

<http://golang.org>



The Expressiveness of Go

Rob Pike
OSCON
July 27, 2011



<http://golang.org>

Google™

Why Go?

A response to Google's internal needs:

- efficient large scale programming
- speed of compilation
- distributed systems
- multicore, networked hardware

And a reaction to: “speed and safety or ease of use; pick one.”

- complexity, weight, noise (C++, Java)
vs.
- no static checking (JavaScript, Python)

Go is statically typed and compiled, like C++ or Java (with no VM), but in many ways feels as lightweight and dynamic as JavaScript or Python.

The surprise

It turned out to be a nice general purpose language.

That was unexpected.

The most productive language I've ever used.
And some people agree.

Acceptance

Go was the 2009 TIOBE "Language of the year" two months after it was released and it won an InfoWorld BOSSIE award.

Year	Winner
2009	Go
2008	C
2007	Python
2006	Ruby
2005	Java
2004	PHP
2003	C++



Why the success?

This acceptance was a pleasant surprise.

But in retrospect, the way we approached the design was important to the expressiveness and productivity of Go.

"I like a lot of the design decisions they made in the language. Basically, I like all of them."

- Martin Odersky, creator of Scala

Principles

The axioms of Go's design



Principles

Go is:

Simple

- concepts are easy to understand
- (the implementation might still be sophisticated)

Orthogonal

- concepts mix cleanly
- easy to understand and predict what happens

Succinct

- no need to predeclare every intention

Safe

- misbehavior should be detected

These combine to give expressiveness.



Respect

Go trusts the programmer to write down what is meant.

In turn, Go tries to respect the programmer's intentions.

It's possible to be safe and fun.

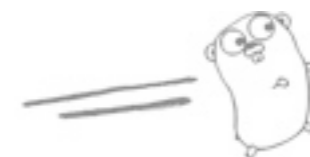
There's a difference between seat belts and training wheels.

Simplicity

Number of keywords is an approximate measure of complexity.

C (K&R)	K&R	32
C++	1991	48
Java	3rd edition	50
C#	2010	77
C++0x	2010	72+11*
JavaScript	ECMA-262	26+16*
Python	2.7	31
Pascal	ISO	35
Modula-2	1980	40
Oberon	1990	32
Go	2010	25


9 * extra count is for reserved words and alternate spellings



Google™

An example

A complete (if simple) web server



Hello, world 2.0

Serving `http://localhost:8080/world`:

```
package main
```

```
import (  
    "fmt"  
    "http"  
)
```


```
func handler(w http.ResponseWriter, r *http.Request) {  
    fmt.Fprintf(w, "Hello, %s.", r.URL.Path[1:])  
}
```

```
func main() {  
    http.ListenAndServe(":8080",  
                        http.HandlerFunc(handler))  
}
```




Stories

A few design tales that illustrate how the principles play out.
Not close to a complete tour of Go.



Basics

Some fundamentals



Starting points

Go started with a few important simplifications relative to C/C++, informed by our experience with those languages:

There are pointers but no pointer arithmetic

- pointers are important to performance, pointer arithmetic not.
- although it's OK to point inside a struct.
- important to control layout of memory, avoid allocation.

Increment/decrement (`i++`) is a statement, not an expression.

- no confusion about order of evaluation.

Addresses last as long as they are needed.

- take the address of a local variable, the implementation guarantees the memory survives while it's referenced.

No implicit numerical conversions (float to integer, etc.).

- C's "usual arithmetic conversions" are a minefield.

Constants are ideal

Implicit conversions often involve constants (`sin(Pi/4)`), but Go mitigates the conversion issue by having nicer constants.

Constants are "ideal numbers": no size or sign or precision limit, hence no `L` or `U` or `UL` endings.

```
077 // octal
```

```
0xFEEDBEEEEEEEEEEEEEEEEEEEF // hexadecimal
```

```
1 << 100
```

```
Pi = 3.1415926535897932384626433832795028841971
```

They are just numbers that can be assigned to variables; no conversions necessary. A typed element in the expression sets the true type of the constant. Here `1e9` becomes `int64`.


```
seconds := time.Nanoseconds()/1e9
```

Fully type safe and orthogonal to type system.



Types and data

Structs, methods, and interfaces



Structs

Structs describe (and control) the layout of data.

Some early proposals included methods in the struct, but that idea was dropped. Instead methods are declared like ordinary functions, outside the type, and with a "receiver".

```
type Coord struct { x, y float64 }
```

```
func (c Coord) Abs() float64 {  
    return math.Sqrt(c.x*c.x + c.y*c.y)  
}
```

The `(c Coord)` declares the receiver (no automatic "this" variable; also notice `c` is not a pointer, although it could be.)

Methods are not mixed with the data definition.

They are orthogonal to types.

Methods are orthogonal to types

Orthogonality of methods allows any type to have them.

```
type Vector []float64
func (v Vector) Abs() float64 {
    sumOfSquares := 0.0
    for i := range v {
        sumOfSquares += v[i]*v[i]
    }
    return math.Sqrt(sumOfSquares)
}
```

It also allows receivers to be values or pointers:

```
func (c *Coord) Scale(ratio float64) {
    c.x, c.y = ratio*c.x, ratio*c.y
}
```

Orthogonality leads to generality.



Interfaces

Interfaces are just sets of methods; work for any type.

```
type Abser interface {  
    Abs() float64  
}
```

```
var a Abser  
a = Coord{3, 4}  
print(a.Abs())  
a = Vector{1, 2, 3, 4}  
print(a.Abs())
```

Interfaces are satisfied implicitly. `Coord` and `Vector` do not declare that they implement `Abser`, they just do.

@mjmoriarity: The way Go handles interfaces is the way I wish every language handled them.

Interfaces are abstract, other types are concrete

In some of our early designs, interfaces could contain data, but this conflated representation and behavior.

Made them distinct:

- concrete types such as structs define data
- interfaces define behavior

As with methods, now anything can satisfy an interface.

```
type Value float64 // basic type
func (v Value) Abs() float64 {
    if v < 0 { v = -v }
    return float64(v)
}
```

```
a = Value(-27)
print(a.Abs())
```

Interfaces are satisfied implicitly

`Coord` and `Vector` and `Value` satisfy `Abser` implicitly; other types might too. A type satisfies an interface simply by implementing its methods. There is no "implements" declaration; interfaces are satisfied implicitly.

It's a form of duck typing, but (usually) checkable at compile time. It's also another form of orthogonality.

An object can (and usually does) satisfy many interfaces simultaneously. For instance, `Coord` and `Vector` satisfy `Abser` and also the empty interface: `interface{}`, which is satisfied by any value (analogous to C++ `void*` or Java `Object`) even of basic type (unlike `void*` or `Object`).

In Go, interfaces are usually one or two (or zero) methods.

Reader

```
type Reader interface {  
    Read(p []byte) (n int, err os.Error) // two returns  
}  
// And similarly for Writer
```

Anything with a **Read** method implements **Reader**.

- Sources: files, buffers, network connections, pipes
- Filters: buffers, checksums, decompressors, decrypters

JPEG decoder takes a **Reader**, so it can decode from disk, network, gzipped HTTP,

Buffering just wraps a **Reader**:

```
var bufferedInput Reader = bufio.NewReader(os.Stdin)
```

Fprintf uses a **Writer**:

```
func Fprintf(w Writer, fmt string, a ...interface{})
```

Interfaces enable retrofitting

Had an existing RPC implementation that used custom wire format. Changed to an interface:

```
type Encoding interface {  
    ReadRequestHeader(*Request) os.Error  
    ReadRequestBody(interface{}) os.Error  
    WriteResponse(*Response, interface{}) os.Error  
    Close() os.Error  
}
```

Two functions (send, receive) changed signature. Before:

```
func sendResponse(sending *sync.Mutex, req *Request,  
    reply interface{}, enc *gob.Encoder, err string)
```

After (and similarly for receiving):

```
func sendResponse(sending *sync.Mutex, req *Request,  
    reply interface{}, enc Encoding, err string)
```

That is almost the whole change to the RPC implementation.

Post facto abstraction

We saw an opportunity: RPC needed only `Encode` and `Decode` methods. Put those in an interface and you've abstracted the codec.

Total time: 20 minutes, including writing and testing the JSON implementation of the interface.

(We also wrote a trivial wrapper to adapt the existing codec for the new `rpc.Encoding` interface.)

In Java, RPC would be refactored into a half-abstract class, subclassed to create `JsonRPC` and `StandardRPC`. You'd have to plan ahead.

In Go it's simpler and the design adapts through experience.

Principles redux

Type and data examples:

Simple

- interfaces are just method sets

Orthogonal

- representation (data) and behavior (methods) are independent
- empty interface can represent any value

Succinct

- no "implements" declarations; interfaces just satisfy


Safe

- static type checking

Expressiveness: implicit satisfaction lets pieces fit together seamlessly.

Names

Visibility



Visibility

With methods not declared as part of the type definition, how to say private/public?

Long design debate with several suggestions:

- placement in the file
- **export** keyword
- name marker (e.g. `Coord{ +x,+y int }`)

Resolution (anguished decision): Don't make it part of the declaration, make it part of the name. After all, that's what you see whenever you use it!

Case of first character determines visibility outside package:

ThisNameIsPublic
thisOneIsNot

Visibility is orthogonal to type

One of the best decisions in the language, yet really hard to make!

Lose control over how to use case, e.g. can't use it to distinguish Types from variables.

In return, though:


- extremely simple, easy rule to understand
- consequences clear
- see a variable, can see whether it's public without going to the declaration
- can make any type, variable, field, method, function or constant public or not with the same mechanism

Orthogonality again!



Concurrency and closures

Goroutines, channels, stacks and closures



The model

Go has native support for concurrent operations in the CSP tradition.

Two styles of concurrency exist in CS: deterministic (well-defined ordering) and non-deterministic (mutual exclusion but order undefined).

Go's goroutines and channels promote* deterministic concurrency (channels with one sender, one receiver), which is easier to reason about.

Simpler for the programmer.

* promote, not require. There are many possibilities.

Go concurrency basics

Start a goroutine:

```
go f()
```

Channel send (arrow points in direction of data flow):

```
ch <- value
```

Channel receive:

```
value = <-ch
```

Channels are unbuffered by default, which combines synchronization with communication.



A simple worker pool

A unit of work:

```
type Work struct { x, y, z int }
```

A worker task:

```
func worker(in <-chan *Work, out chan <- *Work) {  
    for w := range in {  
        w.z = w.x * w.y  
        out <- w  
    }  
}
```

Driver:

```
func Run() {  
    in, out := make(chan *Work), make(chan *Work)  
    for i := 0; i < 10; i++ { go worker(in, out) }  
    go sendLotsOfWork(in)  
    receiveLotsOfResults(out)  
}
```



Secondary support

To make concurrency feasible, need several things:

- language support (axiomatic)
- garbage collection (near axiomatic)

To make it good, you need:

- stack management
- closures

Stacks

Goroutines have "segmented stacks":

```
go f()
```

starts `f()` executing concurrently on a new stack.

Stack grows and shrinks as needed.

No programmer concern about stack size.

No possibility for stack overflow.

A couple of instructions of overhead on each function call, a huge improvement in simplicity and expressiveness.

Concurrent execution is orthogonal to everything else.



Launching a goroutine

Start a service, return a channel to communicate with it:

```
func (s *Service) Start() chan request {  
    ch := make(chan request)  
    go s.serve(ch) // s.serve runs concurrently  
    return ch     // returns immediately  
}
```

A common pattern, given channels as first-class values.

Closures

Closures are just local functions

```
func Compose(f, g func(x float64) float64)
    func(x float64) float64 {
        return func(x float64) float64 {
            return f(g(x))
        }
    }
```

```
func main() {
    print(Compose(sin, cos)(0.5))
}
```

Fit easily into implementation since local variables already move to heap when necessary.

Closures and concurrency

Query servers in replicated database, return first response.

```
func Query(conns []Conn, query string) Result {
    ch := make(chan Result, 1) // buffer of 1 item
    for _, conn := range conns {
        go func(c Conn) {
            select {
            case ch <- c.DoQuery(query):
                // nothing to do
            default: // executes if ch is blocked
                // nothing to do
            }
        }(conn)
    }
    return <-ch
}
```

Principles redux

Concurrency and closure examples:

Simple

- stacks just work; goroutines too cheap to meter

Orthogonal

- concurrency orthogonal to rest of language
- orthogonality of functions make closures easy

Succinct

- `go f()`
- closure syntax clear

Safe

- no stack overflows
- garbage collection avoids many concurrency problems

Expressiveness: complex behaviors easily expressed.



Conclusion

Expressiveness comes from orthogonal composition of simple concepts.



Conclusion

Go is not a small language (goroutines, channels, garbage collection, methods, interfaces, closures, ...) but it is an expressive and comprehensible one.

Expressiveness comes from orthogonal composition of constructs.

Comprehensibility comes from simple constructs that interact in easily understood ways.

Build a language from simple orthogonal constructs and you have a language that will be easy and productive to use.

The surprises you discover will be pleasant ones.

Implementation

The language is designed and usable. Two compiler suites:

Gc, written in C, generates OK code very quickly.

- unusual design based on the Plan 9 compiler suite

Gccgo, written in C++, generates good code more slowly

- uses GCC's code generator and tools

Libraries good and growing, but some pieces are still preliminary.

Garbage collector works well (mark and sweep) but is being rewritten for more concurrency, less latency.

Available for Linux etc., Mac OS X, Windows.

All available as open source; see <http://golang.org>.



Go runtime available on Google App Engine

- A high performance language for App Engine
 - easy and productive to use.
 - scales well.
 - good (and ever-improving) library support.
 - efficient – a true compiled language.
- Full source code for SDK available (open source).
- Full details at available through "Go for Google App Engine" link on Go page at:
 - <http://golang.org>


Try it out

This is a true open source project.

Much more information at

<http://golang.org>

including full source, documentation, links to resources and the App Engine SDK, and even a playground that lets you try Go code right from the browser.




<http://golang.org>



The Expressiveness of Go

Rob Pike
OSCON
July 27, 2011



<http://golang.org>

Google™