

Programming Languages: Application and Interpretation

Version 5.1.1

April 30, 2011

This is the documentation for the software accompanying the textbook **Programming Languages: Application and Interpretation** (PLAI). The full book can be found on the Web at:

<http://www.cs.brown.edu/~sk/Publications/Books/ProgLangs/>

This package contains the following languages:

1 PLAI Scheme

```
#lang plai
```

PLAI Scheme is derived from the `scheme` language. In addition, it includes the `define-type` and `type-case` forms and testing support.

1.1 Defining Types: `define-type`

```
(define-type type-id variant ...)
```

```
variant = (variant-id (field-id contract-expr) ...)
```

Defines the datatype *type-id*. A constructor *variant-id* is defined for each variant. Each constructor takes an argument for each field of its variant.

The value of each field is checked by its associated *contract-expr*. A *contract-expr* may be an arbitrary predicate or a contract.

In addition to the constructors, a `define-type` expression also defines:

- a predicate *type-id?* that returns `true` for instances of the datatype, and `false` for any other value,
- for each variant, a predicate *variant-id?* that returns `true` when applied to a value of the same variant and `false` for any other value,
- for each field of each variant, an accessor *variant-id-field-id* that returns the value of the field, and
- for each field of each variant, a mutator *set-variant-id-field-id!* that set the value of the field.

1.2 Deconstructing Data Structures: `type-case`

```
(type-case datatype-id expr  
  branch ...)
```

```
branch = (variant-id (field-id ...) result-expr ...)  
  | (else result-expr ...)
```

Branches on the datatype instance produced by *expr*, which must be an instance of *datatype-id* (previously defined with *define-type*) Each *branch* extracts the values of the fields, and binds them to *field-id*

If a branch is not specified for each variant, you may use an *else* branch to create a catch-all branch. An *else* branch must be the last branch in the sequence of branches. *type-case* signals a compile-time error if all variants are not covered and the *else* branch is missing. Similarly, *type-case* signals a compile-time error if an *else* branch is unreachable because a branch exists for all variants.

1.3 Testing Infrastructure

PLAI Scheme provides the following syntactic forms for testing.

```
(test result-expr expected-expr)
```

If *result-expr* and *expected-expr* evaluate to the same value, *result-value*, the test prints

```
(good result-expr result-value expected-value location).
```

If they do not evaluate to the same value, the test prints

```
(bad result-expr result-value expected-value location).
```

If evaluating *result-expr* signals an error, the test prints

```
(exception result-expr exception-message <no-expected-value> location)
```

If evaluating *expected-expr* signals an error, the test prints

```
(pred-exception result-expr exception-message <no-expected-value> location)
```

```
(test/pred result-expr pred?)
```

Similar to *test*, but instead of supplying an expected value, the predicate *pred?* is applied to *result-expr*.

If evaluating *pred?* signals an error, the test prints

```
(pred-exception result-expr exception-message <no-expected-value> location)
```

The syntax of *pred?* is considered *expected-value* for the purposes of test reporting.

`error : procedure?`

Like `scheme`'s `scheme:error`, but generates exceptions that are caught by `test/exn`.

`(test/exn result-expr error-message)`

This test succeeds if the expression evaluates to a call to `error`. Moreover, the error message contained in the exception must contain the string `error-message`. Note that `test/exn` only succeeds if the exception was explicitly raised by the user.

For example, the following test succeeds:

```
(test/exn (error "/: division by zero") "by zero")
```

The error message is `"/: division by zero"`, and `"by zero"` is a substring of the error message. However, the following test fails:

```
(test/exn (/ 25 0) "by zero")
```

Although the expression raises an exception and the error string contains `"by zero"`, since the error was not explicitly raised by user-written code, the test fails.

The evaluation of `error-message` is considered *expected-value* for the purposes of test reporting.

`(test/regexp result-expr error-message-regexp)`

This test is similar to `test/exn`, but the error message is matched against a regular expression instead.

The evaluation of `error-message-regexp` is considered *expected-value* for the purposes of test reporting.

1.3.1 Test Flags

`(abridged-test-output [abridge?]) → void?`
`abridge? : boolean? = false`

When this flag is set to `true`, the test forms never prints `result-expr` or `location`.

```
(plai-catch-test-exn [catch?]) → void?  
  catch? : boolean? = true
```

When this flag is set to `true`, exceptions from tests will be caught. By default, exceptions are caught.

```
(halt-on-errors [halt?]) → void?  
  halt? : boolean? = true
```

This flag determines whether the program immediately halts when a test fails. By default, programs do not halt on failures.

```
(print-only-errors [print?]) → void?  
  print? : boolean? = true
```

When this flag is set to `true`, only tests that fail will be printed. By default, the results of all tests are printed.

```
(test-inexact-epsilon epsilon) → void?  
  epsilon : number?
```

When testing inexact values for equality, `test` permits them to differ by `epsilon`. The default value of `epsilon` is `0.01`.

```
(plai-ignore-exn-strings ignore?) → void?  
  ignore? : boolean?
```

If this flag is set to `true`, when testing for exceptions with `test/exn` and `test/regexp`, the message of the exception is ignored. By default, `test/exn` and `test/regexp` only succeed when the message of the exception matches the supplied string or regular expression.

`plai-all-test-results`

This variable is the list of all tests that have been run so far, with the most recent test at the head.

2 GC Collector Scheme

```
#lang plai/collector
```

GC Collector Scheme is based on PLAI Scheme. It provides additional procedures and syntax for writing garbage collectors.

2.1 Garbage Collector Interface

The GC Collector Scheme language provides the following functions that provide access to the heap and root set:

```
(heap-size) → exact-nonnegative-integer?
```

Returns the size of the heap. The size of the heap is specified by the mutator that uses the garbage collector. See `allocator-setup` for more information.

```
(location? v) → boolean?  
v : any/c
```

Determines if `v` is an integer between 0 and `(- (heap-size) 1)` inclusive.

```
(root? v) → boolean?  
v : any/c
```

Determines if `v` is a root.

```
(heap-value? v) → boolean?  
v : any/c
```

A value that may be stored on the heap. Roughly corresponds to the contract `(or/c boolean? number? procedure? symbol? empty?)`.

```
(heap-set! loc val) → void?  
loc : location?  
val : heap-value?
```

Sets the value at `loc` to `val`.

```
(heap-ref loc) → heap-value?
```

`loc : location?`

Returns the value at `loc`.

`(get-root-set id ...)`

Returns the current roots as a list. Local roots are created for the identifiers `id` as well.

`(read-root root) → location?`
`root : root?`

Returns the location of `root`.

`(set-root! root loc) → void?`
`root : root?`
`loc : location?`

Updates the root to reference the given location.

`(procedure-roots proc) → (listof root?)`
`proc : procedure?`

Given a closure stored on the heap, returns a list of the roots reachable from the closure's environment. If `proc` is not reachable, the empty list is returned.

`(with-heap heap expr ...)`

`heap : (vectorof heap-value?)`

Evaluates `(begin expr ...)` in the context of `heap`. Useful in tests:

```
(test (with-heap (make-vector 20)
  (init-allocator)
  (gc:deref (gc:alloc-flat 2)))
  2)
```

2.2 Garbage Collector Exports

A garbage collector must define the following functions:

`(init-allocator) → void?`

`init-allocator` is called before all other procedures by a mutator. Place any requisite initialization code here.

```
(gc:deref loc) → heap-value?  
  loc : location?
```

Given the location of a flat Scheme value, this procedure should return that value. If the location does not hold a flat value, this function should signal an error.

```
(gc:alloc-flat val) → location?  
  val : heap-value?
```

This procedure should allocate a flat Scheme value (number, symbol, boolean, closure or empty list) on the heap, returning its location (a number). The value should occupy a single heap cell, though you may use additional space to store a tag, etc. You are also welcome to pre-allocate common constants (e.g., the empty list). This procedure may need to perform a garbage-collection. If there is still insufficient space, it should signal an error.

Note that closures are flat values. The environment of a closure is internally managed, but contains references to values on the heap. Therefore, during garbage collection, the environment of reachable closures must be updated. The language exposes the environment via the `procedure-roots` function.

```
(gc:cons first rest) → location?  
  first : location?  
  rest : location?
```

Given the location of the `first` and `rest` values, this procedure must allocate a cons cell on the heap. If there is insufficient space to allocate the cons cell, it should signal an error.

```
(gc:first cons-cell) → location?  
  cons-cell : location?
```

If the given location refers to a cons cell, this should return the first field. Otherwise, it should signal an error.

```
(gc:rest cons-cell) → location?  
  cons-cell : location?
```

If the given location refers to a cons cell, this should return the rest field. Otherwise, it should signal an error.

```
(gc:set-first! cons-cell first-value) → void?  
  cons-cell : location?  
  first-value : location?
```

If *cons-cell* refers to a cons cell, set the head of the cons cell to *first-value*. Otherwise, signal an error.

```
(gc:set-rest! cons-cell rest-value) → void?  
  cons-cell : location?  
  rest-value : location?
```

If *cons-cell* refers to a cons cell, set the tail of the cons cell to *rest-value*. Otherwise, signal an error.

```
(gc:cons? loc) → boolean?  
  loc : location?
```

Returns **true** if *loc* refers to a cons cell. This function should never signal an error.

```
(gc:flat? loc) → boolean?  
  loc : location?
```

Returns **true** if *loc* refers to a flat value. This function should never signal an error.

3 GC Mutator Scheme

```
#lang plai/mutator
```

The GC Mutator Scheme language is used to test garbage collectors written with the §2 “GC Collector Scheme” language. Since collectors support a subset of Scheme’s values, the GC Mutator Scheme language supports a subset of procedures and syntax. In addition, many procedures that can be written in the mutator are omitted as they make good test cases. Therefore, the mutator language provides only primitive procedures, such as `+`, `cons`, etc.

3.1 Building Mutators

The first expression of a mutator must be:

```
(allocator-setup collector-module
                 heap-size)
```

heap-size = *exact-nonnegative-integer?*

collector-module specifies the path to the garbage collector that the mutator should use. The collector must be written in the GC Collector Scheme language.

The rest of a mutator module is a sequence of definitions, expressions and test cases. The GC Mutator Scheme language transforms these definitions and statements to use the collector specified in `allocator-setup`. In particular, many of the primitive forms, such as `cons` map directly to procedures such as `gc:cons`, written in the collector.

3.2 Mutator API

The GC Mutator Scheme language supports the following syntactic forms:

```
if and or cond case define define-values let let-
values let* set! lambda λ quote error begin
```

The syntax of `define` is restricted to: `(define (id arg-id ...) body ...+)`. The syntax of `lambda` and `λ` are restricted to: `(lambda (arg-id ...) body ...+)`.

The language also defines the following procedures:

```
add1 sub1 zero? + - * / even? odd? = < > <= >= cons first rest
set-first! set-rest! cons? symbol? symbol=? number? boolean? empty? eq?
```

```
(set-first! c v) → void
  c : cons?
  v : any/c
```

Sets the `first` of the cons cell `c`.

```
(set-rest! c v) → void
  c : cons?
  v : any/c
```

Sets the `rest` of the cons cell `c`.

The identifier `empty` is defined to invoke `(gc:alloc-flat empty)` wherever it is used.

Other common procedures are left undefined as they can be defined in terms of the primitives and may be used to test collectors.

Additional procedures from `scheme` may be imported with:

```
(import-primitives id ...)
```

Imports the procedures `id ...` from `scheme`. Each procedure is transformed to correctly interface with the mutator. That is, its arguments are dereferenced from the mutator's heap and the result is allocated on the mutator's heap. The arguments and result must be `heap-value?`s, even if the imported procedure accepts or produces structured data.

For example, the GC Mutator Scheme language does not define `modulo`:

```
(import-primitives modulo)

(test/value=? (modulo 5 3) 2)
```

3.3 Testing Mutators

GC Mutator Scheme provides two forms for testing mutators:

```
(test/location=? mutator-expr1 mutator-expr2)
```

`test/location=?` succeeds if `mutator-expr1` and `mutator-expr2` reference the same location on the heap.

```
(test/value=? mutator-expr scheme-datum/quoted)
```

`test/value=?` succeeds if `mutator-expr` and `scheme-datum/expr` are structurally equal. `scheme-datum/quoted` is not allocated on the mutator's heap. Furthermore, it must either be a quoted value or a literal value.

```
(printf format mutator-expr ...)
```

```
format = literal-string
```

In GC Mutator Scheme, `printf` is a syntactic form and not a procedure. The format string, `format` is not allocated on the mutator's heap.

3.4 Generating Random Mutators

```
(require plai/random-mutator)
```

This PLAI library provides a facility for generating random mutators, in order to test your garbage collection implementation.

```
(save-random-mutator file
                    collector-name
                    [#:heap-values heap-values
                   #:iterations iterations
                   #:program-size program-size
                   #:heap-size heap-size]) → void?
file : path-string?
collector-name : string?
heap-values : (cons heap-value? (listof heap-value?))
              = (list 0 1 -1 'x 'y #f #t '())
iterations : exact-positive-integer? = 200
program-size : exact-positive-integer? = 10
heap-size : exact-positive-integer? = 100
```

Creates a random mutator that uses the collector `collector-name` and saves it in `file`.

The mutator is created by first making a random graph whose nodes either have no outgoing edges, two outgoing edges, or some random number of outgoing edges and then picking a random path in the graph that ends at one of the nodes with no edges.

This graph and path are then turned into a PLAI program by creating a `let` expression that binds one variable per node in the graph. If the node has no outgoing edges, it is bound to a `heap-value?`. If the node has two outgoing edges, it is bound to a pair and the two edges

are put into the first and rest fields. Otherwise, the node is represented as a procedure that accepts an integer index and returns the destination node of the corresponding edge.

Once the `let` expression has been created, the program creates a bunch of garbage and then traverses the graph, according to the randomly created path. If the result of the path is the expected heap value, the program does this again, up to *iterations* times. If the result of the path is not the expected heap value, the program terminates with an error.

The keyword arguments control some aspects of the generation of random mutators:

- Elements from the *heap-values* argument are used as the base values when creating nodes with no outgoing edges. See also *find-heap-values*.
- The *iterations* argument controls how many times the graph is created (and traversed).
- The *program-size* argument is a bound on how big the program it is; it limits the number of nodes, the maximum number of edges, and the length of the path in the graph.
- The *heap-size* argument controls the size of the heap in the generated mutator.

```
(find-heap-values input) → (listof heap-value?)  
input : (or/c path-string? input-port?)
```

Processes *input* looking for occurrences of *heap-value?*s in the source of the program and returns them. This makes a good start for the *heap-values* argument to *save-random-mutator*.

If *input* is a port, its contents are assumed to be a well-formed PLAI program. If *input* is a file, the contents of the file are used.

4 Web Application Scheme

```
#lang plai/web
```

The Web Application Scheme language allows you to write server-side Web applications for the PLT Web Server.

For more information about writing Web applications, see: *Web: Racket Web Applications*.

When you click on the Run button in DrRacket, your Web application is launched in the Web server.

The application is available at `http://localhost:8000/servlets/standalone.ss`.

The Web Application Scheme language will automatically load this URL in your Web browser.

You may use `no-web-browser` to prevent the browser from being launched and `static-files-path` to serve additional static files.

4.1 Web Application Exports

A Web application must define a procedure `start`:

```
(start initial-request) → response?  
  initial-request : request?
```

The initial request to a Web application is serviced by this procedure.