |  |  | ( $\mathrm{X}_{\mathrm{s}}=\mathrm{a}$ ) | $\mathrm{P} \rightarrow$ a | -1+* | 4 |
| $\begin{aligned} & S_{1} \\ & \left(X_{2}=+\right) \end{aligned}$ | $\mathrm{P} \rightarrow \mathrm{a}$. | -1+* | 0 |  |  |  |  |
|  | $\mathrm{T} \rightarrow \mathrm{P}$. | -1+* | 0 | Ss | $\mathrm{P} \rightarrow \mathrm{a}$. | -1+* | 4 |
|  | $\mathrm{E} \rightarrow \mathrm{T}$. | -1+ | 0 | ( $\mathrm{X}_{6}=-1$ ) | $\mathrm{T} \rightarrow \mathrm{T} * \mathrm{P}$. | -1+* | 2 |
|  | $\mathrm{T} \rightarrow \mathrm{T} . * \mathrm{P}$ | -1+* | 0 |  | $\mathrm{E} \rightarrow \mathrm{E}+\mathrm{T}$. | -1+ | 0 |
|  | $\phi \rightarrow$ E. $\dagger$ | -1 | 0 |  | T $\rightarrow$ T. *P | -1+* | 2 |
|  | $\mathbf{E} \rightarrow \mathrm{E} .+\mathrm{T}$ | $t+$ | 0 |  | $\phi \rightarrow$ E. $\dagger$ |  | 0 |
|  |  |  |  |  | E $\rightarrow$ E. + T | $1+$ | 0 |
| $\begin{aligned} & S_{2} \\ & \left(X_{3}=a\right) \end{aligned}$ | $\mathrm{E} \rightarrow \mathrm{E}+\mathrm{T}$ | $1+$ | 0 |  |  |  |  |
|  | $\mathrm{T} \rightarrow . \mathrm{T} * \mathrm{P}$ | -1** | 2 | So | $\phi \rightarrow \mathrm{E}\}$. | $-1$ | 0 |
|  | $\mathrm{T} \rightarrow$. P | -1+* | 2 |  |  |  |  |
|  | $\mathrm{P} \rightarrow . \mathrm{a}$ | -1* |  |  |  |  |  |

Fig. 1
he uses a stack rather than pointers to keep track of what to do after a parsing decision is made.

Note also that, although it did not develop this way, our algorithm is in effect a top-down parser [3] in which we carry along all possible parses simultaneously in such a way that we can often combine like subparses. This cuts down on duplication of effort and also avoids the leftrecursion problem. (A straightforward top-down parser may go into an infinite loop on grammars containing leftrecursion, i.e. $\mathrm{A} \rightarrow \mathrm{A} \beta$.)

## 4. The Recognizer

The following is a precise description of the recognition algorithm for input string $X_{1} \cdots X_{n}$ and grammar $G$.
Notation. Number the productions of grammar G arbitrarily $1, \cdots, d-1$, where each production is of the form

$$
\mathrm{D}_{p} \rightarrow \mathrm{C}_{p 1} \cdots \mathrm{C}_{p \bar{p}} \quad(1 \leqq p \leqq d-1)
$$

where $\bar{p}$ is the number of symbols on the right-hand side of the $p$ th production. Add a 0th production

$$
\mathrm{D}_{0} \rightarrow \mathrm{R} \dashv
$$

where R is the root of G , and $\dagger$ is a new terminal symbol.
Definition. A state is a quadruple $\langle p, j, f, \alpha\rangle$ where $p, j$, and $f$ are integers $(0 \leqq p \leqq d-1)(0 \leqq j \leqq \bar{p})(0 \leqq f \leqq$ $n+1)$ and $\alpha$ is a string consisting of $k$ terminal symbols. state set is an ordered set of states. A final state is one in which $j=\bar{p}$. We add a state to a state set by putting it last in the ordered set unless it is already a member.
Definition. $\mathrm{H}_{k}(\gamma)=\{\alpha \mid \alpha$ is terminal, $|\alpha|=k$, and $\exists \beta$ such that $\gamma \stackrel{*}{\Rightarrow} \alpha \beta\}$.
$\mathrm{H}_{k}(\gamma)$ is the set of all $k$-symbol terminal strings which begin some string derived from $\gamma$. This is used in forming the look-ahead string for the states.

The Recognizer. This is a function of three arguments $\operatorname{REC}\left(\mathrm{G}, \mathrm{X}_{1} \cdots \mathrm{X}_{n}, k\right)$ computed as follows:
Let $\mathrm{X}_{n+i}=\dagger \quad(1 \leqq i \leqq k+1)$.
Let $S_{i}$ be empty $\quad(0 \leqq i \leqq n+1)$.
Add $\left\langle 0,0,0, \downarrow^{k}\right\rangle$ to $S_{0}$.
For $i \leftarrow 0$ step 1 until $n$ do
Begin
Process the states of $S_{i}$ in order, performing one of the following three operations on each state $\mathrm{s}=\langle p, j, f, \alpha\rangle$.
(1) Predictor: If $s$ is nonfinal and $\mathrm{C}_{p(j+1)}$ is a nonterminal, then for each $q$ such that $\mathrm{C}_{p(j+1)}=\mathrm{D}_{q}$, and for each $\beta \in \mathrm{H}_{k}\left(\mathrm{C}_{p(j+2)} \cdots \mathrm{C}_{p \bar{p} \alpha}\right)$ add $\langle q, 0, i, \beta\rangle$ to $\mathrm{S}_{\boldsymbol{i}}$.
(2) Completer: If s is final and $\alpha=\mathrm{X}_{i+1} \cdots \mathrm{X}_{i+k}$, then for each $\langle q, l, g, \beta\rangle \in \mathrm{S}_{f}$ (after all states have been added to $\mathrm{S}_{f}$ ) such that $\mathrm{C}_{q(l+1)}=\mathrm{D}_{p}$, add $\langle q, l+1, g, \beta\rangle$ to $\mathrm{S}_{i}$.
(3) Scanner: If $s$ is nonfinal and $\mathrm{C}_{p(j+1)}$ is terminal, then if $\mathrm{C}_{p(j+1)}=\mathrm{X}_{i+1}$, add $\langle p, j+1, f, \alpha\rangle$ to $\mathrm{S}_{i+1}$.
If $\mathrm{S}_{i+1}$ is empty, return rejection.
If $i=n$ and $S_{i+1}=\{\langle 0,2,0,-\dagger\rangle\}$, return acceptance.
End
Notice that the ordering imposed on state sets is not important to their meaning but is simply a device which allows their members to be processed correctly by the algorithm. Also note that $i$ cannot become greater than $n$ without either acceptance or rejection occurring because of the fact that $-f$ appears only in production zero. This does not really represent a complete description of the algorithm until we describe in detail how all these operations are implemented on a machine. The following description assumes a knowledge of basic list processing techniques.

## Implementation

(1) For each nonterminal, we keep a linked list of its alternatives, for use in prediction.
(2) The states in a state set are kept in a linked list so they can be processed in order.
(3) In addition, as each state set $S_{i}$ is constructed, we put entries into a vector of size $i$. The fth entry in this vector ( $0 \leqq f \leqq i$ ) is a pointer to a list of all states in $\mathrm{S}_{\boldsymbol{i}}$ with pointer $f$, i.e. states of the form $\langle p, j, f, \alpha\rangle \in S_{i}$ for some $p, j, \alpha$. Thus, to test if a state $\langle p, j, f, \alpha\rangle$ has already been added to $S_{i}$, we search through the list pointed to by the $f$ th entry in this vector. (This takes an amount of time independent of $f$.) The vector and lists can be discarded after $S_{i}$ is constructed.
(4) For the use of the completer, we also keep, for each state set $\mathrm{S}_{i}$ and nonterminal N , a list of all states $\langle\mathrm{p}, j, f, \alpha\rangle$ $\in S_{i}$ such that $\mathrm{C}_{p(j+1)}=\mathrm{N}$.
(5) If the grammar contains null productions ( $\mathrm{A} \rightarrow \lambda$ ), we cannot implement the completer in a straightforward way. When performing the completer on a null state $(A \rightarrow \alpha i)$ we want to add to $S_{i}$ each state in $S_{i}$ with $A$ to the right of the dot. But one of these may not have been added to $S_{i}$ yet. So we must note this and check for it when we add more states to $S_{i}$.

The above implementation description is not meant to be the only way or the best way to implement the algorithm. It is merely a method which does allow the
algorithm to achieve the time and space bounds which we quote in Section 5.
The correctness of this recognizer has been proved in [1]. It requires no restrictions of any kind on the contextfree grammar to be successful.

## 5. Time and Space Bounds

To develop some idea of the efficiency of the algorithm, we can use a formal model of a computer and measure the time as the number of primitive steps executed by this model and the space as the number of storage locations used. We use a random access model (described in the Appendix) because we feel that this model represents most accurately the properties of real computers which are relevant to syntax analysis.
We are interested in upper bounds on the time (and less important the space) as a function of $n$ (the length of the input string) for various classes of context-free grammars. Specifically, an $n^{2}$ algorithm for a subclass A of grammars means that there is some number C (which may depend on the size of the grammar, but not on $n$ ), such that $\mathrm{C} n^{2}$ is an upper bound on the number of primitive steps required to parse any string of length $n$ with respect to a grammar in class A.
The General Case. Our algorithm is an $n^{3}$ recognizer in general. The reasons for this are:
(a) The number of states in any state set $S_{i}$ is proportional to $i(\sim i)$ because the ranges of the $p, j$, and $\alpha$ components of a state are bounded, while only the $f$ component depends on $i$, and it is bounded by $n$.
(b) The scanner and predictor operations each execute a bounded number of steps per state in any state set. So the total time for processing the states in $S_{i}$ plus the scanner and predictor operations is $\sim i$.
(c) The completer executes $\sim i$ steps for each state it processes in the worst case because it may have to add $\sim f$ states for $\mathrm{S}_{f}$, the state set pointed back to. So it takes $\sim i^{2}$ steps in $S_{i}$.
(d) Summing from $i=0, \cdots, n+1$ gives $\sim n^{3}$ steps.

This bound holds even if the look-ahead feature is not used (by setting $k=0$ ). The bound is no better than that obtained by Younger [4] for Cocke's algorithm [5], but our algorithm is better for two reasons. Ours does not require the grammar to be put into any special form (Cocke's required normal form), and ours actually does better than $n^{3}$ on most grammars, as we shall show (Cocke's always requires $n^{3}$ ). Furthermore, although Younger's $n^{3}$ result is obtained on a Turing machine, his algorithm is in no way made faster by putting it on a random access machine.
Unambiguous Grammars. The completer is the only operation which forces us to use $i^{2}$ steps for each state set we process, making the whole thing $n^{3}$. So the question is, in what cases does this operation involve only $i$ steps instead of $i^{2}$ ? If we examine the state set $S_{i}$ after the completer has been applied to it, there are at most proportional to $i$ states in it. So unless some of them were added in more
than one way (this can happen; that's why we must test for the existence of a state before we add it to a state set) then it took at most $\sim i$ steps to do the operation.
In the case that the grammar is unambiguous and reduced, we can show that each such state gets added in only one way. Assume that the state $\langle q, j+1, f, \alpha\rangle$ is added to $S_{i}$ in two different ways by the completer. Then we have

$$
\begin{aligned}
\mathrm{s}_{1} & =\left\langle p_{1}, \bar{p}_{1}, f_{1}, \mathrm{X}_{i+1} \cdots \mathbf{X}_{i+k}\right\rangle \in \mathrm{S}_{i}, \\
\mathrm{~s}_{2} & =\left\langle p_{2}, \bar{p}_{2}, f_{2}, \mathrm{X}_{i+1} \cdots \mathrm{X}_{i+k}\right\rangle \in \mathrm{S}_{i}, \\
\langle q, j, f, \alpha\rangle & \in \mathrm{S}_{f_{1}} \quad \text { and } \quad \mathrm{S}_{f_{2}}, \quad \mathrm{D}_{p_{1}}=\mathrm{C}_{\mathbf{Q}(j+1)}=\mathrm{D}_{p_{2}}
\end{aligned}
$$

and either $p_{1} \neq p_{2}$ or $f_{1} \neq f_{2}$, for otherwise $s_{1}$ and $s_{2}$ would be the same state.
So we have

$$
\begin{array}{r}
\mathrm{X}_{1} \cdots \mathrm{X}_{f} \mathrm{C}_{q 1} \cdots \mathrm{C}_{q(j+1)} \stackrel{*}{\Rightarrow} \mathrm{X}_{1} \cdots \mathrm{X}_{f_{1}} \mathrm{C}_{p_{1} 1} \cdots \mathrm{C}_{p_{1} \bar{p}_{1}} \\
\stackrel{*}{\Rightarrow} \mathrm{X}_{1} \cdots \mathrm{X}_{i}
\end{array}
$$

and

$$
\begin{aligned}
\mathrm{X}_{1} \cdots \mathrm{X}_{j} \mathrm{C}_{q 1} \cdots \mathrm{C}_{q(3+1)} \stackrel{*}{\Rightarrow} \mathrm{X}_{1} \cdots \mathrm{X}_{f_{2}} \mathrm{C}_{p_{2} 1} & \cdots \mathrm{C}_{p_{2} \bar{p}_{2}} \\
& \stackrel{*}{\Rightarrow} \mathrm{X}_{1} \cdots \mathrm{X}_{i}
\end{aligned}
$$

and since $p_{1}=p_{2}$ and $f_{1}=f_{2}$ cannot both be true, the above two derivations of $\mathrm{X}_{1} \cdots \mathrm{X}_{i}$ are represented by different derivation trees. Therefore since the grammar is reduced, every nonterminal generates some terminal string, and so there exists an ambiguous sentence $\mathrm{X}_{1} \cdots \mathrm{X}_{i} \alpha$ for some $\alpha$.
So if the grammar is unambiguous, the completer executes $\sim \mathbf{i}$ steps per state set and the time is bounded by $n^{2}$. Notice that the time is also $n^{2}$ for grammars with bounded ambiguity since each state can then be added by the completer only a bounded number of times. In [1] we show that the time is $n^{2}$ for an even larger class of grammars, and thereby also obtain Younger's $n^{2}$ results for linear and metalinear grammars [6].

Kasami [7] has also obtained independently the result for unambiguous grammars, but his algorithm (which is a modification of Cocke's) has the disadvantage that it requires the grammar in normal form. His algorithm, like ours, achieves its time bound on a random access machine only.
Linear Time. We now characterize the class of grammars which the algorithm will do in time $n$. We notice that for some grammars the number of states in a state set can grow indefinitely with the length of the string being recognized. For some others there is a fixed bound on the size of any state set. We call the latter grammars bounded state grammars. They can be done by the algorithm in time $n$ for the following reason. Let $b$ be the bound on the number of states in any state set. Then the processing of the states in a state set together with the scanner and predictor requires $\sim b$ steps, and the completer requires $\sim b^{2}$ steps. Summing over all the state sets gives us $\sim b^{2} n$ or $\sim n$ steps.
So the class of time $n$ grammars for our algorithm in-
cludes the bounded state grammars, but it actually includes more. We now examine how this class of grammars compares with those that can be done in linear time by other algorithms. Most of the previously mentioned "restricted" algorithms work on some subclass of grammars which they can do in time $n$, and we will henceforth call them time $n$ algorithms. Knuth's LR ( $k$ ) algorithm [2] works on a class of grammars which includes those of just about all the others, so his will be a good one for comparison if we expect to do well.
It turns out that almost all LR $(k)$ grammars are bounded state (except for certain right recursive grammars). And even though some LR (k) grammars may not be bounded state, all of them can be done in time $n$ by our algorithm if a look-ahead of $k$ or greater is used. In fact any finite union of LR $(k)$ grammars (obtained by combining the grammars in a straightforward way in order to generate the union of the languages) is a time $n$ grammar for our algorithm given the proper look-ahead. (This is proved in [1].)
It is here that the look-ahead feature of the algorithm is most obviously useful. We can obtain the $n^{3}$ and $n^{2}$ results without it, but we cannot do all $\operatorname{LR}(k)$ grammars in time $n$ without it. In addition, a look-ahead of $k=1$ is a good practical device for cutting down on a lot of extraneous processing with many common grammars.
The time $n$ grammars for our algorithm, then, include bounded state grammars, finite unions of $\mathrm{LR}(k)$ grammars, and others. They include many grammars which are ambiguous, and some with unbounded degree of ambiguity, but unfortunately there are also unambiguous grammars which require time $n^{2}$.


Fig. 2

The following examples illustrate some of the ideas in this section. Grammar UBDA (Figure 2) actually rerequires time proportional to $n^{3}$. Notice that state

$$
\mathrm{A} \rightarrow \mathrm{AA} . \quad \dagger \mathrm{x} \quad 0^{2}
$$

gets added twice by the completer. This is signified by the superscript on the 0 . One can tell by the looks of the superscripts on states

$$
\begin{array}{lll}
\mathrm{A} \rightarrow \mathrm{AA} . & \dagger \mathrm{x} & 2 \\
\mathrm{~A} \rightarrow \mathrm{AA} . & \dagger \mathrm{x} & 1^{2} \\
\mathrm{~A} \rightarrow \mathrm{AA} . & \dashv \mathrm{x} & 0^{3}
\end{array}
$$

in $\mathrm{S}_{4}$ that there are $\sim i$ states, each of which is added $\sim i$ times, in $\sim n$ state sets, producing the $n^{3}$ behavior.

Grammar BK (Figure 3) has unbounded ambiguity, but it is a time $n$ grammar, and in fact it is bounded state. This can be seen because all the state sets after $S_{1}$ are the same size. Grammar PAL (Figure 4) is an unambiguous grammar which requires time $n^{2}$. $S_{i}$ and $S_{i+1}$ each have $i+4$ states in them, up to and including $\mathrm{S}_{n}$, so the total number of states is $\sim n^{2}$.

Space. Since the space is taken up by $\sim n$ state sets, each containing $\sim n$ states, the space bound is $n^{2}$ in general. This is comparable to Cocke's and Kasami's algorithms, which also require $n^{2}$. However, it has the advantage over Cocke's in that the $n^{2}$ is only an upper bound for ours, while his requires $n^{2}$ all the time.

## GRAMMAR BK



Fig. 3

$$
\begin{aligned}
& \text { GRAMMAR PAL } \\
& A \rightarrow x \mid x A x \\
& \text { sentences: } \mathrm{X}^{n}(n \geqq 1, n \text { odd) } \\
& \text { REC(PAL, } \left.\mathrm{x}^{5}, 0\right) \text { : } \\
& \mathrm{S}_{0} \quad \phi \rightarrow \text {. A- } 0 \quad \mathrm{~S}_{4} \quad \mathrm{~A} \rightarrow \mathrm{xAx} .1 \\
& \mathrm{~A} \rightarrow . \mathrm{x} \quad 0 \\
& A \rightarrow . x A x \quad 0 \\
& \mathrm{~S}_{1} \quad \mathrm{~A} \rightarrow \mathrm{x} . \quad 0 \\
& \mathrm{~A} \rightarrow \mathrm{x} . \mathrm{Ax} \quad 0 \\
& \phi \rightarrow \text { A. } \dagger \quad 0 \\
& \mathrm{~A} \rightarrow \mathrm{x} \quad 1 \\
& A \rightarrow \text {. } \mathrm{xAx} 1 \\
& \mathrm{~S}_{2} \quad \mathrm{~A} \rightarrow \mathrm{x} . \quad 1 \\
& A \rightarrow X . A x \quad 1 \\
& \mathrm{~A} \rightarrow \mathrm{XA} . \mathrm{x} \quad 0 \\
& \mathrm{~A} \rightarrow . \mathrm{x} \quad 2 \\
& A \rightarrow \text {.xAx } 2 \\
& S_{3} \quad A \rightarrow x A x . \quad 0 \\
& A \rightarrow x \\
& A \rightarrow X . A x \quad 2 \\
& \phi \rightarrow \text { A. } \dagger \quad 0 \\
& A \rightarrow x A . x \quad 1 \\
& \mathrm{~A} \rightarrow \text {. } \\
& \mathrm{A} \rightarrow . \mathrm{xAx} \quad 3
\end{aligned}
$$

Fig. 4

## 6. Empirical Results

We have programmed the algorithm and tested it against the top-down and bottom-up parsers evaluated by Griffiths and Petrick [8]. These are the oldest of the context-free parsers, and they depend heavily on backtracking. Perhaps because of this, their upper bounds for time are exponential ( $\mathrm{C}^{n}$ for some constant C ). However, they also can do well on some grammars, and both have been used in numerous compiler-compilers, so it will be interesting to compare our algorithm with them.
The Griffiths and Petrick data is not in terms of actual running times but in terms of "primitive operations." They have expressed their algorithms as sets of nondeterministic rewriting rules for a Turing-machine-like device. Each application of one of these is a primitive operation. We have chosen as our primitive operation the act of adding a state to a state set (or attempting to add one which is already there). We feel that this is comparable to their primitive operation because both are in some sense the most complex operation performed by the algorithm whose complexity is independent of the size of the grammar or input string.
We compare the algorithms on seven different grammars. Two of their examples were not used because the exact grammar was not given. For the first four, Griffiths and Petrick were able to find closed-form expressions for their results, so we did also (Figure 5). BU and TD are the bottom-up and top-down algorithms respectively and SBU and STD are their selective versions. It is obvious from these results that SBU is by far the best of the other algorithms, and the rest of their data bears this out. There-


Fig. 5

PROPOSITIONAL CALCULUS GRAMMAR

| $\text { root: } \begin{aligned} & \mathrm{F} \rightarrow \\ & \mathrm{C} \rightarrow \\ & \mathrm{U} \rightarrow \\ & \mathrm{~L} \rightarrow \\ & \mathrm{~S} \rightarrow \\ & \mathrm{P} \rightarrow \end{aligned}$ | $L \rightarrow L^{\prime}$ $p$ $q$ $r$ <br> $S \rightarrow U \vee S$ $U \vee U$   <br> $P \rightarrow U \wedge$    |  |  |  |
| :---: | :---: | :---: | :---: | :---: |
| Sentence | Length | PA | SBU | Ours |
| p | 1 | 14 | 18 | 28 |
| $(p \wedge q)$ | 5 | 89 | 56 | 68 |
| $\left(p^{\prime} \wedge q\right) \vee r \vee p \vee q^{\prime}$ | 13 | 232 | 185 | 148 |
| $p \supset\left(\left(q \supset \sim\left(r^{\prime} \vee(p \wedge q)\right)\right)\right.$ | 26 | 712 | 277 | 277 |
| $\supset\left(q^{\prime} \vee \mathrm{r}\right)$ |  |  |  |  |
| $\left.\sim\left(\sim p^{\prime} \wedge(q \vee r) \wedge p^{\prime}\right)\right)$ | 17 | 1955 | 223 | 141 |
| $\left((p \wedge q) \vee(q \wedge r) \vee\left(r \wedge p^{\prime}\right)\right)$ | 38 | 2040 | 562 | 399 |
| $\supset \sim\left(\left(p^{\prime} \vee q^{\prime}\right) \wedge\left(r^{\prime} \vee p\right)\right)$ |  |  |  |  |

Fig. 6

| GRAMMAR GRE |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: |
| $\text { root: } \begin{aligned} \mathrm{X} & \rightarrow \mathrm{a}\|\mathrm{Xb}\| \mathrm{Ya} \\ \mathrm{Y} & \rightarrow \mathrm{e} \mid \mathrm{YdY} \end{aligned}$ |  |  |  |  |
|  |  |  |  |  |
| Sentence | Length | PA | SBU | Ours |
| ededea | 6 | 35 | 52 | 33 |
| ededeab ${ }^{4}$ | 10 | 75 | 92 | 45 |
| ededeab ${ }^{10}$ | 16 | 99 | 152 | 63 |
| ededeab ${ }^{200}$ | 206 | 859 | 2052 | 633 |
| (ed) ${ }^{4} \mathrm{eabb}$ | 12 | 617 | 526 | 79 |
| (ed) ${ }^{7} \mathrm{eabb}$ | 18 | 24352 | 16336 | 194 |
| (ed) ${ }^{8} \mathrm{eabb}$ | 20 | 86139 | 54660 | 251 |

Fig. 7

| GRAMMAR NSE |  |  |  |
| :---: | :---: | :---: | :---: |
| root: $\mathrm{S} \rightarrow \mathrm{AB}$ |  |  |  |
|  | $\mathrm{A} \rightarrow \mathrm{a} \mid \mathrm{SC}$ |  |  |
| $\mathrm{B} \rightarrow \mathrm{b} \mid \mathrm{DB}$ |  |  |  |
| $\mathrm{C} \rightarrow \mathrm{c}$ |  |  |  |
| $\mathrm{D} \rightarrow \mathrm{d}$ |  |  |  |
| Sentence | Length | SBU | Ours |
| adbcddb | 7 | 43 | 44 |
| $\mathrm{ad}^{3}$ bebed ${ }^{3} \mathrm{bcd}{ }^{4} \mathrm{~b}$ | 18 | 111 | 108 |
| adbcd ${ }^{2} \mathrm{bc} \mathrm{d}^{5} \mathrm{bcd}{ }^{3} \mathrm{~b}$ | 19 | 117 | 114 |
| $a d^{18} \mathrm{~b}$ | 20 | 120 | 123 |
| $a(b c)^{3} \mathrm{~d}^{8}(\mathrm{bcd})^{2} \mathrm{dbc} \mathrm{d}^{4} \mathrm{~b}$ | 24 | 150 | 141 |
| $a(b c d)^{2} \mathrm{dbc}^{3} \mathrm{bcb}$ | 16 | 100 | 95 |

Fig. 8
fore we compare our algorithm with SBU only. We used our algorithm with $k=0$. The two are comparable on G1, G2, G3, the simple grammars, but on G4, which is very ambiguous, ours is clearly superior- $n$ to $n^{3}$.

For the next three grammars we present only the raw data (Figures 6-8). The data for our algorithm was obtained by programming it and having the program compute the number of primitive operations it performed. We have also included the data from [8] on PA, the predictive analyzer, which is a modified top-down algorithm. On the propositional calculus grammar, PA seems to be running in time $n^{2}$, while both SBU and ours run in time $n$, with ours a little faster. Grammar GRE produces two kinds of behavior. All three algorithms go up linearly with the number of " $b$ "'s, with SBU using a considerably higher constant coefficient. However, PA and SBU go up exponentially with the number of "ed"'s, while ours goes up as the square. Grammar NSE is quite simple, and each algorithm takes time $n$ with the same coefficient.

So we conclude that our algorithm is clearly superior to the backtracking algorithms. It performs as well as the best of them on all seven grammars and is substantially faster on some.

There are at least four distinct general context-free algorithms besides ours-TD, BU, Kasami's $n^{2}$, and Cocke's $n^{3}$. We have shown so far that our algorithm achieves time bounds which are as good as those of any of these algorithms, or better. However, we are also interested in how our algorithm compares with these algorithms in a practical sense, not just at an upper bound.

We have just presented some empirical results in this section which indicate that our algorithm is better than TD and BU. Furthermore, our algorithm must be superior to Cocke's since his always achieves its upper bound of $n^{3}$. This leaves Kasami's. His algorithm [7] is actually described as an algorithm for unambiguous grammars, but it can easily be extended to a general algorithm. In this form we suspect that it will have an $n^{3}$ bound in general and will be $n^{2}$ as often as ours. We are aware of no results about the class of grammars that it can parse in time $n$.

## 7. The Practical Use of the Algorithm

In this section we discuss the question, in what areas and in what form can the algorithm best be put to use?

The Form. Before we can do much with it, we must make the recognizer into a parser. This is done by altering the recognizer so that it builds a parse tree as it does the recognition process. Each time we perform the completer operation adding a state $\mathrm{E} \rightarrow \alpha \mathrm{D} \cdot \beta \mathrm{g}$ (ignoring lookahead) we construct a pointer from the instance of $D$ in that state to the state $\mathrm{D} \rightarrow \gamma$. f which caused us to do the operation. This indicates that D was parsed as $\gamma$. In case D is ambiguous there will be a set of pointers from it, one for each completer operation which caused $\mathrm{E} \rightarrow \alpha \mathrm{D} \cdot \beta \mathrm{g}$ to be added to the particular state set. Each symbol in $\gamma$ will also have pointers from it (unless it is terminal), and so on, thus representing the derivation tree for D .

In this way, when we reach the terminating state $\phi \rightarrow \mathrm{R} \dashv .0$ we will have the parse tree for the sentence hanging from $R$ if it is unambiguous, and otherwise we will have a factored representation of all possible parse trees. In [1] a precise description of this process is given.

The time bounds for the parser are the same as those of the recognizer, while the space bound goes up to $n^{3}$ in general in order to store the parse trees.

We recommend using consistently a look-ahead of $k=1$. In fact it would probably be most efficient to implement the algorithm to do just a look-ahead of 1 . To implement the full look-ahead for any $k$ would be more costly in programming effort and less efficient overall since so few programming languages need the extra look-ahead. Most programming languages use only a one character context to disambiguate their constructs, and if two characters are needed in some cases, our algorithm has the nice property that it will not fail, it may just take a little longer.

Our algorithm has the useful property that it can be modified to handle an extension of context-free grammars which makes use of the Kleene star notation. In this notation: $A \rightarrow\{B C\} * D$ means $A$ may be rewritten as an arbitrary number (including 0 ) of BC 's followed by a D . It generates a language equivalent to that generated by $\mathrm{A} \rightarrow \mathrm{D} \mid \mathrm{BCA}$. However, the parse structure given to the language is different in the two grammars:


Structures like that on the left cannot be obtained using context-free grammars at all, so this extension is useful. The modification to our algorithm which implements it involves two additional operations:
(1) Any state of the form

$$
\mathrm{A} \rightarrow \alpha \cdot\{\beta\}^{*} \gamma \quad \mathrm{f}
$$

is replaced by

$$
\begin{array}{ll}
\mathrm{A} \rightarrow \alpha\{. \beta\}^{*} \gamma & \mathrm{f} \\
\mathrm{~A} \rightarrow \alpha\{\beta\}^{*} \cdot \gamma & \mathrm{f}
\end{array}
$$

indicating that $\beta$ may be present or absent.
(2) Any state of the form

$$
\mathrm{A} \rightarrow \alpha\{\beta .\}^{*} \gamma \quad \mathrm{f}
$$

is replaced by

$$
\begin{array}{ll}
\mathrm{A} \rightarrow \alpha\{. \beta\}^{*} \gamma & \mathrm{f} \\
\mathrm{~A} \rightarrow \alpha\{\beta\}^{*} \cdot \gamma & \mathrm{f}
\end{array}
$$

indicating the $\beta$ may be repeated or not.

The Use. The algorithm will probably be most useful in natural language processing systems where the full power of context-free grammars is used. It should also be useful in compiler writing systems and extendible languages. In most compiler writing systems and extendible languages, the programmer is allowed to express the syntax (or the syntax extension) of his language in something like BNF, and the system uses a parser to analyze subsequent programs in this language. Programming language grammars tend to lie in a restricted subset of context-free grammars which can be processed efficiently, yet some compiler writing systems in fact use general parsers, so ours may be of use here. In addition to its efficiency properties, ours has the advantage that it accepts the grammar in the form in which it is written, so that semantic routines can be associated with productions without fear that the parser will not reflect the original structure of the grammar.

Our algorithm will not compete so favorably with the time $n$ algorithms, however. Certainly ours will do in time $n$ any grammar that a time $n$ parser can do at all, but this does not take into account the constant coefficient of $n$. Most of the time $n$ algorithms really consist of a two-fold process. First they compile from an acceptable grammar a parser for that particular grammar, and then the grammar may be discarded and the compiled parser used directly to analyze strings. This allows the time $n$ algorithms to incorporate much specialized information into the compiled parser, thus reducing the coefficient of $n$ to something quite small-probably an order of magnitude less than that of our algorithm.

Consequently we have developed a compilation process for our algorithm which works only on time $n$ grammars and reduces our coefficient to approximately the same order of magnitude as those of the time $n$ parsers. This may make our algorithm competitive with them, but we have not implemented and tested it, so this is speculation. Some sort of efficient time $n$ parser for a larger class of grammars is needed, however, because most restricted parsers suffer from the problem that the grammar one naturally writes for many programming languages is not acceptable to them, and much fiddling must be done with the grammar to get it accepted. Knuth's algorithm is an exception to this, but it has the problem that the size of the compiled parser is much too great for reasonable programming language grammars (see [1, p. 129]). Unfortunately, our compiled algorithm, since it is similar to Knuth's, may also have these problems.

## 8. Conclusion

In conclusion let us emphasize that our algorithm not only matches or surpasses the best previous results for times $n^{3}$ (Younger), $n^{2}$ (Kasami) and $n$ (Knuth), but it does this with one single algorithm which does not have specified to it the class of grammars it is operating on and does not require the grammar in any special form. In other words, Knuth's algorithm works only 'on LR (k)
grammars and Kasami's (at least in his paper) only on unambiguous ones, but ours works on them all and seems to do about as well as other algorithms automatically.

## Appendix

Random Access Machine. This model has an unbounded number of registers (counters), each of which may contain any nonnegative integer. These registers are named (addressed) by successive nonnegative integers. The primitive operations which are allowed on these registers are as follows:
(1) Store 0 or the contents of one register into another.
(2) Test the contents of one register against 0 or against the contents of another register for equality.
(3) Add 1 or subtract 1 from the contents of a register (taking $0-1=0$ ).
(4) Add the contents of one register to another.

The control for this model is a normal finite state device. The most important property of this machine is that in the above four operations, the register $R$ to be operated on may be specified in two ways:
(1) $R$ is the register whose address is (register $n$ ).
(2) $R$ is the register whose address is the contents of register $n$.
This second mode (sometimes called indirect addressing) plus primitive operation 4 (used for array accessing) gives our model the random access property. The time is measured by the number of primitive operations performed, and the space is measured by the number of registers used in any of these operations.
Acknowledgments. I am deeply indebted to Robert Floyd for his guidance in this research. I also benefited from discussions with Albert Meyer, Rudolph Krutar, and James Gray, and from detailed criticisms by the referees.
Received February, 1969; revised June, 1969

## REFERENCES

1. Earley, J. An efficient context-free parsing algorithm. Ph.D. Thesis, Comput. Sci. Dept., Carnegie-Mellon U., Pittsburgh, Pa., 1968.
2. Knuth, D. E. On the translation of languages from left to right. Information and Control 8 (1965), 607-639.
3. Floyd, R. W. The syntax of programming languages-a survey. IEEE Trans. EC-18, 4 (Aug. 1964).
4. Younger, D. H. Recognition and parsing of context-free languages in time $n^{3}$. Information and Control 10 (1967), 189208.
5. Hays, D. Automatic language-data processing. In Computer Applications in the Behavioral Sciences, H. Borko (Ed.) Prentice Hall, Englewood Cliffs, N.J., 1962.
6. Younger, D. H. Context-free language processing in time $n^{3}$. General Electric R \& D Center, Schenectady, N.Y., 1966.
7. Kasami, T., and Torii, K. A syntax-analysis procedure for unambiguous context-free grammars. J. ACM 16, 3 (July 1969), 423-431.
8. Griffiths, T., and Petrick, S. On the relative efficiencies of context-free grammar recognizers. Comm. ACM 8, 5 (May 1965), 289-300.
9. Feldman, J., and Gries, D. Translator writing systems. Comm. ACM 11, 2 (Feb. 1968), 77-113.

[^0]:    * Computer Science Department. This work was partially supported by the Office of Naval Research under Contract No. NONR3656(23) with the Computer Center, University of California, Berkeley, and by the Advanced Research Projects Agency of the Office of the Secretary of Defense (F44620-67-C-0058), monitored by the Air Force Office of Scientific Research.

