

# Nez: Practical Open Grammar Language

Kimio Kuramitsu

Yokohama National University, JAPAN

kimio@ynu.ac.jp

<http://nez-peg.github.io/>

## Abstract

Nez is a PEG(Parsing Expressing Grammar)-based open grammar language that allows us to describe complex syntax constructs without action code. Since open grammars are declarative and free from a host programming language of parsers, software engineering tools and other parser applications can reuse once-defined grammars across programming languages.

A key challenge to achieve practical open grammars is the expressiveness of syntax constructs and the resulting parser performance, as the traditional action code approach has provided very pragmatic solutions to these two issues. In Nez, we extend the symbol-based state management to recognize context-sensitive language syntax, which often appears in major programming languages. In addition, the Abstract Syntax Tree constructor allows us to make flexible tree structures, including the left-associative pair of trees. Due to these extensions, we have demonstrated that Nez can parse many grammars of practical programming languages.

Nez can generate various types of parsers since all Nez operations are independent of a specific parser language. To highlight this feature, we have implemented Nez with dynamic parsing, which allows users to integrate a Nez parser as a parser library that loads a grammar at runtime. To achieve its practical performance, Nez operators are assembled into low-level virtual machine instructions, including automated state modifications when backtracking, transactional controls of AST construction, and efficient memoization in packrat parsing. We demonstrate that Nez dynamic parsers achieve very competitive performance compared to existing efficient parser generators.

**Categories and Subject Descriptors** D.3.1 [*Programming Languages*]: Formal Definitions and Theory – Syntax; D.3.4 [*Programming Languages*]: Processors – Parsing

**General Terms** Languages, Algorithms

**Keywords** Parsing expression grammars, context-sensitive grammars, parser generators, and grammar specification language

## 1. Introduction

A parser generator is a standard approach to reliable parser development in programming languages and many parser applications. Developers use a formal grammar, such as LALR( $k$ ), LL( $k$ ), or PEG, to specify programming language syntax. Based on the grammar specification, a parser generator tool, such as Yacc[14], ANTLR3/4[29], or Rats[9], produces an efficient parser code that can be integrated with the host language of the compiler and interpreter. This generative approach, obviously, enables developers to avoid error-prone coding for lexical and syntactic analysis.

Traditional parser generators, however, are not entirely free from coding. One particular reason is that the formal grammars used today lack sufficient expressiveness for many of the popular programming language syntaxes, such as typedef in C/C++ and other context-sensitive syntaxes[3, 8, 9]. In addition, a formal grammar itself is a syntactic specification of the input language, not a schematic specification for tree representations that are needed in semantic analysis. To make up for these limitations, most parser generators take an ad hoc approach called *semantic action*, a fragment of code embedded in a formal grammar specification. The embedded code is written in a host language and combined in a way that it is hooked at a parsing context that requires extended recognition or tree constructions.

The problem with arbitrary action code is that they decrease the reusability of the grammar specification[28]. For example, Yacc cannot generate the Java version of a parser simply because Yacc uses C-based action code. There are many Yacc ports to other languages, but they do not reuse the original Yacc grammar without porting of action code. This is an undesirable situation because the well-defined gram-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [Permissions@acm.org](mailto:Permissions@acm.org).

Copyright is held by the owner/author(s). Publication rights licensed to ACM.

*Onward!*'16, November 2–4, 2016, Amsterdam, Netherlands  
ACM. 978-1-4503-4076-2/16/11...\$15.00  
<http://dx.doi.org/10.1145/2986012.2986019>

mar is demanded everywhere in IDEs and other software engineering tools[17].

Nez is a grammar language designed for open grammars. In other words, our aim is that once we write a grammar specification, anyone can use the grammar in any programming language. To achieve the openness of the grammar specification, Nez eliminates any action code from the beginning, and provides a declarative but small set of operators that allow the context-sensitive parsing and flexible Abstract Syntax Tree construction that are needed for popular programming language syntaxes. This paper presents the implementation status of Nez with the experience of our grammar development.

The Nez grammar language is based on *parsing expression grammars (PEGs)*, a popular syntactic foundation formalized by Ford[8]. PEGs are simple and portable and have many desirable properties for defining modern programming language syntax, but they still have several limitations in terms of defining popular programming languages. Typical limitations include typedef-defined name in C/C++ [8, 9], delimiting identifiers (such as <<END) of the Here document in Perl and Ruby, and indentation-based code layout in Python and Haskell [1].

In the first contribution of Nez, we model a single extended parser state, called a *symbol table*, to describe a variety of context-sensitive patterns appearing in programming language constructs. The key idea behind the symbol table is simple. We allow any parsed substrings to be handled as *symbols*, a specialized word in another context. The symbol table is a runtime, growing, and recursive storage for such symbols. If we handle typedef-defined names as symbols for example, we can realize different parsing behavior through referencing the symbol table. Nez provides a small set of symbol operators, including matching, containment testing, and scoping. More importantly, since the symbol table is a single, unified data structure for managing various types of symbols, we can easily trace state changes and then automate state management when backtracking. In addition, the symbol table itself is simply a stack-based data structure; we can translate it into any programming language.

Another important role of the traditional action code is the construction of ASTs. Basically, PEGs are just a syntactic specification, not a schematic specification for tree structures that represent AST nodes. In particular, it is hard to derive the left-associative pair of trees due to the limited left recursion in PEGs. In Nez, we introduce an operator to specify tree structures, modeled on capturing in perl compatible regular expressions (PCREs), and extend tree manipulations including tagging for typing a tree, labeling for sub-nodes, and left-folding for the left-associative structure. Due to these extensions, a Nez parser produces arbitrary representations of ASTs, which would leads to less impedance mismatching in semantic analysis.

To evaluate the expressiveness of Nez, we have performed extensive case studies to specify various popular programming languages, including C, C#, Java8, JavaScript, Python, Ruby, and Lua. Surprisingly, the introduction of a simple symbol table improves the declarative expressiveness of PEGs in a way that Nez can recognize almost all of our studied languages and then produce a practical representation of ASTs. Our case studies are not complete, but they indicate that the Nez approach is promising for a practical grammar specification without any action code.

At last but not least, parsing performance is an important factor since the acceptance of a parser generator relies definitively on practical performance. In this light, Nez can generate three types of differently implemented parsers, based on the traditional source generation, the grammar translation, and the parsing machine[11]. From the perspective of open grammars, we highlight the latter parsing machine that enables a PCRE-style parser library that loads a grammar at runtime. To achieve practical performance, Nez operators are assembled into low-level virtual instructions, including automated state modifications when backtracking, transactional controls of AST construction, and efficient memoization in packrat parsing. We will demonstrate that the resulting parsers are portable and achieve very competitive performance compared to other existing standard parsers for Java, JavaScript, and XML.

The remainder of this paper proceeds as follows. Section 2 is an informal introduction to the Nez grammar specification language. Section 3 is a formal definition of Nez. Section 4 describes parser runtime and implementation. Section 5 reports a summary of our case studies. Section 6 studies the parser performance. Section 7 briefly reviews related work. Section 8 concludes the paper. The tools and grammars presented in this paper are available online at <http://nez-peg.github.io/>

## 2. A Taste of Nez

This section is an informal introduction to the Nez grammar specification language.

### 2.1 Nez and Parsing Expression

Nez is a PEG-based grammar specification language. The basic constructs come from those of PEGs, such as production rules and parsing expressions. A Nez grammar is a set of production rules, each of which is defined by a mapping from a nonterminal  $A$  to a parsing expression  $e$ :

$$A = e$$

Table 1 shows a list of Nez operators that constitute the parsing expressions. All PEG operators inherit the formal interpretation of PEGs[8]. That is, the string 'abc' exactly matches the same input, while [abc] matches one of these characters. The . operator matches any single character. The  $e?$ ,  $e*$ , and  $e+$  expressions behave as in common regular

PEG	Type	Description
' '	Primary	Matches text
[]	Primary	Matches character class
.	Primary	Any character
<i>A</i>	Primary	Non-terminal application
( <i>e</i> )	Primary	Grouping
<i>e</i> ?	Unary suffix	Option
<i>e</i> *	Unary suffix	Zero-or-more repetitions
<i>e</i> +	Unary suffix	One-or-more repetitions
& <i>e</i>	Unary prefix	And-predicate
! <i>e</i>	Unary prefix	Negation
<i>e</i> <sub>1</sub> <i>e</i> <sub>2</sub>	Binary	Sequencing
<i>e</i> <sub>1</sub> / <i>e</i> <sub>2</sub>	Binary	Prioritized Choice
<hr/>		
AST		
{ <i>e</i> }	Construction	Constructor
\$( <i>e</i> )	Construction	Connector
{ \$ <i>e</i> }	Construction	Left-folding
# <i>t</i>	Construction	Tagging
' '	Construction	Replaces a string
<hr/>		
Symbols		
<symbol <i>A</i> >	Action	Symbolize Nonterminal <i>A</i>
<exists <i>A</i> >	Predicate	Exists symbols
<exists <i>A x</i> >	Predicate	Exists <i>x</i> symbol
<match <i>A</i> >	Predicate	Matches symbol
<is <i>A</i> >	Predicate	Equals symbol
<isa <i>A</i> >	Predicate	Contains symbol
<block <i>e</i> >	Action	Nested scope for <i>e</i>
<local <i>A e</i> >	Action	Isolated local scope for <i>e</i>
<hr/>		
Conditional		
<on <i>c e</i> >	Action	<i>e</i> on <i>c</i> is defined
<on ! <i>c e</i> >	Action	<i>e</i> on <i>c</i> is undefined
<if <i>c</i> >	Predicate	If <i>c</i> is defined
<if ! <i>c</i> >	Predicate	If <i>c</i> is undefined

**Table 1.** Nez operators: "Action" stands for symbol mutators and "Predicate" stands for symbol predicates.

expressions, except that they are greedy and match until the longest position. The  $e_1 e_2$  attempts two expressions  $e_1$  and  $e_2$  sequentially, backtracking to the starting position if either expression fails. The choice  $e_1 / e_2$  first attempt  $e_1$  and then attempt  $e_2$  if  $e_1$  fails. The expression  $\&e$  attempts  $e$  without any character consuming. The expression  $!e$  fails if  $e$  succeeds, but fails if  $e$  fails. Furthermore information on PEG operators is detailed in [8].

The expressiveness of PEGs is mostly equal to that of *deterministic* context-free grammars (such as LALR and LL grammars). In general, PEGs are said to express all LALR grammar languages, which are widely used in a standard parser generator such as Lex/Yacc[14]. In addition, PEGs have many desirable properties that are characterized by:

- *deterministic behavior*, avoiding the dangling if-else problem
- *left recursion-free*, preserving operator precedence

- *scanner-less parsing*, avoiding the extensive use of lexer hacks such as in C++ grammars, and
- *unlimited lookahead*, recognizing highly nested structures such as  $\{a_n b_n c_n \mid n > 0\}$ .

Nez inherits all these properties from PEGs, and has extended features based on AST operators, symbol operators and conditional parsing, as listed in Table 1.

## 2.2 AST Operators

The first extension of parsing expressions in Nez is a flexible construction of AST representations. Each node of an AST contains a substring extracted from the input characters. Nez has adopted a PCRE-like capturing operator, denoted by  $\{ \}$ , to capture the substring. The productions `Int` and `Long` below are examples of capturing a sequence of digits (as defined in `NUM`.)

```
NUM = [0-9]+
Int = { NUM }
Long = { NUM } [L1]
```

A string to be captured is the exact one matched with `NUM`. That is, `Long` accepts the input `0L` but captures only the substring `0`. An empty string may be captured by an empty expression.

Tagging is introduced to distinguish the type of nodes. We can just add a `#`-prefixed tag to each of the nodes.

```
Int = { NUM #Int }
Long = { NUM #Long } [L1]
```

A backquote operator `' '` is a string operator that replaces the captured string with the specified string. Using the empty capture, we are allowed to create an arbitrary node at any point of a parsing expression.

```
DefaultValue = { '0' #Int }
```

The connector  $\$(e)$  is used to make a tree structure by connecting two AST nodes in a parent-child relationship. The prefix  $\$$  is used to specify a child node and append it to a node that is constructed on the left-hand side of  $\$(e)$ . As a result, the tree is constructed as the natural order of left-to-right and top-down parsing.

Basically, the AST operators can transform a sequence of nodes into a tree-formed structure. We can specify that a subtree is nested, flattened, or ignored (by dropping the connector). In the following, we make the construction of a flattened list `[#Add 1 2 3]` and a nested right-associative pair `[#Add 1 [#Add 2 3]]` for the same input `1+2+3`.


```
List = { $(Int) ('+' $(int))* #Add }
Binary = { $(Int) ('+' $(Binary)) #Add }
```

The construction of the left-associative structure, however, suffers from the forbidden left-recursion (e.g, `Binary = Binary '+' Int`). Nez specifically provides a left-folding

```

Expr
= Prod { $left '+' #Add / '-' #Sub $right(Prod) } *
Prod
= Val { $left '*' #Mul / '/' #Div $right(Val) } *
Val
= { [0-9]+ #Int }

```



**Figure 1.** Mathematical basic operators and example of ASTs

constructor `{ $ ... }` that allows a left-hand-constructed node to be contained in a new right-hand node. Accordingly, we can make a precedence-preserving form of ASTs by folding from the repetition.

```
Add = Int { $ '+' $(int) #Add } *
```

The connector and the left-folding operator can associate an optional label for the connected child node. The associated label follows the `$` notion:

```
Add = Int { $left '+' $right(int) #Add } *
```

Note that the AST representation in Nez is a so-called common tree, a sort of generalized data structure. However, we can map a tag to a class and then map labels to fields in a mapped class, which leads to automated conversion of concrete tree objects. Although mapping and type checking are an interesting challenge, they are beyond the scope of this paper.

To the end, we show an example of mathematical basic operators written in Nez. Figure 1 shows some of ASTs constructed when evaluate the nonterminal `Expr`.

### 2.3 Symbol Table

PEGs are very expressive, but they sometimes suffer from insufficient expressiveness for context-sensitive syntax, appearing even in popular programming languages such as:

- Typedef-defined name in C/C++ [8, 9]
- Here document in Perl, Ruby, and many other scripting languages
- Indentation-based code layout in Python and Haskell [1]
- Contextual keywords used in C# and other evolving languages[5]

The problem with context-sensitive syntax is that PEGs are assumed stateless parsing and cannot handle state changes

in the parser context. However, most of the state changes above can be modeled by a *string specialization*; the meaning of words is changed in a certain context. We call such a specialized string a *symbol*. Nez is designed to provide a symbol table and related symbol operators to manage symbols in the parser context.

To illustrate what a symbol is, let us suppose that the production `NAME` is defined:

```
NAME = [A-Za-z] [A-Za-z0-9] *
```

The symbolization operator takes the form of `<symbol A>` to declare that a substring matched at the nonterminal `A` is a symbol. We call such a symbol an *A-specialized symbol*, which is stored with the association of the nonterminal in the symbol table.

Here is a symbolization of `NAME`, where `NAME` first attempts to match and, if matched, `<match NAME>` adds the matched symbol to the symbol table.

```
<symbol NAME>
```

In Nez, a *symbol predicate* is defined to refer to the symbol by the nonterminal name. The following `<match NAME>` is one of the symbol predicates that exactly match the `NAME`-specialized symbol for the input characters.

```
<match NAME>
```

Compared to the nonterminal `NAME`, the result of `<match NAME>` varies, depending on the past result at `<symbol NAME>`. That is, if the `<symbol NAME>` accepts `Apple` previously, `<match NAME>` only accepts `'Apple'`. In this way, the symbol operators handle context-dependent parsing with the state changes in the symbol table.

The uniqueness of Nez is that the state changes are limited to a single symbol table. That is, we are allowed to add multiple symbols with `<symbol NAME>` or different kinds of symbols with another nonterminal in the same table. Nez offers various types of symbol predicates, which can apply different patterns of the symbol references.

- `<exists A>` – checks if the symbol table has any *A*-specialized symbols
- `<exists A s>` – checks if the symbol table has an *A*-specialized symbol that is equal to the given string *s*
- `<match A>` – matches the last *A*-specialized symbol over the input characters
- `<is A>` – compares the last *A*-specialized symbol with an *A*-matched substring.
- `<isa A>` – checks the containment of an *A*-matched substring in a set of *A*-specialized symbols stored in the table.

Note that the two symbol predicates `<match A>` and `<is A>` are quite similar, but they significantly differ in terms of the input acceptance. To illustrate the difference, suppose

```

USERTYPE = [A-Za-z_] !W*
W = [A-ZA-z_0-9]
S = [ \t\r\n]

TypeDef
  = 'typedef' S* TypeName S* <symbol USERNAME> S*

TypeName
  = BuildInTypeName / <isa USERNAME>

BuiltInType
  = 'int' !W / 'long' !W / 'float' !W ...

```

**Figure 2.** Definition of typedef-name syntax with symbol operators

that the previous `<match NAME>` accepts and stores the symbol 'in'. `<match NAME>` accepts the input 'include' (i.e., the input 'clude' is unconsumed), while `<is NAME>` does not because it compares the stored symbol 'in' with the NAME-matched string 'include'.

Figure 2 shows a simplified example of the typedef-name syntax, excerpted from the C grammar. The production `TypeDef` describes the syntax of the typedef statement, which includes the symbolization of `USERTYPE`. In `TypeName`, we first match built-in type names and then match one of the `USERTYPE`-specialized symbols. As a result, we can express context-sensitive patterns in `TypeName`.

### 2.3.1 Scoping

In principle, symbols in the symbol table are globally referable from any production. However, many programming languages have their own scope rule for identifiers, which may require us to restrict the scope of symbols in parallel with a scope construct of the language. Nez provides the explicit scope declaration for that purpose.

Here, we consider a simple case where XML tags need to match closed tags with the same open tags.

```
<A><B> ... </B> </A>
```

Briefly, an XML element can be specified as follows:

```

ELEMENT =
  <symbol TAG> ELEMENT <is TAG>

```

As described above, the symbol `TAG` enables us to ensure the same name in both open and closed tags. However, the `ELEMENT` involves nested symbolization inside, resulting in repeated symbolization. On the other hand, `<is TAG>` refers to the latest `TAG`-specialized symbol. As a result, the `ELEMENT` above can only accept the last tag, such as `<A><B> ... </B> </B>`, which is a not desirable result.

The notation `<block e>` is used to declare a nested scope of symbols. That is, any symbols defined inside `<block e>` are *not* referable outside the block. Here is a scoped version of the `ELEMENT`, which works as expected

```

ELEMENT = <block
  <symbol TAG> ELEMENT <is TAG> >

```

In Nez, we have adopted a single nested scope for multiple nonterminal symbols. In other words, different kinds of symbols are equally controlled with a single scope. One could consider that there is a language that has a more complex scoping rule, while our focus is only on names that directly influence the syntactic analysis. One important exception is an explicit isolation of the specific symbols by `<local A e>`. In this scope, all `A`-specialized symbols before the local scope are isolated and not referable by the subexpression `e`.

## 2.4 Conditional Parsing

The idea of conditional parsing is inspired by *conditional compilation*, where the `#ifdef ... #endif` directives switch the compilation behavior. Likewise, Nez uses `<if>` to switch the parsing behavior, but it differs in that the condition, as well as the semantic predicate, can be controlled by `<on>` in the parser context.

Nez supports multiple parsing conditions, identified by the user-specified condition name. Let `c` be a condition name. A parsing expression `<if c> e` means that `e` is evaluated only if `c` is *true*. We allow the `!` predicate for the negation of the condition `c`. That is, `<if !c>` is a syntax sugar of `!<if c>`. Two expressions `e1` and `e2` are distinctly switched by `<if c> e1 / <if !c> e2`.

In addition, Nez provides the scoped-condition controller for arbitrary parsing expressions. `<on c e>` means that the subexpression `e` is evaluated under the condition that `c` is true, while `<on !c e>` means that the subexpression `e` is evaluated under the condition that `c` is false.

Here is an example of conditional parsing; the acceptance of `Spacing` below depends on the condition `IgnoreNewLine`:

```
Spacing = <if !IgnoreNewLine> [\n\r] / [ \t]
```

The conditions, as well as the symbols, are a global state across productions. That is, the condition `IgnoreNewLine` affects all nonterminals and parsing expressions that involve the `Spacing` nonterminal. The expression `<on IgnoreNewLine e>` is used to declare the static condition of the inner subexpression. The following are the `IgnoreNewLine` version and non-`IgnoreNewLine` version of the `Expr` nonterminal.

```

<on IgnoreNewLine Expr>
<on !IgnoreNewLine Expr>


```

Note that conditional parsing is a Boolean version of the symbol table. Let `C` be an empty expression (i.e., `C = ''`). `<on c e>` is a syntax sugar of `<block <symbol C> e>` and `<if c>` is a syntax sugar of `<exists C ''>`. However, the conditional parsing can be eliminated from arbitrary Nez grammars, since the possible states are at most 2. This is why Nez provides specialized operators for conditional parsing.

$A, B, C \in N$	Nonterminals
$a, b \in \Sigma$	Characters
$x, y, z \in \Sigma^*$	Sequence of characters
$xy$	Concatenation of $x$ and $y$
$e$	Parsing expressions
$A = e$	Production
$T$	Tree node
$[A, x]$	$A$ -specialized symbol $x$
$S = [A, x][A', y]$	State, sequence of labeled symbols
$\epsilon$	empty string or sequence

**Figure 3.** Notations uses throughout this paper

**Input:** `if(a > b) return a; else { return b; }`



**Figure 4.** Pictorial notation of ASTs

### 3. Grammar and Language Design

This section describes the formal definition of Nez. Figure 3 is a list of notations used throughout this section.

#### 3.1 ASTs and Symbol Table

We start by defining the model of the AST representation and the symbol table in Nez.

An AST is a tree representation of the abstract structure of parsed results. The tree is "abstract" in the sense that it contains no unnecessary information such as white spaces or grouping parentheses. Based on the AST operators in Nez, the syntax of the AST representation, denoted by  $v$ , is defined inductively:

$$v ::= \#t[v] \mid \#t['\dots'] \mid vv$$

where  $\#t$  is a tag to identify the type of  $T$  and a captured string written by  $'\dots'$ . A whitespace concatenates two or more nodes as a sequence. Note that we ignore the labeling of subnodes for simplicity.

Here is an example of an AST tree, shown in Figure 4.

```

#If[
  #GreaterThan[
    #Variable['a']
    #Variable['b']
  ]
  #Return[#Variable['a']]
  #Return[#Variable['b']]
]

```

```

#Return[#Variable['b']]
]

```

The initial state of the AST node is an empty tree.  $\text{new}(x)$  is a function that instantiates a new AST node including a given string  $x$ . Let  $v$  be a reference of a node of ASTs. That is, we say  $v = v'$  if  $v'$  is the mutated  $v$  with the following tree manipulation functions:

- $\text{tag}(v, \#t)$  – replaces the tag of  $v$  with the specified  $\#t$
- $\text{replace}(v, x)$  – replaces the string of  $v$  with the specified string  $x$
- $\text{link}(v, v')$  – appends a child node  $v'$  to the parent node  $v$

Next, we will turn to the model of the symbol table. Let  $[A, x]$  denote a symbol  $x$  that is specialized by nonterminal  $A$ . The symbol table  $S$  is represented by a sequence of symbols, such as  $[A, x][B, y]\dots$ . Suppose two following tables:

$$S_1 = [A, x][B, y]$$

$$S_2 = [A, x][B, y][A, z]$$

The initial state of the symbol table is an empty sequence, denoted by  $\epsilon$ . We write  $[A, x] \in S$  for the containment of  $[A, x]$  in  $S$ . In addition, we write  $S[A, z]$  to represent an explicit addition of  $[A, z]$  into  $S$ . ( $SS'$  is a concatenation of  $S$  and  $S'$ ) That is, we say that  $S_1[A, z] = S_2$  and  $S_1 \subset S_2$ . Other operations for  $S$  are represented by the following two functions:

- $\text{top}(S, A)$  – a  $A$ -specialized symbol that is recently added in the table. That is, we say  $\text{top}(S_1, A) = x$  and  $\text{top}(S_2, A) = z$
- $\text{del}(S, A)$  – a new table that removes all  $A$ -specialized entries from  $S$ .

#### 3.2 Grammars

A Nez  $G = (N, \Sigma, P, e_s, \mathcal{T})$  has elements, where  $N$  is a set of nonterminals,  $\Sigma$  is a set of characters,  $P$  is a set of productions,  $e_s$  is a start expression, and  $\mathcal{T}$  is a set of tags.

Each production  $p \in P$  is a mapping from  $A$  to  $e$ ; we write as  $A = e$ , where  $A \in N$  and  $e$  is a parsing expression. Figure 5 is an abstract syntax of parsing expressions in Nez. Due to space constraints, we focus on core parsing expressions. Note that  $e?$  is the syntax sugar of  $e/\epsilon$ ,  $e^*$  the sugar of  $A' = eA'/\epsilon$ , and  $e+$  the sugar of  $ee^*$  (as defined in [8]).

The transition form  $(S, xy, v) \xrightarrow{e} (S', x, v')$  may be read: In symbol table  $S$ , the input sequence  $xy$  is reduced to characters  $y$  while evaluating  $e$  (i.e,  $e$  consumes  $x$ ). Simultaneously, the semantic value  $v$  is evaluated to  $v'$ . The symbol table is unchanged if  $(S, x, v) \xrightarrow{e} (SS', y, v')$  and  $S'$  is empty. Likewise, the AST node is not mutated if  $(S, x, v) \xrightarrow{e} (S', y, v')$  and  $v = v'$ .

$e$	$::=$	$\epsilon$	: empty
		$A$	: nonterminal
		$a$	: character
		$e e'$	: sequence
		$e / e'$	: prioritized choice
		$e^*$	: repetition
		$\&e$	: and predicate
		$!e$	: not predicate
		$\{ e \}$	: constructor
		$\{ \$ e \}$	: left-folding
		$\$(e)$	: connector
		$\#t$	: tag
		$'x'$	: string replacement
		$\langle symbol A \rangle$	: symbol definition
		$\langle exists A \rangle$	: symbol existence
		$\langle exists A x \rangle$	: symbol existence
		$\langle match A \rangle$	: symbol match
		$\langle is A \rangle$	: symbol equivalence
		$\langle isa A \rangle$	: symbol containment
		$\langle block e \rangle$	: nested scope
		$\langle local A e \rangle$	: isolated local scope

**Figure 5.** An abstract syntax of Nez expressions

We write  $\bullet$  for a special failure state, suggesting backtracking to the alternative if it exists. Here is an example failure transition.

$$(S, bx, v) \xrightarrow{a} \bullet \quad (a \neq b)$$

Figure 6 is the definition of the operational semantics of PEG operators. Due to the space constraint, we omit the failure transition case. Instead, undefined transitions are regarded as the failure transition. The semantics of Nez is conservative and shares with the same interpretation of PEGs. Notably, the and-predicate  $\&e$  allows the state transitions of  $S \mapsto SS'$  and  $v \mapsto v'$ , meaning that the lookahead definition of symbols and the duplication of local AST nodes.

Figure 7 and Figure 8 show the semantics of AST operators and symbol operators. The recognition of Nez is AST-independent, because any values  $v$  are not premise for the transitions. Accordingly, an AST-eliminated grammar accepts the same input with the original Nez grammar.

Formally, the language generated by the expression  $e$  is  $L(S, e) = \{x | (S, xy) \xrightarrow{e} (S', y)\}$ , and the language of grammar  $L(G)$  is  $L(G) = \{x | (\epsilon, xy) \xrightarrow{e_s} (S, y)\}$ . Note that the semantic value is unnecessary in the language definition. As suggested in the study of predicated grammars[29], the language class of  $L(G)$  is considered to be *context-sensitive* because predicates can check both the left and right context.

#### 4. Parser Runtime and Implementation

One of the advantages in open grammar is the freedom from parser implementation methods. In fact, the Nez parser tool, which we have developed with the Nez language, can generate three types of differently implemented parsers. First,

$\epsilon$	$(S, x, v) \xrightarrow{\epsilon} (S, x, v)$
$a$	$(S, ax, v) \xrightarrow{a} (S, x, v)$
$A = e$	$\frac{(S, xy, v) \xrightarrow{e} (SS', y, v')}{(S, xy) \xrightarrow{A} (SS', y, v')}$
$e e'$	$\frac{(S, xyz, v) \xrightarrow{e} (SS', yz, v') \quad (SS', yz, v') \xrightarrow{e'} (SS'S'', z, v'')}{(S, x, v) \xrightarrow{ee'} (SS'S'', z, v')}$
$e/e'$	$\frac{(S, xy, v) \xrightarrow{e} (SS', y, v')}{(S, xy, v) \xrightarrow{e/e'} (SS', y, v')}$
	$\frac{(S, xy, v) \xrightarrow{e} \bullet \quad (S, xy, v) \xrightarrow{e'} (SS', y, v')}{(S, xy, v) \xrightarrow{e/e'} (SS', y, v')}$
$\&e$	$\frac{(S, xy, v) \xrightarrow{e} (SS', y, v')}{(S, xy, v) \xrightarrow{\&e} (SS', xy, v')}$
$!e$	$\frac{(S, x, v) \xrightarrow{e} \bullet}{(S, x, v) \xrightarrow{!e} (S, x, v)}$

**Figure 6.** Semantics of PEG operators

$\{ e \}$	$\frac{(S, xy, v) \xrightarrow{e} (SS', y, v')}{(S, xy, T) \xrightarrow{\{ \$ e \}} (SS', y, \mathbf{new}(x))}$
$\$(e)$	$\frac{(S, xy, v) \xrightarrow{e} (S', y, v')}{(S, xy, v) \xrightarrow{\$(e)} (S', y, \mathbf{link}(v, v'))}$
$\{ \$ e \}$	$\frac{(S, xy, v) \xrightarrow{e} (S', y, v')}{(S, xy, v) \xrightarrow{\{ \$ e \}} (S', y, \mathbf{link}(\mathbf{new}(x), v))}$
$\#t$	$(S, x, v) \xrightarrow{\#t} (S, x, \mathbf{tag}(v, \#t))$
$'x'$	$(S, x, v) \xrightarrow{'x'} (S, x, \mathbf{replace}(v, x))$

**Figure 7.** Semantics of AST operators

as well as traditional parser generators, Nez produces parser source code written in the target language of the parser application. Second, Nez provides a grammar translator into other existing PEG-based grammars, (such as Rats!, PEGjs, or PEGTL). Third, Nez itself works as an efficient parser interpreter that loads a grammar at runtime. In this section, we briefly describe these three implementations.

$$\begin{array}{c}
\boxed{\langle \text{symbol } A \rangle} \\
\frac{(S, xy, T) \xrightarrow{A} (S, y, T')}{(S, xy, T) \xrightarrow{\langle \text{symbol } A \rangle} (S[A, x], y, T')} \\
\\
\boxed{\langle \text{block } e \rangle} \\
\frac{(S, xy, T) \xrightarrow{e} (S', y, T')}{(S, xy, T) \xrightarrow{\langle \text{block } e \rangle} (S, y, T')} \\
\\
\boxed{\langle \text{local } A e \rangle} \\
\frac{S_{\bar{A}} = \text{del}(S, A) \quad (S_A, xy, T) \xrightarrow{e} (S', y, T')}{(S, xy, T) \xrightarrow{\langle \text{local } A e \rangle} (S, y, T')} \\
\\
\boxed{\langle \text{exists } A \rangle} \\
\frac{\exists z[A, z] \in S}{(S, x, T) \xrightarrow{\langle \text{exists } A \rangle} (S, x, T)} \\
\\
\boxed{\langle \text{exists } A x \rangle} \\
\frac{\exists[A, z] \in S}{(S, x, T) \xrightarrow{\langle \text{exists } A x \rangle} (S, x, T)} \\
\\
\boxed{\langle \text{match } A \rangle} \\
\frac{\text{top}(S, A) = z}{(S, zx, T) \xrightarrow{\langle \text{match } A \rangle} (S, x, T)} \\
\\
\boxed{\langle \text{is } A \rangle} \\
\frac{(S, zx, T) \xrightarrow{A} (S, x, T') \quad \text{top}(S, A) = z}{(S, zx, T) \xrightarrow{\langle \text{is } A \rangle} (S, x, T')} \\
\\
\boxed{\langle \text{isa } A \rangle} \\
\frac{(S, zx, T) \xrightarrow{A} (S, x, T') \quad \exists[A, z] \in S}{(S, zx, T) \xrightarrow{\langle \text{isa } A \rangle} (S, x, T')}
\end{array}$$

**Figure 8.** Semantics of Nez’s symbol operators

#### 4.1 Parser Runtime and Parser Generation

The parser runtime required in Nez is fundamentally lightweight and portable. Nez parsers, as well as PEG parsers, can be implemented with recursive descent parsing with backtracking. All productions are simply implemented with parse functions that compute on the input characters, and then nonterminal calls are computed by function calls. Backtracking is simply implemented over the call stacks.

Notoriously, backtracking might cause an exponential time parsing in the worst cases, while packrat parsing[7] is well established to guarantee the linear time parsing with PEGs. The idea behind packrat parsing is a memoization whereby all results of nonterminal calls are stored in the memoization table to avoid redundant nonterminal calls. Although the memoization table may require some complex-

ity for its efficiency, we use a simple and constant-memory memoization table, presented in [21].

In addition to the standard PEG parser runtime, Nez parsers require two additional state management to handle ASTs and symbols. Note that all conditional expressions are eliminated by creating true/false versions of productions for each conditional flags. This elimination produces  $2^N$  versions for  $N$  different conditions in the worst case. Our empirical study (as described in Section 5) suggests that practical grammars involves not so many conditions that it would cause such an exponential increase of eliminated productions.

The AST construction runtime provides a parser with APIs that are based on AST operators in Nez. Importantly, Nez parsers are speculative parsing in a way that some of the AST operations may be discarded when backtracking. To handle the consistency of ASTs, we support transactional operations (such as commit and abort) and all operations are stored as logs to be committed. This transactional structure can be easily implemented with stack-based logging, and the AST construction is partially aborted at the backtracking time. Note that efficient packrat parsing with ASTs requires additional transactional management for ASTs. A detailed mechanism is reported in [22].

The symbol table requires another state management. As with the operation logs in the AST constructor, we use a stack-based structure to control both the symbol scoping and backtracking consistency. The symbol table runtime provides a parser with APIs that enable adding symbols, eliminating symbols, and testing symbols to match the input.

Originally, we write the AST runtime and the symbol table runtime in Java. There is no use of functional data structure; instead, both operation logs and symbols are stored in a linked list, which is available in any programming language. Indeed, we have already ported Nez runtime into several languages, including C and JavaScript. The C version of the AST runtime is at most 500 lines in code and the symbol table runtime is at most 200 lines in code, suggesting its high portability.

#### 4.2 Grammar Translation

A Nez grammar is specified with a declarative form of the AST operators and symbol operators, which are performed as a kind of action in the parsing context. These actions are limited in number and can be statically translated into code fragments written in any programming language. This means that a Nez grammar is convertible into PEGs with embedded semantic action code.

Grammar translation is another approach to the Nez parser implementation. The advantage is that we enjoy optimized parsers that existing PEG tools generate. On the other hand, impedance mismatch occurs between two grammars. In particular, many existing PEG-based parser generators have their own supports for AST representations. For example, PEGjs[24] produces a JSON object as a semantic value



```

e / e'      {symbolTable.startBlock()}
            e { symbolTable.commit() }
            / {symbolTable.abort()} e'
<symbol A>  { symbolTable.add(A, capture(A)) }
<block e>   { symbolTable.startBlock() }
            e
            { symbolTable.abortBlock() }
<local A e> { symbolTable.startBlock().mask(A) }
            e
            { symbolTable.abortBlock() }
<exists A>  &{ symbolTable.count(A) > 0 }
<exists A x> &{ symbolTable.top(A) == x }
<match A>   &{ match(symbolTable.top(A)) }
<is A>      &{ symbolTable.top(A) == capture(A) }
<isa l>     &{ symbolTable.contains(A, capture(A)) }

```

**Figure 9.** Implementing Nez symbol operators with action code

of all nonterminal calls. In those cases, the translation of AST operators is unnecessary in practice. The symbol operators are always translatable if the symbol runtime that run on a host language is readily available.

Figure 9 shows an example of converting symbol operators into semantic actions `{ ... }` and semantic predicates `&{ ... }`. The rollback of the symbol table at the backtracking time is automated before attempting alternatives. As shown, a Nez grammar is recursively convertible into PEGs with action code using the Nez runtime.

Currently, the Nez tool provides the grammar translation into PEGjs, PEGTL, and LPeg, although the translation of the AST construction might be restricted due to the impedance mismatch.

### 4.3 Virtual Parsing Machine

The simplicity of PEGs makes it easier to achieve dynamic parsing in a way that a parser interpreter loads a grammar at runtime to parse the input. Since dynamic parsing requires no source compilation process, dynamic parsing makes it easier to import the Nez parser functionality as a parser library into parser applications. We consider that dynamic parsing is better suitable for many use cases of open grammars. Accordingly, the standard Nez parser is based on dynamic parsing and then implemented on top of a virtual parsing machine, an efficient implementation of dynamic parsing.

A Nez parsing machine is a stack-based virtual machine that runs with a set of bytecode instructions, specialized for PEG, AST, symbol, and memoization operators. A parsing expression  $e$  in Nez is converted to bytecode by using a compile function  $\tau(e, L)$ , where  $L$  is a label representing the next code point. The byte compilation is simply inductive, while Nez supports additional super instructions to generate optimized bytecode.

Currently, the Nez parser tool includes a bytecode compiler and a parsing machine, both written in Java. In addition,

Language	#G	#P	Symbols and Conditions
C	102	102	TypeName
C#5.0	897	1325	<i>AwaitKeyword</i> , <i>Global</i>
CoffeeScript	121	121	Indent
Java8	185	185	
JavaScript	153	166	<i>InOperator</i>
Konoha	163	190	<i>Offside</i>
Lua5	102	105	<i>MultiLineBracketString</i>
Python3	139	155	Indent, <i>Offside</i>
Ruby	420	580	<i>Delimiter</i> , <i>Primary</i> , <i>DoExpr</i>

**Table 2.** Grammar summary: the typewriter font stands for a symbol name and the *italic* font stands for a condition name

the parsing machine is ported into C and JavaScript. Especially, the C version of the Nez parsing machine is highly optimized with indirect thread code and SEE4 instructions. The performance evaluation is studied in Section 6.

## 5. Case Studies and Experiences

We have developed many Nez grammars, ranging from programming languages to data formats. All developed grammars are available online at <http://github.com/nez-peg/nez-grammar>. Table 2 shows a summary of the major developed grammars. The column labeled “#G” indicates the number of *defined* productions in a grammar, while the column labeled “#P” indicates the number of *parser* productions through eliminating the conditional operators. The same numbers of productions in #G and #P means no use of conditional parsing. The column labeled “Symbols and Conditions” indicates nonterminals and conditions that are used in the grammar.

Here are some short comments on each of the interesting developed grammars.

- C – Based on two PEG grammars written in Mouse. The semantic code embedded to handle the typedef-defined name is converted into the TypeDef symbol. The nested typedef statements are not implemented.
- C# – Developed from scratch, referencing C#5.0 Language Specification (written in a natural language). The condition *AwaitKeyword* is used to express contextual keyword *await*.
- Java8 – Ported from Java8 grammar written in ANTLR4<sup>1</sup>. Java can be specified without any Nez extensions.
- JavaScript – Based on JavaScript grammar for PEG.js<sup>2</sup>. We simplify the grammar specification using the parsing condition *InOperator* that distinguishes expression rules from the *for/in* context.

<sup>1</sup> <https://github.com/antlr/grammars-v4/blob/master/java8/Java8.g4>

<sup>2</sup> <https://github.com/pegjs/pegjs/blob/master/examples/javascript.pegjs>

- Konoha – Developed from scratch. Konoha is a statically typed scripting language, which we have designed in [20]. The most syntactic constructs come from Java, but we use the parsing condition `SemicolonInsertion` for expressing the conditional end of a statement. More importantly, Konoha is a language implementation that uses ASTs constructed by the Nez parser.
- Python3 – Ported from Python 3.0 abstract grammar<sup>3</sup>. The `Indent` is used to express the indentation-based code block by capturing white spaces.
- Ruby – Developed from scratch. The `Delimiter` is used to express the context-sensitive delimiting identifier in the `Here` document.
- Haskell – Postponed. We didn’t express the indent symbol for Haskell’s code layout since the indentation starts in the middle of expressions. This is mainly because of a limitation of a symbol operator in Nez; if Nez provided Haskell-specialized symbol operator such as `<haskell-indent>`, we could handle the indentation-based code layout through the symbol tables.

Two findings are confirmed throughout our case studies. First, many of the programming languages, as listed in Table 2, can not be well expressed with pure PEG. The introduction of symbol tables significantly improve the expressiveness, although some corner cases might remain. Even in corner cases, we consider that the symbol table approach is still applicable. Second, the parsing condition, although it does not directly improve the expressiveness of PEGs, simplifies the specification task. The growth of `#P/#G` indicates how the conditional parsing reduce the grammar developer’s tasks. If the conditional parsing were not supported in Nez, the grammar developers have to specify as many productions as indicated at `#P`. Furthermore detailed discussions and earlier reports on grammar developments with examples are available in [23].

All grammars we have developed include the AST construction. At the moment, the only working example of a language implementation with Nez is with Konoha, a reimplemented version of [20] as a JVM-based language implementation, including type checking and code generation on top of the ASTs that are constructed by a Nez parser. Since each AST node records the source location when capturing a string, Konoha can carry out error reporting, although it is very primitive.

## 6. Performance Study

This section is an experimental report on our performance and comparative study. We start by describing common setups. All tests in this paper are measured on DELL XPS-8700, with 3.4GHz Intel Core i7-4770, 8GB of DDR3 RAM, and running on Linux Ubuntu 14.04.1 LTS. All C programs

are compiled with clang-3.6, all Java parsers run on Oracle JDK 1.8, and all JavaScript parsers run on node.js version 5.0. Tests are run several times, and we record the average time for each iteration. The parsing time is measured in millisecond by `System.nanoTime()` in Java APIs and `gettimeofday()` in Linux. Tested source files are randomly collected from major open source repositories. All tested Nez grammars are available online as described in Section 5. For convenience, we label tested grammars: `java.nez`, `js.nez`, `python3.nez`, `xml.nez`, and `xmlsym.nez`. Note that `xmlsym.nez` is a symbol version of `xml.nez`, which checks whether the closed tag is equal to the open tag, as described in Section 2. The two tested grammars `python3.nez` and `xmlsym.nez` contain symbol operators in a grammar.

Using Nez grammars, we generate Nez parsers for Java, C, and JavaScript. The generated parsers are grouped by the base implementation with the following labels:

- `nez-j` - an original Nez parsing machine on JVM, incorporated as a standard Nez parser,
- `nez-c` - a C-ported Nez parsing machine,
- `nez-js` - a JavaScript-ported Nez parsing machine, and
- `cnez` - generated C parsers from the Nez tool


Linear time parsing is a primary concern of parser implementations. Figure 10 shows the parsing time plotted against file sizes in, respectively, Java, JavaScript, and Python. The lines are the linear regression line fit of `nez-j`, `cnez`, and `nez-c` parsers, suggesting the parsing time is linear. Since Nez parsers are integrated with packrat parsing, we have not observed any super-linear parser behavior throughout this experiment.

To make detailed analysis, we have chosen some of tested source files from the collection. The chosen data are labeled as:


- `processing` – a JavaScript version of the Processing runtime
- `jquery` – a very popular JavaScript library. We used a compressed source.
- `java20` – 20 Java files collection, consisting of relatively large files (more than 10KB size). The parsing time is recorded as the summation of all 20 parsing times.
- `xmark` – a synthetic and scalable XML files that are provided by XMark benchmark program [32]. We used a newly generated file in 10MB.

Figure 11 shows the parse time of Nez parsers. The data point `parse` indicates the parsing time including the AST construction, while `match` indicates the parsing time without the AST construction. The time scale is plotted as log base 10 due to the variety of results. The `cnez` parsers are fastest as easily predicted. In comparisons among the parsing machines, the `nez-c` parser is almost 2.5 times faster than `nez-java`, and almost 10 times faster than the `nez-js`.

<sup>3</sup> <https://docs.python.org/3.0/library/ast.html>



**Figure 10.** Parsing time of the open source files for Java, JavaScript, and Python against the input size plotted as log-log base 10.



**Figure 11.** Performance comparison of differently implemented Nez parsers: nez-j, nez-c, cnez, and nez-js

Next, we will turn to the performance comparison with existing standard parsers. Here, we have chosen Yacc[14], ANTLR4[30], Rats![9], and PEGjs[24], since they notably produce efficient parsers and are accepted in many serious projects. We set up these parser generators as follows.

- `yacc` – a standard LALR parser generator for C. In this experiment, we suffered from the limited availability of grammars, and could only run a JavaScript parser without AST construction.
- `antlr4` – a standard parser generator based on the extended  $LL(k)$  with some packrat parsing feature. All tested grammars are derived from the official ANTLR4 grammar repository.
- `rats!` – a PEG-based parser generator[9] for Java. Java grammar is derived from a part of the `xtc` implementation, while JavaScript parser is generated from a Nez grammar `js.nez` by the grammar translation.

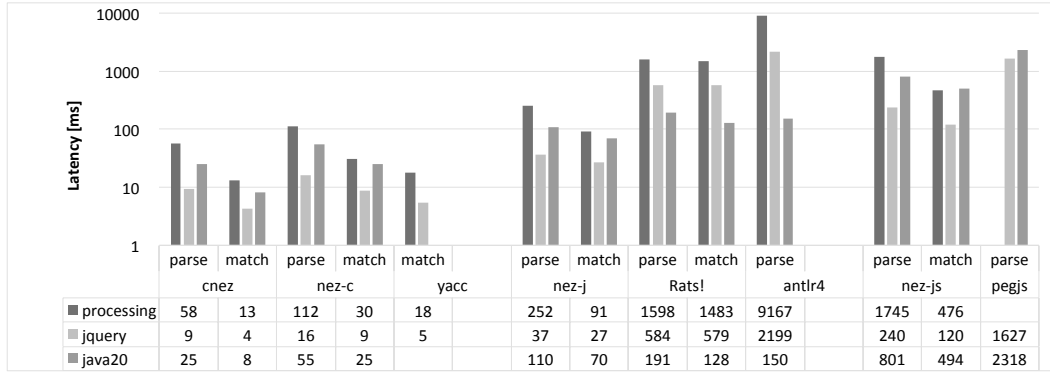
- `pegjs` – a PEG-based parser generator[24] for JavaScript. JavaScript grammar is derived from its official site, and Java grammar is translated from `java.nez` by Nez tools.

Figure 12 shows the parsing time of processing, jquery, and java20 in each of the examined parsers. All Nez parsers show better performance with competitive parsers in each of the host environments. In ANTLR4 and Rats!, some of results are surprisingly bad. We run fairly in terms of the time measurement and the JIT condition, but we are still not sure that Rats! and ANTLR4 parsers are best optimized in the experiment. As a result, we conclude that Nez parsers achieve *competitive* performance compared to these practical parser generators. Likewise, we compare Nez parsers for xmark with standard XML parsers such as libxml and xerces, which are highly optimized by hand. Figure 13 shows the comparison of parsing time in xmark. The data points match and parse in XML parsers correspond to SAX and DOM parsers, respectively. Although we have observed the overhead of AST construction in Nez parsers, they are also competitive.


Finally, we focus on the performance effect of symbol operators. Figure 14 shows the performance comparison of `xml.nez` and `xmlsym.nez`. We confirm that the costs of symbol operators, requiring `<block>`, `<symbol>`, and `<is>` at each closed tag, are small at the acceptable level.

## 7. Related Work


Developing parsers is ubiquitous in many applications, and hand-written parsers are obviously prone to errors. Parser generation from a formal grammar specification has a long history in programming language research. A common key challenge is balancing the expressiveness of formal grammars and how to implement them in practical manners. In the early 1970s, Stephen C. Johnson developed Yacc[14], based on a LALR(1) grammar with embedded C code to perform user-defined actions. Since Yacc successfully hits a sweet spot between the expressiveness and practical parser gener-



**Figure 12.** Performance comparison of processing, jquery, and java20 in various parsers. The pegjs parser cannot end parsing process due to its out of memory errors.



**Figure 13.** Parsing time in xmark



**Figure 14.** Effects of symbol operators

ation, the use of action code has been broadly adopted in most of the major parser generators [16, 18, 29] across various grammar foundations, including  $LL(k)$ , GLR, and PEG.

In parallel, the problems of action code have been pointed out, for example, the lack of grammar reuse[28], decreased maintainability [19], and ad hoc behavior[1]. There have been several attempts to decrease the disadvantage of action code. Terence Parr, the author of ANTLR, has proposed a prototype grammar, inspired by the revision control system[28]. Elkhound[25], for instance, allows the users to write action code written in multiple programming lan-

guages. Despite these attempts, grammar reuse is still limited and arbitrary action inevitably requires porting to other languages.

The idea of open grammars has been strongly inspired by the article: *Pure and declarative syntax definition: paradise lost and regained* [18]. SDF+ADF and Stratego [4, 6] intend their users to write a syntactic analysis using an algebraic transformation. The actions are not arbitrary, but they are expressive enough. However, the actions differ from those required in Nez, since SDF is based on Generalized LR [33], where actions are mainly used to resolve grammar ambiguity. Nez, on the other hand, is unambiguous as a PEG, and symbol operators are used for handling state changes.

Data-dependent grammars [3, 13] share similar ideas in terms of recognizing context-sensitive syntax. Briefly, these grammars use a variable to be bound to a parsed result, which seems to be the addition of symbols in Nez. Data-dependent grammars are more expressive in a way that they express more complex condition such as  $([n>0] OCTET\{n:=n1\})$ , while Nez is more declarative and leads to better readability. In addition, the better expressiveness of data-dependent grammars suggests a future promising extension of Nez, including numerical values in the symbol table.

Among many formal grammars such as  $LALR(k)$  and  $LL(k)$ , PEGs are simple and seemingly a suitable foundation for implementing portable parser runtime. In the remainder of this section, we would like to focus on PEG-specific related work.

**Extensions.** With Rats!, Robert Grimm has shown an elegant integration of PEGs with action code despite its speculative parsing[9]. Many other PEG-based parser generators [10, 24, 31] have adopted the semantic action approach for expressing syntax constructs that PEGs can not recognize. Yet, there have been several declarative attempts for PEG extensions. Adams has recently formalized Indent-Sensitive CFGs [1] for the indentation-based code layout and has extended the PEG-version of IS-CFGs [2]. Iwama et al have

extended PEGs with the  $e U e$  operator to combine a black-boxed parser for a natural language[12]. Compared to these existing extensions, Nez more broadly covers various parsing aspects.

**AST Construction.** Traditionally, a grammar developer usually writes action code to construct some forms of ASTs as a semantic value[15]. However, writing the AST construction is tedious, and many parser generators have some annotation-based supports to automate the AST construction. The underlying idea is to filter trees that are derived from structural parse results of nonterminal calls. However, the filtering approach is problematic in PEGs since PEGs disallows the left-recursive structure of derived trees. Recently, handling the left recursion in PEGs has been established in [27], although another annotation is needed for preserving operator precedence. Nez allows users to specify an explicit structure of ASTs (including tags and labels), resulting in a more general string-to-tree transducer. Notably, a similar capturing idea appears in LPeg[11], a PEG-based pattern matching, but Nez enables more structural complexity for capturing structured data.

**Parser Runtime.** Nez parser runtime is based on many of previous works reported in the literature[9, 21, 26, 31]. In particular, the state management in speculative parsing is built on the partial transaction management of Rats!. The Nez virtual machine is said to be an extended version of the LPeg machine. Notably, the Nez machine has instruction supports for packrat parsing [7], which is a major lack in the LPeg machine. The originality of Nez parsers is that we combine these established techniques in a way that the parser runtime is still simple and portable.

## 8. Conclusion

Nez is a simple, portable, and declarative grammar specification language. As an alternative to action code, Nez provides a small set of extended PEG operators that allow the grammar developer to make AST constructions, as well as context-sensitive and conditional parsing. Due to these extensions, Nez can recognize a language that cannot be well expressed by pure PEGs. In addition, the Nez parser runtime is also simple and portable. Based on the Java-based implementation, we have developed C and JavaScript ports, suggesting that the Nez parser runtime is implementable in any modern programming languages. More importantly, Nez parsers achieve practical and competitive performance in each of the ported environments.

This paper is an initial report on Nez and the idea of open grammars. Further evaluations are obviously necessary on the expressiveness of Nez with extensive case studies. To achieve the aim of open grammars, a lot of interesting challenges remain. Future work that we will investigate includes grammar modularity and inference, a tree checker for better connectability of parser applications, and a DFA-based machine for more efficient parsing. In the end, we hope that Nez

parsers will be available in many programming languages. Our developed tools and grammars are available online at <http://nez-peg.github.io/>.

## Acknowledgments

The C version of Nez parsers (cnez and virtual parsing machines) have been developed by Masahiro Ide and Shun Honda. The JavaScript version has been developed by Takeru Sudo. Finally, the anonymous reviewers are acknowledged for their helpful suggestions for improving an earlier draft of this paper.

## References

- [1] MICHAEL D. ADAMS. Principled parsing for indentation-sensitive languages: Revisiting landin's offside rule. In *Proceedings of the 40th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL '13*, pages 511–522, New York, NY, USA, 2013. ACM. ISBN 978-1-4503-1832-7. .
- [2] MICHAEL D. ADAMS & ÖMER S. AĞACAN. Indentation-sensitive parsing for parsec. In *Proceedings of the 2014 ACM SIGPLAN Symposium on Haskell, Haskell '14*, pages 121–132, New York, NY, USA, 2014. ACM. ISBN 978-1-4503-3041-1. .
- [3] ALI AFROOZEH & ANASTASIA IZMAYLOVA. One parser to rule them all. In *2015 ACM International Symposium on New Ideas, New Paradigms, and Reflections on Programming and Software (Onward!)*, Onward! 2015, pages 151–170, New York, NY, USA, 2015. ACM. ISBN 978-1-4503-3688-8. .
- [4] MARTIN BRAVENBOER & EELCO VISSER. Concrete syntax for objects: Domain-specific language embedding and assimilation without restrictions. In *Proceedings of the 19th Annual ACM SIGPLAN Conference on Object-oriented Programming, Systems, Languages, and Applications, OOPSLA '04*, pages 365–383, New York, NY, USA, 2004. ACM. ISBN 1-58113-831-8. .
- [5] MARTIN BRAVENBOER, ÉRIC TANTER, & EELCO VISSER. Declarative, formal, and extensible syntax definition for aspectj. In *Proceedings of the 21st Annual ACM SIGPLAN Conference on Object-oriented Programming Systems, Languages, and Applications, OOPSLA '06*, pages 209–228, New York, NY, USA, 2006. ACM. ISBN 1-59593-348-4. .
- [6] MARTIN BRAVENBOER, KARL TRYGVE KALLEBERG, ROB VERMAAS, & EELCO VISSER. Stratego/xt 0.17. a language and toolset for program transformation. *Sci. Comput. Program.*, 72(1-2):52–70, June 2008. ISSN 0167-6423. .
- [7] BRYAN FORD. Packrat parsing:: Simple, powerful, lazy, linear time, functional pearl. In *Proceedings of the Seventh ACM SIGPLAN International Conference on Functional Programming, ICFP '02*, pages 36–47, New York, NY, USA, 2002. ACM. ISBN 1-58113-487-8. .
- [8] BRYAN FORD. Parsing expression grammars: A recognition-based syntactic foundation. In *Proceedings of the 31st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL '04*, pages 111–122, New York, NY, USA, 2004. ACM. ISBN 1-58113-729-X. .

- [9] ROBERT GRIMM. Better extensibility through modular syntax. In *Proceedings of the 2006 ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI '06*, pages 38–51, New York, NY, USA, 2006. ACM. ISBN 1-59593-320-4. .
- [10] COLIN HIRSCH & DANIEL FREY. Parsing expression grammar template library. <https://code.google.com/p/pegtl/>, 2014.
- [11] ROBERTO IERUSALIMSCHY. A text pattern-matching tool based on parsing expression grammars. *Softw. Pract. Exper.*, 39(3):221–258, March 2009. ISSN 0038-0644. .
- [12] FUTOSHI IWAMA, TAIGA NAKAMURA, & HIRONORI TAKEUCHI. Constructing parser for industrial software specifications containing formal and natural language description. In *Proceedings of the 34th International Conference on Software Engineering, ICSE '12*, pages 1012–1021, Piscataway, NJ, USA, 2012. IEEE Press. ISBN 978-1-4673-1067-3.
- [13] TREVOR JIM, YITZHAK MANDELBAUM, & DAVID WALKER. Semantics and algorithms for data-dependent grammars. In *Proceedings of the 37th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL '10*, pages 417–430, New York, NY, USA, 2010. ACM. ISBN 978-1-60558-479-9. .
- [14] STEPHEN C JOHNSON. *Yacc: Yet another compiler-compiler*, volume 32. ell Laboratories Murray Hill, NJ, 1975.
- [15] ADRIAN JOHNSTONE & ELIZABETH SCOTT. Tear-insert-fold grammars. In *Proceedings of the Tenth Workshop on Language Descriptions, Tools and Applications, LDTA '10*, pages 6:1–6:8, New York, NY, USA, 2010. ACM. ISBN 978-1-4503-0063-6. .
- [16] MANOHAR JONNALAGEDDA, THIERRY COPPEY, SANDRO STUCKI, TIARK ROMPF, & MARTIN ODERSKY. Staged parser combinators for efficient data processing. In *Proceedings of the 2014 ACM International Conference on Object Oriented Programming Systems Languages & Applications, OOPSLA '14*, pages 637–653, New York, NY, USA, 2014. ACM. ISBN 978-1-4503-2585-1. .
- [17] LENNART C. L. KATS, ROB VERMAAS, & EELCO VISSER. Integrated language definition testing: Enabling test-driven language development. In Kathleen Fisher, editor, *Proceedings of the 26th Annual ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA 2011)*, Portland, Oregon, USA, 2011. ACM.
- [18] LENNART C.L. KATS, EELCO VISSER, & GUIDO WACHSMUTH. Pure and declarative syntax definition: Paradise lost and regained. In *Proceedings of the ACM International Conference on Object Oriented Programming Systems Languages and Applications, OOPSLA '10*, pages 918–932, New York, NY, USA, 2010. ACM. ISBN 978-1-4503-0203-6. .
- [19] PAUL KLINT, TIJS VAN DER STORM, & JURGEN VINJU. On the impact of dsl tools on the maintainability of language implementations. In *Proceedings of the Tenth Workshop on Language Descriptions, Tools and Applications, LDTA '10*, pages 10:1–10:9, New York, NY, USA, 2010. ACM. ISBN 978-1-4503-0063-6. .
- [20] KIMIO KURAMITSU. Konoha: implementing a static scripting language with dynamic behaviors. In *S3 '10: Workshop on Self-Sustaining Systems*, pages 21–29, New York, NY, USA, 2010. ACM. ISBN 978-1-4503-0491-7. .
- [21] KIMIO KURAMITSU. Packrat parsing with elastic sliding window. *Journal of Information Processing*, 23(4):505–512, 2015. .
- [22] KIMIO KURAMITSU. Fast, flexible, and declarative construction of abstract syntax trees with PEGs. *Journal of Information Processing*, 24(1):123–131, 2016.
- [23] KIMIO KURAMITSU & TETSURO MATSUMURA. A declarative extension of parsing expression grammars for recognizing most programming languages. *Journal of Information Processing*, 24(2):256–264, 2016.
- [24] DAVID MAJDA. Peg.js - parser generator for javascript, 2015.
- [25] SCOTT MCPEAK & GEORGE C. NECULA. Elkhound: A fast, practical glr parser generator. In Evelyn Duesterwald, editor, *Compiler Construction*, volume 2985 of *Lecture Notes in Computer Science*, pages 73–88. Springer Berlin Heidelberg, 2004. ISBN 978-3-540-21297-3.
- [26] SÉRGIO MEDEIROS & ROBERTO IERUSALIMSCHY. A parsing machine for pegs. In *Proceedings of the 2008 Symposium on Dynamic Languages, DLS '08*, pages 2:1–2:12, New York, NY, USA, 2008. ACM. ISBN 978-1-60558-270-2. .
- [27] SÉRGIO MEDEIROS, FABIO MASCARENHAS, & ROBERTO IERUSALIMSCHY. Left recursion in parsing expression grammars. *Sci. Comput. Program.*, 96(P2):177–190, December 2014. ISSN 0167-6423. .
- [28] TERENCE PARR. The reuse of grammars with embedded semantic actions. In *Proceedings of the 2008 The 16th IEEE International Conference on Program Comprehension, ICPC '08*, pages 5–10, Washington, DC, USA, 2008. IEEE Computer Society. ISBN 978-0-7695-3176-2. .
- [29] TERENCE PARR & KATHLEEN FISHER. Ll(\*): The foundation of the antlr parser generator. In *Proceedings of the 32nd ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI '11*, pages 425–436, New York, NY, USA, 2011. ACM. ISBN 978-1-4503-0663-8. .
- [30] TERENCE PARR, SAM HARWELL, & KATHLEEN FISHER. Adaptive ll(\*) parsing: The power of dynamic analysis. In *Proceedings of the 2014 ACM International Conference on Object Oriented Programming Systems Languages & Applications, OOPSLA '14*, pages 579–598, New York, NY, USA, 2014. ACM. ISBN 978-1-4503-2585-1. .
- [31] ROMAN R. REDZIEJOWSKI. Parsing expression grammar as a primitive recursive-descent parser with backtracking. *Fundam. Inf.*, 79(3-4):513–524, August 2007. ISSN 0169-2968.
- [32] ALBRECHT SCHMIDT, FLORIAN WAAS, MARTIN KERSTEN, MICHAEL J. CAREY, IOANA MANOLESCU, & RALPH BUSSE. Xmark: A benchmark for xml data management. In *Proceedings of the 28th International Conference on Very Large Data Bases, VLDB '02*, pages 974–985. VLDB Endowment, 2002.
- [33] EELCO VISSER. *Syntax Definition for Language Prototyping*. PhD thesis, University of Amsterdam, September 1997.