

Extend

## Language Reference Manual

Ishaan Kolluri(isk2108), Kevin Ye(ky2294) Jared Samet(jss2272), Nigel Schuster(ns3158)

December 20, 2016

# Contents

|          |   |          |
|----------|---|----------|
| <b>1</b> | <b>Introduction to Extend</b>           | <b>3</b> |
| <b>2</b> | <b>Structure of an Extend Program</b>   | <b>3</b> |
| 2.1      | Import Statements . . . . .             | 3        |
| 2.2      | Function Definitions . . . . .          | 4        |
| 2.3      | Global Variables . . . . .              | 4        |
| 2.4      | External Library Declarations . . . . . | 4        |
| 2.5      | main function . . . . .                 | 5        |
| 2.6      | Scoping and Namespace . . . . .         | 5        |
| 2.7      | Identifiers . . . . .                   | 5        |
| <b>3</b> | <b>Types and Literals</b>               | <b>6</b> |
| 3.1      | Primitive Data Types . . . . .          | 6        |
| 3.2      | Ranges . . . . .                        | 6        |
| 3.2.1    | Range Literals . . . . .                | 7        |
| <b>4</b> | <b>Expressions</b>                      | <b>7</b> |
| 4.1      | Arithmetic Operators . . . . .          | 8        |
| 4.2      | Boolean Operators . . . . .             | 9        |
| 4.3      | Conditional Expressions . . . . .       | 11       |
| 4.3.1    | Ternary Expressions . . . . .           | 11       |
| 4.3.2    | Switch Expressions . . . . .            | 11       |
| 4.4      | Additional Operators . . . . .          | 12       |
| 4.5      | Function Calls . . . . .                | 13       |
| 4.6      | Range Expressions . . . . .             | 13       |
| 4.6.1    | Slices . . . . .                        | 14       |
| 4.6.2    | Selections . . . . .                    | 14       |

|          |  |           |
|----------|--|-----------|
| 4.6.3    | Corresponding Cell . . . . .                                   | 14        |
| 4.6.4    | Selection Examples . . . . .                                   | 14        |
| 4.7      | Precedence Expressions . . . . .                               | 15        |
| <b>5</b> | <b>Functions</b>   | <b>16</b> |
| 5.1      | Format . . . . .   | 16        |
| 5.2      | Variable Declarations . . . . .                                | 16        |
| 5.3      | Formula Assignment . . . . .                                   | 17        |
| 5.3.1    | Combined Variable Declaration and Formula Assignment . . . . . | 17        |
| 5.3.2    | Formula Assignment Errors . . . . .                            | 17        |
| 5.4      | Parameter Declarations . . . . .                               | 18        |
| 5.5      | Application on Ranges . . . . .                                | 18        |
| 5.6      | Lazy Cell Evaluation and Circular References . . . . .         | 19        |
| 5.7      | External Libraries . . . . .                                   | 20        |
| <b>6</b> | <b>Standard Library Reference</b>                              | <b>20</b> |
| 6.1      | File I/O . . . . .   | 20        |
| 6.2      | Math Functions - Imported straight from C . . . . .            | 20        |
| 6.3      | Math Functions - Not imported from C . . . . .                 | 20        |
| 6.4      | String Functions . . . . .                                     | 21        |
| 6.5      | Plotting . . . . .   | 21        |
| 6.6      | Range Functions . . . . .                                      | 21        |

# 1. Introduction to Extend

Extend is a dynamically-typed, statically-scoped, declarative language that brings the semantics of an interactive spreadsheet application to a compiled language. Extend features immutable values and allows the developer to write code without explicitly specifying the order of computations. It goes beyond typical spreadsheet applications by allowing the developer to encapsulate related sets of computations in reusable functions. In order to offer the best performance, Extend compiles down to LLVM.

Extend's syntax is meant to provide clear punctuation and easily understandable cell range access specifications, while borrowing elements from languages with C-style syntax for ease of development. Despite these syntactic similarities, the semantics of an Extend program have more in common with a spreadsheet such as Microsoft Excel than imperative languages such as C, Java or Python.

## 2. Structure of an Extend Program

An Extend program consists of one or more source files. A source file can contain any number of import directives, function definitions, global variable declarations, and external library declarations, in any order.

### 2.1. Import Statements

Import statements in Extend are written with `import`, followed by the name of a file in double quotes, and terminated with a semicolon. The syntax is as follows:

```
import "string.xtnd";
```

Extend imports act like `#include` in C, except that multiple imports of the same file are ignored. The imports are all aggregated into a single namespace.

## 2.2. Function Definitions

Function definitions comprise the bulk of an Extend program. In short, a function consists of a set of variable declarations, formula assignments, and a return expression. Each variable consists of cells; the values of each cell are, if necessary, calculated according to formulas which each apply to a specified subset of the cells. Each cell value, once calculated, is immutable. A couple examples follow for context; functions are described in detail in section 5.

```
isNumber(x) {
    return typeof(x) == "Number";
}

sum_column([m,1] rng) {
    /* Returns the sum of the values in the column, skipping any values that are non-
       numeric */
    [m,1] running_sum;
    running_sum[0,0] = #rng;
    running_sum[1:,0] = running_sum[-1,] + (isNumber(#rng) ? #rng : 0);
    return running_sum[-1];
}
```

## 2.3. Global Variables

In essence, global variable declarations function as constants in Extend. They are written with the keyword `global`, followed by a variable declaration in the combined variable declaration and assignment format described in section 5.3.1. As with local variables, the cell values of a global variable, once computed, are immutable. A few examples follow:

```
global pi := 3.14159265359;
global num_points := 24;
global [num_points,1]
    circle_x_vals := cos(2 * pi * row() / num_points),
    circle_y_vals := sin(2 * pi * row() / num_points);
```

## 2.4. External Library Declarations

An external library is declared with the `extern` keyword, followed by the name of an object file in double quotes, followed by a semicolon-delimited list of external function declarations enclosed by curly braces. A library declaration informs the compiler of the functions' names and signatures and instructs the compiler to link the object file when producing an executable. An external function declared as `foo` will call an appropriately written C function `extend_foo`. An example follows:

```
extern "mylib.o" {
    foo(arg1, arg2);
    bar();
}
```

This declaration would cause the compiler to link `mylib.o` and would make the C functions `extend_foo` and `extend_bar` available to Extend programs as `foo` and `bar` respectively. The required signature and format of the external functions is specified precisely in section 5.7.

## 2.5. main function

When a compiled Extend program is executed, the `main` function is evaluated. All computations necessary to calculate the return value of the function are performed, after which the program terminates. The `main` function must be a function of a single argument, conventionally denoted `args`, which is guaranteed to be a 1-by-n range containing the command line arguments.

## 2.6. Scoping and Namespace

For functions and for global variables, there is a single namespace that is shared between all files composing an Extend program, and they are visible throughout the entire program. Functions declared in external libraries share this namespace as well. For a local variable, the scope is the entire body of the function in which it is defined. Functions may declare local variables sharing a name with a global variable; inside that function, the name will refer to the local variable.

```
global x := "I'm a global";

foo() {
    y := x; // Scope of x is entire function
    x := "In here I'm a local";
    return y; // Returns "In here I'm a local"
}

bar(x) {
    return x; // Parameters mask globals; returns argument
}

baz() {
    return x; // Returns "I'm a global"
}
```

## 2.7. Identifiers

A function or variable name must begin with a lowercase or uppercase letter and can be followed by any number of letters, digits, or underscores.

### 3. Types and Literals

Extend has three primitive data types, **Number**, **String**, and **Empty**, and one composite type, **Range**.

#### 3.1. Primitive Data Types

A **Number** is an immutable primitive value corresponding to a double-precision 64-bit binary format IEEE 754 value. Numbers can be written in an Extend source file as either integer or floating point constants; both are represented internally as floating-point values. There is no separate type representing an integer.

A **String** is an immutable primitive value that is internally represented a C-style null-terminated byte array corresponding to ASCII values. A String can be written in an Extend source file as a sequence of characters enclosed in double quotes, with the usual escaping conventions. Extend does not allow for slicing of strings to access specific characters; access to the contents of a string will only be available through standard library functions.

The **Empty** type can be written as the keyword `empty`, and serves a similar function to `NULL` in SQL; it represents the absence of a value.

| Primitive Data Type | Examples                         |
|---------------------|----------------------------------|
| Number              | 42 or -5 or 2.71828 or 314159e-5 |
| String              | "Hello, World!\n" or "foo" or "" |
| Empty               | empty                            |

#### 3.2. Ranges

Extend has one composite type, **Range**. A range is a subset of the cells of a variable, as described in section 5.2. Ranges can be nested arbitrarily deeply and can be used to represent (immutable) lists, matrices, or more complicated data structures. For convenience, the range literal syntax can be used to implicitly declare an anonymous variable and assign the range to the entire contents of this variable.

### 3.2.1. Range Literals

A range literal is a semicolon-delimited list of rows, enclosed in curly brackets. Each row is a comma-delimited list of numbers, strings, or range literals. A few examples follow:

```
legal_ranges() {
  r1 := {"Don't"; "Panic"}; // two rows, one column
  r2 := {"Don't", "Think", "Twice"}; // one row, three columns
  r3 := {1,2,3;4,5,6;7,8,9}; // three rows, three columns
  r4 := {"Hello";0,1,2,3,4}; // two rows, five columns
  r5 := {{{{{1}}}}}; // one row, one column
  r7 := {-1.5,-2.5,{-2,"nested"},-3.5}; // one row, four columns
  return
  print_endline(r1) ->print_endline(r2) ->print_endline(r3) ->
  print_endline(r4) -> print_endline(r5) -> print_endline(r7);
}

main(args) {
  return legal_ranges();
}
```

## 4. Expressions

Expressions in Extend allow for arithmetic and boolean operations, function calls, conditional branching, extraction of contents of other variables, string concatenation, and determination of the location of the cell containing the expression. The sections for boolean and conditional operators refer to truthy and falsey values: the Number 0 is the only falsey value; all other values are truthy. As empty represents the absence of a value, it is neither truthy nor falsey.

## 4.1. Arithmetic Operators

The arithmetic operators listed below take one or two expressions and return a number, if both expressions are Numbers, or empty otherwise. Operators grouped within the same inner box have the same level of precedence, and are listed from highest precedence to lowest precedence. All of the binary operators are infix operators, and, with the exception of exponentiation, are left-associative. Exponentiation, bitwise negation, and unary negation are right-associative. All of the unary operators are prefix operators. The bitwise operators round their operands to the nearest signed 32-bit integer (rounding half to even) before performing the operation and evaluate to a Number.

| Operator | Description    | Definition  |
|----------|----------------|---|
| ~        | Bitwise NOT    | Performs a bitwise negation on the binary representation of an expression.                                  |
| -        | Unary negation | A simple negative sign to negate expressions.   |
| **       | Power          | Returns the first expression raised to the power of the second expression                                   |
| *        | Multiplication | Multiplies two expressions  |
| /        | Division       | Divides first expression by second.   |
| %        | Modulo         | Finds the remainder by dividing the expression on the left side of the modulo by the right side expression. |
| «        | Left Shift     | Performs a bitwise left shift on the binary representation of an expression.                                |
| »        | Right Shift    | Performs a sign-propagating bitwise right shift on the binary representation of an expression.              |
| &        | Bitwise AND    | Performs a bitwise AND between two expressions.   |
| +        | Addition       | Adds two expressions together.  |
| -        | Subtraction    | Subtracts second expression from first.   |

|   |             |  |
|---|-------------|--|
|   | Bitwise OR  | Performs a bitwise OR between two expressions.           |
| ^ | Bitwise XOR | Performs a bitwise exclusive OR between two expressions. |

```

easy() {
    return 3 - -3 ** 2 %5; //-1
}
g_easy() {
    return (((1 << 2 | 1) << 2) | 1) << 1; //42
}

```

## 4.2. Boolean Operators

These operators take one or two expressions and evaluate to empty, 0 or 1. Operators grouped within the same inner box have the same level of precedence and are listed from highest precedence to lowest precedence. All of these operators besides logical negation are infix, left-associative operators. The logical AND and OR operators feature short-circuit evaluation. Logical NOT is a prefix, right-associative operator. Besides logical NOT, all boolean operators have lower precedence than all arithmetic operators. For Strings, the boolean operators <, <=, >, and >= implement case-sensitive lexicographic comparison.

| Operator | Description | Definition  |
|----------|-------------|---|
| !        | Logical NOT | Evaluates to 0 or 1 given a truthy or falsey value respectively. !empty evaluates to empty. It has equal precedence with and unary minus. |

|                   |                                  |  |
|-------------------|----------------------------------|--|
| <p>==</p>         | <p>Equals</p>                    | <p>Always evaluates to 0 if the two expressions have different types. If both expressions are primitive values, evaluates to 1 if they have the same type and the same value, or 0 otherwise. If both expressions are ranges, evaluates to 1 if the two ranges have the same dimensions and each cell of the first expression == the corresponding cell of the second expression. <code>empty == empty</code> evaluates to 1. Strings are compared by value.</p> |
| <p>!=</p>         | <p>Not equals</p>                | <p><code>x != y</code> is equivalent to <code>!(x == y)</code>.</p>  |
| <p>&lt;</p>       | <p>Less than</p>                 | <p>If the expressions are both Numbers or both Strings and the first expression is less than the second, evaluates to 1. If the expressions are both Numbers or both Strings and the first expression is greater than or equal to the second, evaluates to 0. Otherwise, evaluates to empty.</p>   |
| <p>&gt;</p>       | <p>Greater than</p>              | <p>Equivalent rules about typing as for &lt;.</p>  |
| <p>&lt;=</p>      | <p>Less than or equal to</p>     | <p>Equivalent rules about typing as for &lt;.</p>  |
| <p>&gt;=</p>      | <p>Greater than or equal to</p>  | <p>Equivalent rules about typing as for &lt;.</p>  |
| <p>&amp;&amp;</p> | <p>Short-circuit Logical AND</p> | <p>If the first expression is falsey or empty, evaluates to 0 or empty respectively. Otherwise, if the second expression is truthy, falsey, or empty, evaluates to 1, 0, or empty respectively.</p>  |

|  |                                 |  |
|--|---------------------------------|--|
|  | Short-circuit Logical<br><br>OR | If the first expression is truthy or empty, evaluates to 1 or empty respectively. Otherwise, if the second expression is truthy, falsey, or empty, evaluates to 1, 0, or empty respectively. |
|--|---------------------------------|--|

```

somethings_false() {
    return !1 != !1 || 4 <= 3;
}
somethings_empty() {
    return empty || empty <= !3 || 5 > 3;
}
somethings_true() {
    return 6 > 2 && !(1 == !1);
}

```

### 4.3. Conditional Expressions

There are two types of conditional expressions: a simple ternary if-then-else expression and a switch expression which can represent more complex logic.

#### 4.3.1. Ternary Expressions

A ternary expression, written either as `cond-expr ? expr-if-true : expr-if-false` or, equivalently, `if(cond-expr, expr-if-true, expr-if-false)` evaluates to `expr-if-true` if `cond-expr` is truthy, or `expr-if-false` if `cond-expr` is falsey. If `cond-expr` is empty, the expression evaluates to empty. Both `expr-if-true` and `expr-if-false` are mandatory. `expr-if-true` is only evaluated if `cond-expr` is truthy, and `expr-if-false` is only evaluated if `cond-expr` is falsey. If `cond-expr` is empty, neither expression is evaluated. The ternary operator `? :` has the lowest precedence level of all operators.

#### 4.3.2. Switch Expressions

A switch expression takes a optional condition, and a list of cases and expressions that the overall expression should evaluate to if the case applies. In the event that multiple cases are true, the expression of the first matching case encountered will be evaluated. An example is provided below:

```

switch_example(foo) {
  return switch (foo) {
    case 2: "foo is 2";
    case 3,4: "foo is 3 or 4";
    default: "none of the above";
  };
}

alternate_format(foo) {
  return switch {
    case foo == 2:
      "foo is 2";
    case foo == 3, foo == 4:
      "foo is 3 or 4";
    default:
      "none of the above";
  };
}

```

The format for a `switch` statement is the keyword `switch`, optionally followed by pair of parentheses containing an expression `switch-expr`, followed by a list of case clauses enclosed in curly braces and delimited by semicolons. A case clause consists of the keyword `case` followed by a comma-separated list of expressions `case-expr1` [, `case-expr2`, [...]], a colon, and an expression `match-expr`, or the keyword `default`, a colon, and an expression `default-expr`. If `switch-expr` is omitted, the `switch` expression evaluates to the `match-expr` for the first case where one of the `case-exprs` is truthy, or `default-expr` if none of the `case-exprs` apply. If `switch-expr` is present, the `switch` expression evaluates to the `match-expr` for the first case where one of the `case-exprs` is equal (with equality defined as for the `==` operator) to `switch-expr`, or `default-expr` if none of the `case-exprs` apply.

The `switch` expression can be used to compactly represent what in most imperative languages would require a long string such as `if (cond1) {...} else if (cond2) {...}`. The `switch` operator is internally converted to an equivalent (possibly nested) ternary expression; as a result, it features short-circuit evaluation throughout.

#### 4.4. Additional Operators

There are four additional operators available to determine the size and type of other expressions. In addition, the infix `+` operator is overloaded to perform string concatenation.

| Operator                  | Description          | Definition  |
|---------------------------|----------------------|---|
| <code>size(expr)</code>   | Dimensions           | Evaluates to a Range consisting of one row and two columns; the first cell contains the number of rows of <code>expr</code> and the second contains the number of columns. If <code>expr</code> is a Number, a String, or Empty, both cells will contain 1. |
| <code>typeof(expr)</code> | Value Type           | Evaluates to "Number", "String", "Range", or "Empty".   |
| <code>row()</code>        | Row Location         | No arguments; returns the row of the cell that is being calculated  |
| <code>column()</code>     | Column Location      | No arguments; returns the column of the cell that is being calculated   |
| <code>+</code>            | String concatenation | "Hello, " + "World!\n" == "Hello, World!\n"   |

Given `[5,5]foo`, then `foo[1,4] = row() * 2 + col()` will evaluate to 6.

## 4.5. Function Calls

A function expression consists of an identifier and an optional list of expressions enclosed in parentheses and separated by commas. The value of the expression is the result of applying the function to the arguments passed in as expressions. Extend is an applicative language: the arguments are evaluated from left to right before the function is called. For more detail, see section 5.

## 4.6. Range Expressions

Range expressions are used to select some or all of the cells of a variable or another range. A range expression consists of a bare identifier, a bare range literal, or an expression and a selector. If a range expression has exactly 1 row and 1 column, the value of the expression is the value of the single cell of the range. If it has more than 1 row or more than 1 column, the value of the expression is the selected range. If the range has zero or fewer rows or zero or fewer columns, the value of the expression is empty. If a range expression with a selector would access a row index or column index greater than the number of rows or columns of the range, or a negative row or column index, the value of the expression is empty.

#### 4.6.1. Slices

A slice consists of an optional integer literal or expression `start`, a colon, and an optional integer literal or expression `end`, or a single integer literal or expression `index`. If `start` is omitted, it defaults to 0. If `end` is omitted, it defaults to the length of the dimension. A single `index` with no colon is equivalent to `index:index+1`. Enclosing `start` or `end` in square brackets is equivalent to the expression `row() + start` or `row() + end`, for a row slice, or `column() + start` or `column() + end` for a column slice. The slice includes `start` and excludes `end`, so the length of a slice is `end - start`. A negative value is interpreted as the length of the dimension minus the value. As mentioned above, the value of a range that is not 1 by 1 is a range, but the value of a 1 by 1 range is essentially dereferenced to the result of the cell formula.

#### 4.6.2. Selections

A selection expression consists of an expression and a pair of slices separated by a comma and enclosed in square brackets, i.e. `[row_slice, column_slice]`. If one of the dimensions of the range has length 1, the comma and the slice for that dimension can be omitted. If the comma is present but a slice is omitted, that slice defaults to `[0]` for a slice corresponding to a dimension of length greater than one, or `0` for a slice corresponding to a dimension of length one.

#### 4.6.3. Corresponding Cell

A very common selection to make is the cell in the "corresponding location" of a different variable. Since this case is so common, `#var` is syntactic sugar for `var[,]`. As a result, if `var` has more than column and more than one row, `#var` is equivalent to `var[row(), column()]`. If `var` has multiple rows and one column, it is equivalent to `var[row(), 0]`. If `var` has one row and multiple columns, it is equivalent to `var[0, column()]`; and if `var` has one row and one column, it is equal to `var[0, 0]`.

#### 4.6.4. Selection Examples

```

selection_examples() {
  foo :=
    {"Alpha", "Bravo", "Charlie", "Delta", "Echo";
     "Foxtrot", "Golf", "Hotel", "India", "Juliett";
     "Kilo", "Lima", "Mike", "November", "Oscar";
     "Papa", "Quebec", "Romeo", "Sierra", "Tango"};

  [3,3] bar;
  bar[0,0] = foo[0,2]; // "Charlie"
  bar[0,1] = foo[0,:]; // {"Alpha", "Bravo", "Charlie", "Delta", "Echo"}
  bar[0,2] = foo[:,2]; // {"Charlie"; "Hotel"; "Mike"; "Romeo"}
  bar[1,1] = foo[[1],[2]]; // "November" - the [1] indicates relative
                          // In this case, works out to foo[2,3]

  bar[1,2] = foo[3,]; // "Romeo" since foo has multiple columns
  bar[2,2] = foo[2:[2],[-1]]; // {"Lima"; "Quebec"}

  /* In this example, each cell of spam would be equal to the cell
   * in ham in the equivalent location plus 1. */
  ham := {2,4,6; 10,11,12; 20,30,40};
  [3,3] spam := #ham + 1; // {3,5,7; 11,12,13; 21,31,41}

  /* In this example, more_cookies would be a 3x4 range where in each row,
   * the value is equal to the value in cookies in the same column.
   * In other words, each row of more_cookies would be a copy of cookies. */
  cookies := {"Chocolate","Oatmeal","Vanilla","Peanut Butter"};
  [3,4] more_cookies := #cookies;

  /* In this example, the values of baz would be
   * 11, 12, 13 in the first row;
   * 21, 22, 23 in the second row;
   * 31, 32, 33 in the third row. */
  ones := {1,2,3}; // 1 row, 3 columns
  tens := {10;20;30}; // 3 rows, 1 column
  [3,3] nums := #ones + #tens; // Equivalent to ones[0,[0]] + tens[[0],0]

  return 0;
}

```

## 4.7. Precedence Expressions

A precedence expression is used to force the evaluation of one expression before another, when that order of operation is required for functions with side-effects. It consists of an expression `prec-expr`, the precedence operator `->`, and an expression `succ-expr`. The value of the expression is `succ-expr`, but the value of `prec-expr` will be calculated first and the result ignored. All functions written purely in Extend are free of side effects. However, some of the external functions provided by the standard library, such as for file I/O and plotting, do have side effects. The precedence operator has the second-lowest grammatical precedence of all operators, higher only than the ternary operator.

## 5. Functions

The bulk of an Extend program consists of functions. Although Extend has some features, such as immutable and lazily evaluated cell values, that are inspired by functional languages, its functions are not *first class objects*. By default, the standard library is automatically compiled and linked with a program, but there are no functions built into the language itself.

### 5.1. Format

As in most programming languages, the header of the function declares the parameters it accepts. The body of the function consists of an optional set of variable declarations and formula assignments, which can occur in any order, and a return statement, which must be the last statement in the function body. All variable declarations and formula assignments, in addition to the return statement, must be terminated by a semicolon. This very simple function returns whatever value is passed into it:

```
foo(arg) {  
    return arg;  
}
```

### 5.2. Variable Declarations

A variable declaration associates an identifier with a set of cells of the specified dimensions, which are listed in square brackets before the identifier. For convenience, if the square brackets and dimensions are omitted, the identifier will be associated with a single cell. In addition, multiple identifiers, separated by commas, can be listed after the dimensions; all of these identifiers will be separate variables, but with equal dimension sizes. The dimensions can be specified as any valid expression that evaluates to a Number, which will be rounded to the nearest signed 32-bit integer. If either dimension is zero or negative, or if the expression does not evaluate to a Number, a runtime error causing the program to halt will occur.

```
[2, 5] foo; // Declares foo as a variable with 2 rows and 5 columns  
[m, n] bar; // Declares bar as a variable with m rows and n columns  
[3, 3] ham, eggs, spam; // Declares ham, eggs and spam as distinct 3x3 variables  
baz; // Declares baz as a variable with single cell
```

### 5.3. Formula Assignment

A formula assignment assigns an expression to a subset of the cells of a variable. Unlike most imperative languages, this expression is not immediately evaluated, but is instead only evaluated if and when it is needed to calculate the return value of the function. A formula assignment consists of an identifier, an optional pair of slices enclosed in square brackets specifying the subset of the cells that the assignment applies to, an =, and an expression, followed by a semicolon. As with the expressions specifying the dimensions of a variable, these slices specifying the cell subset can contain arbitrary expressions, as long as the expression taken as a whole evaluates to a Number, which will be rounded to the nearest signed 32-bit integer. Negative numbers are legal in these slices, and correspond to (dimension length + value).

```
[5, 2] foo, bar, baz; // Declares foo, bar, and baz as distinct 5x2 variables
foo[0,0] = 42; // Assigns the expression 42 to the first cell of the first row of foo
foo[0,1] = foo[0,0] * 2; // Assigns (foo[0,0] * 2) to the 2nd cell of the 1st row of foo
bar = 3.14159; // Assigns pi to every cell of every row of bar
baz[1:-1,0:1] = 2.71828; // Assigns e to cells (1,0) through (3,1), inclusive, of baz

/* The next line assigns foo[[-1],0] + 2 to every cell in
   both columns of foo, besides the first row */
foo[1:,: ] = foo[[-1],0] + 2;
```

The last line of the source snippet above demonstrates the idiomatic Extend way of simulating an imperative language's loop; `foo[4,0]` would evaluate to  $42+2+2+2+2 = 50$  and `foo[4,1]` would evaluate to  $(42*2)+2+2+2+2 = 92$ .

#### 5.3.1. Combined Variable Declaration and Formula Assignment

For convenience, a variable declaration and a formula assignment to all cells of that variable can be combined on a single line by inserting a `:=` and an expression after the identifier. Multiple variables and assignments, separated by commas, can be declared on a single line as well. All global variables must be defined using the combined declaration and formula assignment syntax.

```
/* Creates two 2x2 variables; every cell of foo evaluates to 1 and every cell of
   bar evaluates to 2. */
[2,2] foo := 1, bar := 2;
```

#### 5.3.2. Formula Assignment Errors

If the developer writes code in such a way that more than one formula applies to a cell, a runtime error will occur if the cell's value is required to compute the return expression. If there is

no formula assigned to a cell, the cell will evaluate to empty.

## 5.4. Parameter Declarations

Parameters can be declared with or without dimensions. If dimensions are declared, they can either be specified as integer literals or as identifiers. If a dimension is specified as an integer literal, the program will verify the dimension of the argument before beginning to evaluate the return expression; if it does not match, a runtime error will occur causing the program to halt. If it is specified as an identifier, that variable will contain the dimension size and will be available inside the function body. If the same identifier is repeated in the function declaration, the program will verify that every parameter dimension with that identifier has equal dimension size; if they differ, a runtime error will occur causing the program to halt. A few examples follow:

```
number_of_cells([m,n] arg) {
    return m*n; // m and n are initialized with the dimensions of arg
}

die_unless_primitive([1,1] arg) {
    return 0; // If arg is not a primitive value, a runtime error will occur
}

num_cells_if_column_vector([m,1] arg) {
    // If arg has one column, return number of cells; otherwise runtime error
    return m;
}

die_unless_square([m,m] arg) {
    return 0; // Runtime error if number of rows != number of columns
}

num_cells_if_same_size([m,n] arg1, [m,n] arg2) {
    // If arguments are the same size, return # of cells, otherwise runtime error
    return m*n;
}

main(args) {
    [3,4] foo;
    [3,5] bar;
    return print_endline(num_cells_if_same_size(foo,bar));
}
```

## 5.5. Application on Ranges

Extend gives the developer the power to easily apply operations in a functional style on ranges.

For example, the following function performs cell wise addition:

```
foo([m,n] arg1, [m,n] arg2) {
    [m,n] bar := #arg1 + #arg2;
    return bar;
}
```

This function normalizes a column vector to have unit norm:

```

normalize_column_vector([m,1] arg) {
  [m,1] squared_lengths := #arg * #arg, normalized := #arg / vector_norm;
  vector_norm := sqrt(sum(squared_lengths));
  return normalized;
}

```

## 5.6. Lazy Cell Evaluation and Circular References

All cell values and variable dimensions are evaluated lazily if and when they are needed to calculate the return expression. Using lazy evaluation ensures that the cell values are calculated in a valid topological sort order and allows for detection of circular references; internally this is accomplished by constructing a function for each formula which is called the first time the cell's value is needed, and marking the cell as "in-progress" once it starts being evaluated and as "complete" once the value has been calculated. A cell's value is needed when a range expression consists of that single cell, or when the cell belongs to a range that is assigned as the value for another cell. In other words, an intermediate range expression that consists of multiple cells will not cause the constituent cells to be evaluated; however, a range expression that has one row and one column will cause that one cell's value to be evaluated. In conditional expressions and in short-circuiting operator expressions, only the predicate and the relevant conditional branch will be evaluated. In an expression using the precedence operator, the preceding expression will be evaluated before the succeeding expression. If a program is written in such a way as to cause a circular dependency of one cell on another, and the return expression is dependent on that cell's value, a runtime error will occur. For example, in the following function:

```

maybeCircular(truth_value) {
  x := x;
  return truth_value ? x : 0;
}

main(args) {
  foo :=
  print_endline("To be or not to be?") ->
  print_endline("Enter \"Not to be\" to attempt to evaluate a circular reference.") ->
  readline(STDIN);

  return
  maybeCircular(foo == "Not to be" || foo == "\"Not to be\"") ->
  print_endline("Good thing I didn't look at the value of x.");
}

```

A runtime error will occur if `maybeCircular(1)` is called; but if `maybeCircular(0)` is called, the function will simply return 0.

## 5.7. External Libraries

Using the following library declaration:

```
extern "mylib.o" {
    foo(arg1, arg2);
    bar();
}
```

will make the functions `foo` (taking two arguments) and `bar` (taking zero arguments) available within `Extend`. In LLVM, the compiler will declare external functions `extend_foo` and `extend_bar` as functions of two and zero arguments respectively. All arguments must have the type `value_p`, and the function must have return type `value_p`, declared in the `Extend` standard library header file. In other words, the C file compiled to generate the library must have defined:

```
value_p extend_foo(value_p arg1, value_p arg2) {
    /* function body here; */
}

value_p extend_bar() {
    /* function body here; */
}
```

## 6. Standard Library Reference

### 6.1. File I/O

`open(filename, mode)` - returns a file handle for use with the other file I/O functions  
`close(file_handle)` - close a file handle  
`read(file_handle, num_bytes)` - reads `num_bytes` from a file; 0 reads entire file  
`readline(file_handle)` - read until the first newline  
`write(file_handle, buffer)` - write the contents of `buffer` (a `String`) to the handle  
`STDIN, STDOUT, STDERR` - global variables initialized to the file handles associated with `stdin, stderr, and stdout`  
`print_endline(val)` - convert `val` to a string and write to `STDOUT`

### 6.2. Math Functions - Imported straight from C

`sin(x), cos(x), tan(x), acos(x), asin(x), atan(x), sinh(x), cosh(x), tanh(x),`  
`exp(x), log(x), log10(x), sqrt(x), ceil(x), fabs(x), floor(x), isNaN(x)`  
`random()` - Just for fun - very non-random.

### 6.3. Math Functions - Not imported from C

`isInfinite(x)` - returns -1 for -infinity, 0 for finite, or 1 for +infinity  
`round(val, number_of_digits)`;  
`gcd(m, n)` - returns the GCD of two numbers  
`lcm(m, n)` - returns the LCM of two numbers  
`sign(arg)` - returns -1, 0, or 1  
`sum(rng)` - adds all the numbers in `rng`  
`nmax(n1, n2)` - returns the max of two numbers

`max(rng)` - returns the largest number in a range  
`nmin(n1, n2)` - returns the min of two numbers  
`min(rng)` - returns the smallest number in a range  
`avg([m,n] rng)` - return the average of the numbers in a range  
`stdev([m,n] rng)` - return the standard deviation of the numbers in a range  
`sumsq(rng)` - returns the sum of the squares of the numbers in rng  
`sumproduct([m,n] rng1, [m,n] rng2)` - returns the inner product of rng1 and rng2  
`sumxmy2([m,n] rng1, [m,n] rng2)` - returns the sum of squared differences between the elements of rng1 and rng2  
`mmult([m,n] rng1, [n,p] rng2)` - multiplies two matrices  
`linest([p,q] known_ys, [p,q] known_xs)` - performs a linear regression with known\_ys as the dependent variables and known\_xs as the independent variables  
`normalize([m,n] arg)` - return the unit norm vector in the same direction as arg

## 6.4. String Functions

`len(str)` - returns the length of a String  
`toASCII(val)` - returns a 1 x n range of the ASCII values of a String  
`fromASCII(val)` - converts a 1 x n range of ASCII values into a String  
`parseFloat(str)` - wrapper around C atof()  
`toUpper(text)` - converts a string to uppercase  
`toLower(text)` - converts a string to lowercase  
`left(str, num_chars)` - returns the leftmost num\_chars of str  
`right(str, num_chars)` - returns the rightmost num\_chars of str  
`substring(str, start, length)` - returns a substring of str  
`repeat(str, num)` - repeat a string, num times.  
`toString(arg)` - convert any value into a String representation  
`ltrim(s)` - remove whitespace at the beginning of s  
`rtrim(s)` - remove whitespace at the end of s  
`trim(s)` - remove whitespace on both ends of s  
`reverse(s)` - reverses a string  
`padLeft(str, pad_char, total_length)` - for a string shorter than total\_length, pad on the left with pad\_char  
`charAt(str, i)` - return the ASCII code of the ith character of str  
`parseString(s)` - best efforts to convert a string into the correct value

## 6.5. Plotting

```

bar_chart(file_handle, labels, vals);
line_chart(file_handle, labels, x_vals);

```

## 6.6. Range Functions

`transpose([m,n] rng)` - transpose a matrix; works with any dimensions  
`flatten([m,n] rng)` - turn a rectangular range into a long row vector  
`isNumber(x)` - equal to `typeof(x) == "Number"`  
`isEmpty(x)` - equal to `typeof(x) == "Number"`  
`colRange(start, end)` - return a column vector with the integers from start to (end-1)  
`rowRange(start, end)` - return a row vector with the integers from start to (end-1)  
`match(list, val)` - finds the first occurrence of val in list; list can be either a row or a column vector and does not need to be sorted  
`bsearch(list, val)` - finds the first occurrence of val in list; list must be a sorted column vector  
`join([m,n] cells, joiner)` - concatenate the string representation of either a column or a row vector, using joiner as the delimiter  
`joinRange([m,n] cells, rowJoiner, colJoiner)` - concatenate a range, joining rows with rowJoiner and columns with colJoiner  
`numRows(arg)` - return the number of rows in arg  
`numCols(arg)` - return the number of columns in arg  
`split(string, splitter)` - returns a row vector of strings using splitter (which must be a one-character String) as a delimiter  
`splitToRange(string, row_splitter, col_splitter)` - returns a range of strings using row\_splitter as the row delimiter and col\_splitter as the column delimiter  
`case charAt(trimmed,0) == toASCII("{") && charAt(trimmed,-1) == toASCII("}"):`  
`append([m,n] rg1, [p,q] rg2)` - concatenate two ranges, horizontally  
`stack(rg1, rg2)` - concatenate two ranges, vertically

```
mergesort([m,n] rng, sort_col) - return a sorted copy of rng, using sort_col for  
comparisons
```