

PGo Manual

For commit d0241d3

August 25, 2017

1 Introduction

TLA+ is a formal specification language, built on the mathematical concepts of set theory and the temporal logic of actions. Using the TLC model checker, TLA+ specifications can be checked exhaustively for specific properties, and the TLA proof system (TLAPS) allows for machine-checked proofs. PlusCal is an algorithm language which can be translated to TLA+ and uses TLA+ as its expression language. It is easy to specify and verify distributed algorithms in PlusCal, thanks to its simple constructs for nondeterminism, concurrency primitives, and rich mathematical constructs. However, PlusCal is not a programming language so it cannot be run, which limits its utility.

PGo aims to correspond a verified PlusCal specification with an executable Go implementation. PGo compiles an annotated PlusCal specification into a Go program which preserves the semantics of the specification. The main goal of PGo is to create a compiled Go program which is easy to read and edit. The intended use case for PGo is for the developer to create and verify a PlusCal spec of the system, compile it with PGo, then edit the compiled program to fully implement details not included in the specification. Programs compiled by PGo are type safe (they do not contain type assertions which panic).

2 Using PGo

2.1 Installation

Requirements: Eclipse or Ant 1.9

- Git clone the source at <https://bitbucket.org/whiteoutblackout/pgo/>
 - Option 1: Import as an Eclipse project
 - Option 2: Execute `ant pgo` assuming the project is in the `pgo/` directory

Dependencies:

- The Plume options library.
- The TLA+ tools.
- The Go Data Structures library.

PGo was tested on JRE8 and Go 1.8.3.

2.2 Execution

To run PGo, run the Eclipse project or run `pgo.sh`. The command-line usage is `pgo [options] pcalfile gofolder gofile`.

Optional command line arguments:

```
-h --help=<boolean> - Print usage information [default false]
-q --logLvlQuiet=<boolean> - Reduce printing during execution [default false]
-v --logLvlVerbose=<boolean> - Print detailed information during execution [default false]
-i --infile=<string> - the input pluscal file to transpile [default ]
-o --outfile=<string> - the output file to generate [default ]
-d --outfolder=<string> - the output folder to put additional packages [default ]
--writeAST=<boolean> - write the AST generated and skip the rest [default false]
```

2.3 Annotations

PGo requires annotation of the PlusCal source file to properly compile the PlusCal source. Annotations must be placed in comments within the body of the PlusCal algorithm, but there can be multiple annotations per comment. Annotations may appear in any part of the comment block containing the PlusCal algorithm. Annotations are of the form `@PGo{ <annotation> }@PGo`.

2.3.1 Annotating types

Many annotations require specifying a type name. PGo uses the following format for type names:

Primitive types	
Go type	Annotation keyword
int	int
bool	bool
uint64	uint64
float64	float64
string	string
(void type)	void
Collection types	
[]<type>	[]<type>
set with <type> elements	set [<type>]
map with <K> keys and <V> values	map [<K>] <V>
tuple with all elements of type <E>	chan [<E>]
tuple with <i>i</i> th element of type <E _{<i>i</i> <td>tuple [<E₁>, <E₂>, ...]</td>}	tuple [<E ₁ >, <E ₂ >, ...]

Type names cannot contain spaces.

2.3.2 TLA+ definitions and constants declared outside of the algorithm block (Required)

The PlusCal algorithm can make use of TLA+ definitions that are found outside the algorithm block. These are not parsed by PGo and need to be in an annotation for PGo to parse them. Any TLA+ definitions found in `define` blocks also need to be annotated. The annotation for a TLA+ definition is of the form `def <name>(<params>)? <type>? == <TLA expression>`. The definition can be copied almost verbatim from the TLA+, but a parameter is specified by `<name> <type>` so typing information needs to be added to the parameters. The type that the expression should evaluate to may also be specified, if desired. A TLA+ definition without parameters is compiled into a variable, and a definition with parameters is compiled into a function.

```

1  \* outside the algorithm body
2  foo(i, j) == i + j
3  bar == {1, 3, 5}
4  \* ...
5  \* inside the algorithm body
6  (*
7     @PGo{ def foo(i int, j int) int == i + j}@PGo
8     @PGo{ def bar = {1, 3, 5} }@PGo
9  *)

```

PlusCal

```

1  func foo(i int, j int) int {
2      return i + j
3  }
4
5  var bar pgoutil.Set = pgoutil.NewSet(1, 3, 5)

```

Compiled Go

There will be constants in PlusCal that are not declared in the PlusCal algorithm (e.g., constant N for model checking). These variables will need to be specified using PGo annotations either as constants in the compiled Go program, or command line arguments to the Go program. Constants are specified as `const <varname> <type> <val>` indicating that `varname` is a Go constant of type `<type>` with initial Go value `<val>`. Command line argument variables of Go are specified as `arg <varname> <type> (<flagname>)?` indicating that variable `<varname>` is of type `<type>` and is going to be specified through a command line argument to the Go program. If no `<flagname>` is specified, then the variable will be a positional argument in the order that it appeared in the annotation. If `<flagname>` is specified, then the variable will be set on the command line using `-<flagname>=<value>`.

If a constant is not a primitive type, it cannot be declared as constant or as a command line argument in Go. The constant can instead be annotated as a TLA+ definition, where the expression is the desired constant value. This will be compiled to a global variable that is initialized with the given value. PGo provides a compile-time guarantee that the constant indeed remains constant (it is not assigned to or mutated).

```

1  \* outside the algorithm body
2  CONSTANT N, M, ProcSet
3  \* ...
4  \* inside the algorithm body
5  (*
6     @PGo{ arg N int }@PGo
7     @PGo{ arg M int PGoFlag }@PGo
8     @PGo{ def ProcSet set[string] == {"a", "b", "c"} }@PGo
9  *)
10 print << N, M >>;

```

PlusCal

```

1 var ProcSet pgoutil.Set = pgoutil.NewSet("a", "b", "c")
2 var N int
3 var M int
4
5 func main() {
6     flag.IntVar(&M, "PGoFlag", 0, "")
7     flag.Parse()
8     N, _ = strconv.Atoi(flag.Args()[0])
9
10    fmt.Printf("%v %v\n", N, M)
11 }

```

Compiled Go

```

$ go run CompiledGo.go -PGoFlag=1 2
2 1

```

Command-line usage

2.3.3 PlusCal procedure return values (Optional)

PlusCal has no return values, so procedures can only return values by writing to a global variable. PGo requires the developer to indicate which variable serves this purpose. In PGo, the annotation is of the form `ret <varname>`. PGo automatically scans all function definitions for the one where the variable is used. Note that using this feature will remove the specified variable from the global variables. If you rely on global state of the variable for more than just the function return value, do not specify it as a return value and use the global variable instead.

2.3.4 Variable types (Optional)

PGo will automatically infer types for variables declared in PlusCal. However, you may want to specify the types rather than using the inferred types (e.g., you want a variable to be a `uint64` rather than an `int`). This is possible by specifying `var <varname> <type>`. Annotations are required for variables that involve PlusCal tuples, since these may compile to slices or tuples depending on context. If no type annotation is provided for a variable, PGo will indicate the type it inferred in the output Go code.

2.3.5 Function signatures

Similar to specifying variable types. The annotation `func <rtype>? <funcname>() <p1type>?+ <p2type>?+ <p3type>?+ <funcname>()` having a return type of `<rtype>` if specified, and parameters of type `<p1type>`, `<p2type>`,... If any types are specified, all return types or parameters must be specified.

```

1 /* @PGo{ func foo() int string }@PGo
2 procedure foo(a, b) {
3     print << a + 1, b >>;
4 }

```

PlusCal

```

1 func foo(a int, b string) {
2     fmt.Println("%v %v\n", a + 1, b)
3 }

```

Compiled Go

2.3.6 Locking (Optional)

By default, PGo adds locks when compiling a multiprocess PlusCal algorithm. The locking behaviour is described in more detail in 3.8. The locking annotation is of the form `lock [true|false]`. PGo adds or omits locks based on the annotation, overriding the default behaviour.

3 Translation

PlusCal has many language constructs that are translated to Go in a non-trivial way. These constructs and their translations are described in this section. The `pgoutil` Go helper library, which contains some useful data structures and helper methods, is also described.

3.1 Variable declarations

In addition to the simple variable declaration `var = <val>`, PlusCal supports the declaration `var \in <set>`. This asserts that the initial value of `var` is an element of `<set>`. This is translated into a loop in which each set element is assigned to `var`. This is the same as the behaviour of TLC in model-checking mode, and is the most likely use case.

```
1 variables
2   S = {1, 3};
3   v \in S;
4 {
5   \* algorithm body...
6 }
```

PlusCal

```
1 import "pgoutil"
2
3 var S pgoutil.Set = pgoutil.NewSet(1, 3) // PGo inferred type set[int]
4 var v int // PGo inferred type int
5
6 func main() {
7   for v_interface := range S.Iter() {
8     v = v_interface.(int)
9     // algorithm body...
10  }
11 }
```

Compiled Go

(The set methods of the `pgoutil` package are described in 3.4.1.)

If no type annotation was provided for a variable, PGo prints a comment which contains the type that it inferred. If the variable needs thread-safe access, then PGo will also print the lock group that the variable belongs to (a description of locking is in 3.8).

3.2 Variable assignment

PlusCal supports multiple variable assignment statements: the statement `x := a || y := b` evaluates the right-hand sides first, then assigns the values to the left-hand sides. A common use is swapping the variables `x` and `y` with the statement `x := y || y := x`.

Go has a multiple assignment construct, but a PlusCal assignment might not translate to a Go assignment (for example, assigning a value to a particular map key translates to the `Put` method). The right-hand sides are first evaluated and assigned to temporary variables, then assigned to their corresponding variables.

```
1 x := y || y := x || mapping[i] := a
```

PlusCal

```
1 x_new := y
2 y_new := x
3 mapping_new := a
4 x = x_new
5 y = y_new
6 mapping.Put(i, a)
```

Compiled Go

3.3 Macros

PlusCal macros have the same semantics as C/C++ `#define` directives. PGo expands the PlusCal macro wherever it occurs.

```
1 variables p = 1, q = 2;
2 macro add(a, b) {
3     a := a + b;
4 }
5 {
6     add(p, q);
7     print p;
8 }
```

PlusCal

```
1 import "fmt"
2
3 var p int = 1
4 var q int = 2
5
6 func main() {
7     p = p + q
8     fmt.Println("%v", p)
9 }
```

Compiled Go

3.4 Data types

PlusCal supports maps, sets, and tuples, which are implemented in the `pgoutil` library.

3.4.1 Sets

PlusCal sets are implemented in Go by the `pgoutil.Set` type. This container is an ordered, thread-safe set which has methods equivalent to PlusCal set operations.

```
1 variables
2     A = {1, 2, 3};
3     B = 3 .. 6; /* the set {3, 4, 5, 6}
4     C = A \union B;
5     D = SUBSET C; /* powerset of C
6 {
7     print A = C;
8 }
```

```
1 import (
2     "fmt"
3     "pgoutil"
4 )
5
6 var A pgoutil.Set = pgoutil.NewSet(1, 2, 3)
7 var B pgoutil.Set = pgoutil.Sequence(3, 6)
8 var C pgoutil.Set = A.Union(B)
9 var D pgoutil.Set = C.PowerSet()
10
11 func main() {
12     fmt.Printf("%v\n", A.Equal(C))
13 }
```

PlusCal also supports the typical mathematical set constructor notations, which are offloaded to library methods:

```
1 variables
2     S = {1, 5, 6};
3     T = {2, 3};
4     U = {x \in S : x > 3}; /* equivalent to {5, 6}
5     V = {x + y : x \in S, y \in T}; /* equivalent to {3, 4, 7, 8, 9}
6     /* ...
```

PlusCal

```

1 import "pgoutil"
2
3 var S pgoutil.Set = pgoutil.NewSet(1, 5, 6)
4 var T pgoutil.Set = pgoutil.NewSet(2, 3)
5 var U pgoutil.Set = pgoutil.SetConstructor(func(x int) bool {
6     return x > 3
7 }, S)
8 var V pgoutil.Set = pgoutil.SetImage(func(x int, y int) int {
9     return x + y
10 }, S, T)
11 // ...

```

Compiled Go

While not as concise as the PlusCal, the output Go code is still readable.

The Cartesian product operator $\backslash X$ in PlusCal is not associative. For example,

- $\langle\langle 1, 2, 3 \rangle\rangle$ is in $\text{Nat} \backslash X \text{Nat} \backslash X \text{Nat}$,
- $\langle\langle \langle 1, 2 \rangle, 3 \rangle\rangle$ is in $(\text{Nat} \backslash X \text{Nat}) \backslash X \text{Nat}$, and
- $\langle\langle 1, \langle\langle 2, 3 \rangle\rangle \rangle\rangle$ is in $\text{Nat} \backslash X (\text{Nat} \backslash X \text{Nat})$,

and none of these are equivalent. In Go, `S.CartesianProduct(T, U, V, ...)` is equivalent to $S \backslash X T \backslash X U \backslash X \dots$. PGo composes these function calls to create the parenthesized expressions above.

```

1 variables
2   S = {1, 2, 3};
3   T = S \X S \X S; \* << 1, 3, 2 >> \in T
4   U = S \X (S \X S); \* << 1, << 3, 2 >> >> \in U

```

PlusCal

```

1 import "pgoutil"
2
3 var S pgoutil.Set = pgoutil.NewSet(1, 2, 3)
4 var T pgoutil.Set = S.CartesianProduct(S, S)
5 var U pgoutil.Set = S.CartesianProduct(S.CartesianProduct(S))

```

Compiled Go

3.4.2 Maps

PlusCal “functions” are just maps in Go. `pgoutil` implements an ordered, thread-safe map in Go, which supports using slices and tuples as keys (unlike the builtin Go map). There are several syntaxes to declare maps, which are handled by library methods similar to the set methods.

A PlusCal function can be indexed by multiple indices. This is syntactic sugar for a map indexed by a tuple whose components are the indices.

```

1 variables
2   S = {1, 2};
3   f = [x \in S, y \in S |-> x + y];
4   a = f[2, 2]; \* a = 4

```

PlusCal

```

1 import "pgoutil"
2
3 var S pgoutil.Set = pgoutil.NewSet(1, 2)
4 var f pgoutil.Map = pgoutil.MapsTo(func(x int, y int) int {
5     return x + y
6 }, S, S)
7 var a int = f.Get(pgoutil.NewTuple(2, 2))

```

Compiled Go

(Behind the scenes, the `pgoutil.Map` uses tuples as keys.)

3.4.3 Tuples

PlusCal tuples are used in several different contexts, so variables involving tuples must be annotated with their types. Tuples can store homogeneous data, in which case they correspond to Go slices. Tuple components may be of different types, which correspond to `pgoutil.Tuples`. Most often in multiprocessing algorithms, tuples are used for communication, in which case `pgoutil.Tuples` are used. PlusCal tuples are 1-indexed, but Go tuples and slices are 0-indexed, so 1 is subtracted from all indices in Go.

```

1 (* @PGo{ var slice []string }@PGo
2   @PGo{ var tup tuple[int,string] }@PGo *)
3 variables
4   slice = << "a", "b", "c" >>;
5   tup = << 1, "a" >>;
6 {
7   print slice[2]; \* "b"
8 }

```

PlusCal

```

1 import (
2   "fmt"
3   "pgoutil"
4 )
5
6 var slice []string = []string{"a", "b", "c"}
7 var tup pgoutil.Tuple = pgoutil.NewTuple(1, "a")
8
9 func main() {
10   fmt.Printf("%v", slice[2 - 1])
11 }

```

Compiled Go

3.5 Predicate operations

PlusCal supports the mathematical quantifiers \forall and \exists . PGo compiles these to the library methods `ForAll` and `Exists`.

```

1 variables
2   S = {1, 2, 3};
3   T = {4, 5, 6};
4   b1 = \E x \in S : x > 2 \* TRUE
5   b2 = \A x \in S, y \in T : x + y > 6 \* FALSE

```

PlusCal


```

1 import "pgoutil"
2
3 var S pgoutil.Set = pgoutil.NewSet(1, 2, 3)
4 var T pgoutil.Set = pgoutil.NewSet(4, 5, 6)
5 var b1 bool = pgoutil.Exists(func(x int) bool {
6     return x > 2
7 }, S)
8 var b2 bool = pgoutil.ForAll(func(x int, y int) bool {
9     return x + y > 6
10 }, S, T)

```

Compiled Go

Related is Hilbert's ε operator (called **CHOOSE** in PlusCal). In PlusCal, **CHOOSE** $x \setminus \text{in } S : P(x)$ is defined to be some arbitrary value in S such that $P(x)$ is true. If there is no such x , the expression is unspecified. Furthermore, **CHOOSE** is deterministic and always selects the same x given equal sets and equivalent predicates. The Go equivalent function `pgoutil.Choose` selects the least element that satisfies the predicate. This guarantees the determinism that the operator provides, and is the same as TLC's behaviour. If no element satisfies the predicate P , `pgoutil.Choose` panics, similar to TLC returning an error when it encounters unspecified behaviour.

```

1 variables
2     S = {1, 2, 3};
3     a = CHOOSE x \in S : x >= 1
4     b = CHOOSE y \in S : y < 0 \* TLC error

```

PlusCal

```

1 import "pgoutil"
2
3 var S pgoutil.Set = pgoutil.NewSet(1, 2, 3)
4 var a int = pgoutil.Choose(func(x int) bool {
5     return x >= 1
6 }, S).(int) // a == 1
7 var b int = pgoutil.Choose(func(y int) bool {
8     return y < 0
9 }, S).(int) // panic (no element satisfies predicate)

```

Compiled Go

3.6 With

The PlusCal `with` statement has the syntax

```

1 variables S_1 = {1, 2, 3}, a = "foo";
2 \* ...
3 with (x_1 \in S_1, x_2 = a, ...)
4 {
5     \* do stuff with the x_i ...
6 }

```

PlusCal

This construct selects a random element from each S_i and assigns them to the local variables x_i . If the syntax `x_i = a` is used, this simply assigns a to x_i . In Go, this translates to

```

1 import (
2     "math/rand"
3     "pgoutil"
4     "time"
5 )
6
7 var S_1 pgoutil.Set = pgoutil.NewSet(1, 2, 3)
8 var a string = "foo"
9 // ...
10 func main() {
11     rand.Seed(time.Now().UnixNano())
12     // ...
13     {
14         var x_1 int = S_1.ToSlice()[rand.Intn(S_1.Size())].(int)
15         var x_2 string = a
16         // ...
17         // do stuff with the x_i ...
18     }
19 }

```

Compiled Go

In Go the random number generator is seeded with the current Unix time at the beginning of the `main` function. The local variables declared by the `with` and its body are contained in a separate code block, to preserve the variables' local scope.

3.7 Processes

The PlusCal algorithm body can either contain statements in a uniprocess algorithm, or process declarations in a multiprocess algorithm. Processes can be declared with the syntax `process (Name \in S)` or `process (Name = Id)`. The first construct spawns a set of processes, one for each ID in the set `S`, and the second spawns a single process with ID `Id`. A process can refer to its identifier with the keyword `self`. In TLC, multiprocess algorithms are simulated by repeatedly selecting a random process then advancing it one atomic step (if it does not block).

PGo converts each process body to a function and spawns a goroutine per process. The function takes a single parameter `self`, the process ID. There are two semantic considerations: the main goroutine should not exit before all goroutines finish executing, and the time at which all child goroutines begin executing should be synchronized. To preserve these semantics, PGo uses a global waitgroup which waits on all goroutines, and each process body pulls from a dummy channel before beginning execution. When all goroutines have been initialized, the dummy channel is closed so that the channel pull no longer blocks, allowing for a synchronized start.

The following is a simple example:

```

1 variables
2     idSet = {1, 2, 3};
3     id = "id";
4
5 process (PName \in idSet)
6     variable local;
7     {
8         local := self;
9     }
10
11 process (Other = id) {
12     print self;
13 }

```

```

1 import (
2     "fmt"
3     "pgoutil"
4     "sync"
5 )
6
7 var idSet pgoutil.Set = pgoutil.NewSet(1, 2, 3) // PGo inferred type set[int]
8 var id string = "id" // PGo inferred type string
9 var PGoWait sync.WaitGroup
10 var PGoStart chan bool = make(chan bool)
11
12 func PName(self int) {
13     var local int // PGo inferred type int
14     defer PGoWait.Done()
15     <-PGoStart
16     local = self
17 }
18 func Other(self string) {
19     defer PGoWait.Done()
20     <-PGoStart
21     fmt.Printf("%v\n", self)
22 }
23
24 func main() {
25     // spawn goroutines for PName processes
26     for procId := range idSet.Iter() {
27         PGoWait.Add(1)
28         go PName(procId.(int))
29     }
30     PGoWait.Add(1)
31     close(PGoStart)
32     PGoWait.Wait()
33 }

```

Compiled Go

Note that all processes use the same waitgroup (PGoWait) and dummy channel (PGoStart).

3.8 Labels

In PlusCal, labels are used as targets for goto statements and also to specify atomic operations. If a statement is labeled, all statements until the next label are considered to be a single atomic operation. In Go, unused labels cause compile errors so PGo only inserts labels when they are used in goto statements.

To deal with atomicity, PGo divides the global variables into groups and guards each group with a sync.RWMutex. PGo groups variables by performing a disjoint-set union, merging two variable sets when two variables in them can be accessed in the same label. The following is a simple example:

```

1 variables a = 0, b = 1, c = 2, d = 3;
2 process (P \in {1, 2, 3}) {
3     lab1: a := 1;
4         b := 2;
5     lab2: b := 3;
6         c := 4;
7     lab3: d := 5;
8 }

```

PlusCal

```

1 var PGoLock []sync.RWMutex = make([]sync.RWMutex, 2)
2 var a int = 0 // PGo inferred type int; Lock group 0
3 var b int = 1 // PGo inferred type int; Lock group 0
4 var c int = 2 // PGo inferred type int; Lock group 0
5 var d int = 3 // PGo inferred type int; Lock group 1
6 var PGoWait sync.WaitGroup
7 var PGoStart chan bool = make(chan bool)
8
9 func P(self int) {
10     defer PGoWait.Done()
11     <-PGoStart
12     PGoLock[0].Lock()
13     a = 1
14     b = 2
15     PGoLock[0].Unlock()
16     PGoLock[0].Lock()
17     b = 3
18     c = 4
19     PGoLock[0].Unlock()
20     PGoLock[1].Lock()
21     d = 5
22     PGoLock[1].Unlock()
23 }
24 // ...

```

Compiled Go

The variable `b` may be accessed atomically with `a` (in the label `lab1`) and also with `c` (in the label `lab2`) so all three of `a`, `b`, and `c` must be grouped together to prevent data races. PGo locks the correct group before each atomic operation and unlocks it afterwards. Even single operations must use the lock, since there are no atomicity guarantees for most Go statements. If the atomic operations are specified to be smaller in PlusCal by adding more labels, PGo will compile smaller variable groups, allowing for more parallelism.

Locking may be turned off with the annotation `lock false`, and it can similarly be enforced using the annotation `lock true`.

3.9 Limitations

Not all PlusCal specifications can be compiled by PGo. This is an overview of some important PlusCal features that are currently unsupported.

- The `await` keyword
- The `either` construct
- Referencing `self` in a procedure call
- TLA+ features:
 - Alignment of boolean operators in bulleted lists determining precedence
 - Records (structs)
 - Bags (multisets)
 - `IF ... THEN ... ELSE` and `CASE` expressions
 - The `LET .. IN` construct
 - Temporal logic operators
 - Recursive definitions
 - Operator definition (use definitions with parameters instead)

The following are features that may be desirable for a developer working with the compiled program but are not yet considered by PGo:

- Support for networking
- Interfacing with other programs
- Reading input from sources other than the command line

4 Notes on the Go library

The `pgoutil` Go package provides some useful data structures and methods that correspond to PlusCal language features. The data structures are generic and provide no runtime type safety. PGo guarantees at compile-time that the compiled code does not contain any type assertions that will panic. This section contains some useful information about these containers for developers wishing to modify and optimize the compiled Go code.

4.1 Maps

The Go library map is mutable, but the PlusCal map is not. An assignment to a component of a PlusCal map does not change the value of other references to that map:

```
1 variables
2   S = {1, 2, 3};
3   M = [x \in S |-> 2 * x];
4   T = {M};
5 {
6   print M \in T; \* TRUE
7   M[2] := 3;
8   print M \in T; \* FALSE
9 }
```

The `pgoutil.Map` clones any map that is to be inserted into a map or a set. This preserves semantics but may unnecessarily slow the program, if the map is never mutated. Similarly, since slices are mutable but PlusCal tuples are immutable, slices are cloned before being used as map keys or values.

The Go map provides three iterators: `Keys` and `Values` range over the keys and values stored in the map ordered by keys, and `Iter` ranges over `KVPairs`, which are `structs` with two fields: `Key` and `Val`.

If a PlusCal map is indexed by multiple indices, the Go map internally uses a `Tuple` of the indices as its key. The value `M[a, b]` in PlusCal can be recovered in Go by calling `M.Get(pgoutil.NewTuple(a, b))`.

The `pgoutil.MapsTo` method makes extensive use of reflection. Performance-critical sections of the application should be rewritten to avoid the use of this method. The `pgoutil.Map`'s comparator function also uses reflection, which may slow the program especially when using slices or structs as map keys.

The `pgoutil` package exposes a `Map` interface. A custom data type which implements `pgoutil.Map` is interoperable with the builtin type.

4.2 Sets and predicates

The `pgoutil.Set` is implemented using the `pgoutil.Map` as a base, so much of the same semantics apply.

The sets in the Go set library are mutable, whereas TLA+ sets are immutable. However, all operations equivalent to TLA+ set operations are implemented as methods which do not mutate the Go set. For performance purposes, sets are not cloned when inserted into a map or a set. The developer modifying the compiled Go program should exercise caution when making sets of sets or maps with sets if mutating the sets that are inserted. The sets should be copied with the `Clone` method before insertion.

Simple optimizations that can be made by the developer when dealing with sets include replacing the PlusCal `S := S \union {elt}` with simply adding `elt` to `S` and similarly replacing `S := S \ {elt}` with removing `elt` from `S`. Since this mutates the set, the set must be cloned if it is inserted into a set or used with a map.

Elements inserted into a set must be of the same type; otherwise, the program will panic at runtime. This is the same behaviour as TLC, which throws an error when one set contains elements of different types. At compile-time, PGo will throw an error if elements in a set are not of the same type.

The `SetConstructor` and `SetImage` methods, along with the quantifiers `ForAll` and `Exists` and the `Choose` method, make use of the `reflect` package. This drastically improves the readability of the compiled Go program, but it also incurs a significant performance penalty. Performance-critical sections of the application should be rewritten to avoid use of these methods.

Since slices are mutable, but PlusCal tuples are immutable, slices are cloned before insertion into a set.

The `pgoutil` library exposes a `Set` interface. The developer modifying the Go code may wish to write a custom set container (say, which is specialized for a specific data type). If this container implements the `pgoutil.Set` interface, it is interoperable with the builtin set.

4.3 Tuples

The `pgoutil.Tuple` implements an immutable tuple type which holds elements of arbitrary type, backed by a slice. The `At` and `Set` tuple methods panic if the index accessed is out of range. The `Set` and `Append` tuple methods create a copy of the tuple, to preserve immutability. (However, the `Head` and `Tail` methods do not.)

The `pgoutil` library exposes a `Tuple` interface. Any custom container implementing this interface is *not* interoperable with the builtin tuple.

5 Example programs

5.1 Euclidean algorithm

The Euclidean algorithm is a simple algorithm that computes the greatest common divisor of two integers, and is a good example PlusCal algorithm.

```

1 ----- MODULE Euclid -----
2 EXTENDS Naturals, TLC
3 CONSTANT N
4
5 (*
6 --algorithm Euclid {
7 (* @PGo{ arg N int }@PGo *)
8   variables u = 24;
9         v \in 1 .. N;
10        v_init = v;
11
12   {
13     a: while (u # 0) {
14       if (u < v) {
15         u := v || v := u;
16       };
17     b: u := u - v;
18   };
19   print <<24, v_init, "have gcd", v>>
20 }
21 *)
22 =====

```

PlusCal

```

1 package main
2
3 import (
4   "flag"
5   "fmt"
6   "pgoutil"
7   "strconv"
8 )
9
10 var v int // PGo inferred type int
11 var u int // PGo inferred type int

```

```

12 var v_init int // PGo inferred type int
13 var N int
14
15
16 func main() {
17     flag.Parse()
18     N, _ = strconv.Atoi(flag.Args()[0])
19
20     for v_interface := range pgoutil.Sequence(1, N).Iter() {
21         v = v_interface.(int)
22         u = 24
23         v_init = v
24         for u != 0 {
25             if u < v {
26                 u_new := v
27                 v_new := u
28                 u = u_new
29                 v = v_new
30             }
31             u = u - v
32         }
33         fmt.Printf("%v %v %v %v\n", 24, v_init, "have gcd", v)
34     }
35 }

```

Compiled Go

The constant N needs to be annotated since its declaration does not appear in the comment containing the algorithm. Note the code to swap u and v in lines 26-29 of the Go program.

5.2 N-Queens problem

This PlusCal algorithm computes all possible ways to place n queens on an $n \times n$ chessboard such that no two queens attack each other. It demonstrates the expressive power of PlusCal's set constructs, as the algorithm is very concise.

```

1  ----- MODULE QueensPluscal -----
2  EXTENDS Naturals, Sequences
3  (*****
4  (* Formulation of the N-queens problem and an iterative algorithm to solve *)
5  (* the problem in TLA+. Since there must be exactly one queen in every row *)
6  (* we represent placements of queens as functions of the form *)
7  (* queens \in [ 1..N -> 1..N ] *)
8  (* where queens[i] gives the column of the queen in row i. Note that such *)
9  (* a function is just a sequence of length N. *)
10 (* We will also consider partial solutions, also represented as sequences *)
11 (* of length \leq N. *)
12 (*****
13
14 CONSTANT N \** number of queens and size of the board
15 ASSUME N \in Nat \ {0}
16
17 (* The following predicate determines if queens i and j attack each other
18    in a placement of queens (represented by a sequence as above). *)
19 Attacks(queens,i,j) ==
20   \ / queens[i] = queens[j] \** same column
21   \ / queens[i] - queens[j] = i - j \** first diagonal
22   \ / queens[j] - queens[i] = i - j \** second diagonal
23
24 (* A placement represents a (partial) solution if no two different queens
25    attack each other in it. *)
26 IsSolution(queens) ==

```

```

27 \A i \in 1 .. Len(queens)-1 : \A j \in i+1 .. Len(queens) :
28   ~ Attacks(queens,i,j)
29
30 (*****
31 (* We now describe an algorithm that iteratively computes the set of *)
32 (* solutions of the N-queens problem by successively placing queens. *)
33 (* The current state of the algorithm is given by two variables: *)
34 (* - todo contains a set of partial solutions, *)
35 (* - sols contains the set of full solutions found so far. *)
36 (* At every step, the algorithm picks some partial solution and computes *)
37 (* all possible extensions by the next queen. If N queens have been placed *)
38 (* these extensions are in fact full solutions, otherwise they are added *)
39 (* to the set todo. *)
40 (*****
41
42 (* --algorithm QueensPluscal
43   (** @PGO{ arg N int }@PGO
44     @PGO{ var todo set[[]int] }@PGO
45     @PGO{ var sols set[[]int] }@PGO
46     @PGO{ def Attacks(queens []int,i int,j int) ==
47       \ / queens[i] = queens[j] \ ** same column
48       \ / queens[i] - queens[j] = i - j \ ** first diagonal
49       \ / queens[j] - queens[i] = i - j \ ** second diagonal }@PGO
50     @PGO{ def IsSolution(queens []int) ==
51       \A i \in 1 .. Len(queens)-1 : \A j \in i+1 .. Len(queens) :
52         ~ Attacks(queens,i,j) }@PGO **)
53
54     variables
55       todo = { << >> };
56       sols = {};
57
58     begin
59     nxtQ: while todo # {}
60       do
61         with queens \in todo,
62           nxtQ = Len(queens) + 1,
63           cols = { c \in 1..N : ~ \E i \in 1 .. Len(queens) :
64             Attacks( Append(queens, c), i, nxtQ ) },
65           exts = { Append(queens,c) : c \in cols }
66         do
67           if (nxtQ = N)
68             then todo := todo \ {queens}; sols := sols \union exts;
69             else todo := (todo \ {queens}) \union exts;
70             end if;
71           end with;
72         end while;
73       end algorithm
74 *)
=====

```

PlusCal

```

1 package main
2
3 import (
4   "flag"
5   "math/rand"
6   "pgoutil"
7   "strconv"
8   "time"
9 )
10

```



```

11 var todo pgoutil.Set = pgoutil.NewSet([]int{})
12 var sols pgoutil.Set = pgoutil.NewSet()
13 var N int
14
15 func Attacks(queens []int, i int, j int) bool {
16     return queens[i - 1] == queens[j - 1]
17         || queens[i - 1] - queens[j - 1] == i - j
18         || queens[j - 1] - queens[i - 1] == i - j
19 }
20 func IsSolution(queens []int) bool {
21     return pgoutil.ForAll(func(i int) bool {
22         return pgoutil.ForAll(func(j int) bool {
23             return !Attacks(queens, i, j)
24         }, pgoutil.Sequence(i + 1, len(queens)))
25     }, pgoutil.Sequence(1, len(queens) - 1))
26 }
27
28 func main() {
29     rand.Seed(time.Now().UnixNano())
30     flag.Parse()
31     N, _ = strconv.Atoi(flag.Args()[0])
32
33     for !pgoutil.NewSet().Equal(todo) {
34         {
35             var queens []int // PGo inferred type []int
36                 = todo.ToSlice()[rand.Intn(todo.Size())].([]int)
37             var nxtQ int = len(queens) + 1 // PGo inferred type int
38             var cols pgoutil.Set // PGo inferred type set[int]
39                 = pgoutil.SetConstructor(pgoutil.Sequence(1, N), func(c int) bool {
40                 return !pgoutil.Exists(func(i int) bool {
41                     return Attacks(append(queens, c), i, nxtQ)
42                 }, pgoutil.Sequence(1, len(queens)))
43             })
44             var exts pgoutil.Set // PGo inferred type set[[]int]
45                 = pgoutil.SetImage(func(c int) []int {
46                 return append(queens, c)
47             }, cols)
48             if (nxtQ == N) {
49                 todo = todo.Difference(pgoutil.NewSet(queens))
50                 sols = exts.Union(sols)
51             } else {
52                 todo = exts.Union((todo.Difference(pgoutil.NewSet(queens))))
53             }
54         }
55     }
56 }

```

Compiled Go

(Some lines were broken in order to fit on the page.)

Note that the variables `cols` and `exts` do not require a type annotation, though they involve PlusCal tuples, since PGo can infer their type based on the type of `todo`.

5.3 Dijkstra's mutex algorithm

This is a multiprocess algorithm which only allows one process to be in the critical section at one time.

```

1 ----- MODULE DijkstraMutex -----
2 EXTENDS Integers
3 (*****
4 (* There is no reason why the processes need to be numbered from 1 to N. *)
5 (* So, we assume an arbitrary set Proc of process names. *)

```

```

6  (*****
7  CONSTANT Proc
8
9  (*****
10 Here is the PlusCal version of this algorithm.
11 The algorithm was modified from the original by adding the variable temp2,
12 to avoid a type consistency conflict when temp changes type at Li4a.
13 (* @PGo{ def Proc == 1 .. 10 }@PGo
14   @PGo{ lock false }@PGo
15 *)
16
17 --algorithm Mutex
18 { variables b = [i \in Proc |-> TRUE], c = [i \in Proc |-> TRUE], k \in Proc;
19   process (P \in Proc)
20     variable temp, temp2 ;
21     { Li0: while (TRUE)
22       { b[self] := FALSE;
23         Li1: if (k # self) { Li2: c[self] := TRUE;
24           Li3a: temp := k;
25             Li3b: if (b[temp]) { Li3c: k := self } ;
26               Li3d: goto Li1
27             };
28           Li4a: c[self] := FALSE;
29             temp2 := Proc \ {self};
30             Li4b: while (temp2 # {})
31               { with (j \in temp2)
32                 { temp2 := temp2 \ {j};
33                   if (~c[j]) { goto Li1 }
34                 }
35               };
36             cs: skip; /* the critical section
37             Li5: c[self] := TRUE;
38             Li6: b[self] := TRUE;
39             ncs: skip /* non-critical section ("remainder of cycle")
40           }
41         }
42     }
43 *****
44 =====

```

PlusCal

```

1  package main
2
3  import (
4    "math/rand"
5    "pgoutil"
6    "sync"
7    "time"
8  )
9
10 var Proc pgoutil.Set = pgoutil.Sequence(1, 10)
11 var k int // PGo inferred type int
12 var b pgoutil.Map // PGo inferred type map[int]bool
13 var c pgoutil.Map // PGo inferred type map[int]bool
14 var PGoWait sync.WaitGroup
15 var PGoStart chan bool = make(chan bool)
16
17 func P(self int) {
18   var temp int // PGo inferred type int
19   var temp2 pgoutil.Set // PGo inferred type set[int]

```

```

20     defer PGoWait.Done()
21     <-PGoStart
22     for true {
23         b.Put(self, false)
24         Li1:
25         if k != self {
26             c.Put(self, true)
27             temp = k
28             if b.Get(temp).(bool) {
29                 k = self
30             }
31             goto Li1
32         }
33         c.Put(self, false)
34         temp2 = Proc.Difference(pgoutil.NewSet(self))
35         for !pgoutil.NewSet().Equal(temp2) {
36             {
37                 var j int = temp2.ToSlice()[rand.Intn(temp2.Size())].(int) // PGo inferred type int
38                 temp2 = temp2.Difference(pgoutil.NewSet(j))
39                 if !c.Get(j).(bool) {
40                     goto Li1
41                 }
42             }
43         }
44         // TODO skipped from pluscal
45         c.Put(self, true)
46         b.Put(self, true)
47         // TODO skipped from pluscal
48     }
49 }
50
51 func main() {
52     rand.Seed(time.Now().UnixNano())
53     for k_interface := range Proc.Iter() {
54         k = k_interface.(int)
55         b = pgoutil.MapsTo(func(i int) bool {
56             return true
57         }, Proc)
58         c = pgoutil.MapsTo(func(i int) bool {
59             return true
60         }, Proc)
61         for procId := range Proc.Iter() {
62             PGoWait.Add(1)
63             go P(procId.(int))
64         }
65         close(PGoStart)
66         PGoWait.Wait()
67     }
68 }

```

Compiled Go

The Go statements used in compilation are atomic on a usual configuration, so locking is turned off. (All `pgoutil.Map` and `pgoutil.Set` operations are atomic.)

For reference, here is the compiled code when locking is turned on:

```

1 package main
2
3 import (
4     "math/rand"
5     "pgoutil"
6     "sync"

```

```

7     "time"
8 )
9
10 var PGoLock []sync.RWMutex = make([]sync.RWMutex, 3)
11 var Proc pgoutil.Set = pgoutil.Sequence(1, 10)
12 var k int // PGo inferred type int; Lock group 2
13 var b pgoutil.Map // PGo inferred type map[int]bool; Lock group 0
14 var c pgoutil.Map // PGo inferred type map[int]bool; Lock group 1
15 var PGoWait sync.WaitGroup
16 var PGoStart chan bool = make(chan bool)
17
18 func P(self int) {
19     var temp int // PGo inferred type int
20     var temp2 pgoutil.Set // PGo inferred type set[int]
21     defer PGoWait.Done()
22     <-PGoStart
23     for true {
24         PGoLock[0].Lock()
25         b.Put(self, false)
26         PGoLock[0].Unlock()
27         Li1:
28         if k != self {
29             PGoLock[1].Lock()
30             c.Put(self, true)
31             PGoLock[1].Unlock()
32             temp = k
33             if b.Get(temp).(bool) {
34                 PGoLock[2].Lock()
35                 k = self
36                 PGoLock[2].Unlock()
37             }
38             goto Li1
39         }
40         PGoLock[1].Lock()
41         c.Put(self, false)
42         temp2 = Proc.Difference(pgoutil.NewSet(self))
43         PGoLock[1].Unlock()
44         for !pgoutil.NewSet().Equal(temp2) {
45             {
46                 var j int = temp2.ToSlice()[rand.Intn(temp2.Size())].(int) // PGo inferred type int
47                 temp2 = temp2.Difference(pgoutil.NewSet(j))
48                 if !c.Get(j).(bool) {
49                     goto Li1
50                 }
51             }
52         }
53         // TODO skipped from pluscal
54         PGoLock[1].Lock()
55         c.Put(self, true)
56         PGoLock[1].Unlock()
57         PGoLock[0].Lock()
58         b.Put(self, true)
59         PGoLock[0].Unlock()
60         // TODO skipped from pluscal
61     }
62 }
63
64 func main() {
65     rand.Seed(time.Now().UnixNano())
66     for k_interface := range Proc.Iter() {
67         k = k_interface.(int)
68         b = pgoutil.MapsTo(func(i int) bool {

```

```

69         return true
70     }, Proc)
71     c = pgoutil.MapsTo(func(i int) bool {
72         return true
73     }, Proc)
74     for procId := range Proc.Iter() {
75         PGoWait.Add(1)
76         go P(procId.(int))
77     }
78     close(PGoStart)
79     PGoWait.Wait()
80 }
81 }

```

Compiled Go with locks

Note that the loop at line 44 does not require any locking, since it only involves local variables.

The constant `Proc` is defined to be `1 .. 10` in the annotation. If the process set needs to be changed, only the annotation needs to be edited.