

Go-HEP: writing concurrent software with ease and Go

S Binet¹

¹ Laboratoire de Physique de Clermont-Ferrand, CNRS/IN2P3, FR

E-mail: `sebastien.binet@clermont.in2p3.fr`

Abstract. High Energy and Nuclear Physics (HENP) libraries are now required to be more and more multi-thread-safe, if not multi-thread-friendly and multi-threaded. This is usually done using the new constructs and library components offered by the C++11 and C++14 standards. These components are however quite low-level and hard to use and compose.

However, Go provides a set of better building blocks for tackling concurrency: goroutines and channels. This language is now used by the whole cloud industry: docker/moby, rkt, Kubernetes are obvious flagships for Go. But to be able to perform any meaningful physics analysis, one needs a set of basic libraries (matrix operations, linear algebra, plotting, I/O, ...) We present Go-HEP, a set of packages to easily write concurrent software to interface with legacy HENP C++ physics libraries.

1. Introduction

With the end of Dennard Scaling, High Energy and Nuclear Physics (HENP) libraries and applications are now required to be multi-thread safe. This is to ensure the scaling of performances with new CPU architectures. The new C++ standards introduced new constructs and library components: `std::thread`, `std::mutex`, `std::lock`, ... These components are however quite low-level and hard to use and compose, or easy to misuse. Moreover, C++ is still plagued with scalability issues. Development speed is hindered by the interplay between the compilation model (based on `#include`) and C++ templates. C++ is a very large language, hard and subtle to teach and understand: this has an impact on code maintenance and the ability to bring newcomers up-to-speed on any given project. On the user side, C++ is showing deficiencies as well: installing a project's many dependencies can be quite difficult. Python users will be surprised to learn there is still no pip-like mechanism to install, automatically and recursively, dependencies for a given project. Finally, most HEP software stacks rely on applications dynamically linked to hundreds shared libraries: packaging and deployment can be time consuming tasks.

Fixing the whole C++ ecosystem sounds like a rather daunting task. It might be easier instead to start from a blank page, with a new language that addresses all the current deficiencies of C++, and improves the day-to-day life of a typical software programmer in a multi-core environment.

Go [1] was created in 2009 to address these issues. We will first describe briefly the Go programming language, its concurrency primitives and the features of its tooling that make Go a great fit for HENP applications. We will then introduce Go-HEP, a set of libraries and applications that aims to provide physicists with the basic tools to perform HENP analyses and integrate them in already existing C++/Python analysis pipelines.

2. Introduction to Go

Go is an open source language, released under the BSD-3 license in 2009. It is a compiled language with a garbage collector and builtin support for reflection. The language is reminiscent of C/C++; it shows a syntax similar to its older peers but adds first-class functions, closures and object-oriented programming via the concept of interfaces. Go is already available, via the `gc` toolchain, on all major platforms (Linux, Windows, macOS, Android, iOS, ...) and for many architectures (`amd64`, `arm64`, `i386`, `s390x`, `mips64`, ...) There is also a Go frontend to GCC, aptly named `gccgo`, that can target all the platforms and architectures that the GCC toolchain supports. Griesemer, Pike and Thompson created Go to replace the multi-threaded C++ web servers that were hard to develop and slow to compile, even when using Google's massive infrastructure. As such, Go exposes two builtin concurrency primitives: the *goroutines* – very lightweight green threads – and the *channels* – typed conduits that connect goroutines together.

Launching goroutines is done by prepending a function (or method) call with the keyword `go`. Goroutines can be described as green threads: very lightweight execution contexts that are multiplexed on native threads. Each goroutine starts with a small stack (around *4KB*) that can grow and shrink as needed. This allows to write applications with thousands of goroutines without the need to buy high-end servers. Writing real world concurrent programs needs synchronization and communication: the ability to exchange data between goroutines. This is achieved *via* channels. Not only do channels allow to exchange data in a type safe manner but they are also a synchronization point: a goroutine trying to send a token of data through a channel will block until there is another goroutine on the other end of the channel trying to extract data from the channel, and *vice versa*. The last piece that makes building concurrent programs in Go efficient is the keyword `select`. This keyword is like a `switch` statement for controlling the data flow between multiple channels and thus between multiple goroutines. The concurrency builtin tools of Go are easier to reason about and more composable than mutexes and locks.

Moreover, Go comes with a package system and a builtin tool to build Go code. There are no header files like in C++ as they require to recursively process all dependent headers and thus slow the build. Once compiled, Go packages are completely self-contained and do not require a client of package `p1` – which itself depends on `p2` and `p3` – to know anything about these packages. The client only needs `p1`. This greatly improves the scalability of the build process as well as its speed. Thanks to the way third-party packages are imported in Go, *e.g.* `import "github.com/pkg/errors"`, a build tool only needs to parse and understand Go code to discover (recursively) dependencies of a given package. Coupled to the fact that package import strings contain the URL to the repository (GitHub, BitBucket, ...) holding the actual code, this allows a simple command – `go get` – to fetch code from the internet, (recursively) discover and fetch its dependencies and, finally build the whole artefact. The instructions to install code are valid on all platforms that are supported by a given Go toolchain: one just needs to point `go get` at a repository. Finally, Go code being very quick to compile makes static compilation manageable and enables simple deployment scenarii that boil down to `scp`-ing the resulting binary. Go exposes a productive development environment that is concurrency friendly. It is being used by many companies ¹ beside Google, and in a variety of scenarii: from rocket telemetry to cloud systems to container orchestration. But for Go to be useful in HENP, physicists need libraries and tools to carry analyses. Moreover, these tools need to be interoperable with existing analyses pipelines.

3. Go-HEP

Go-HEP is a set of libraries and applications released under the BSD-3 license. They allow High Energy and Nuclear Physicists to write efficient analysis code in the Go programming language.

¹ A non-exhaustive list is maintained here: <https://github.com/golang/go/wiki/GoUsers>.

The `go-hep.org/x/hep` organization provides a set of pure-Go packages and building blocks to:

- write analyses in Go,
- write data acquisition and monitoring systems in Go,
- write control frameworks in Go.

As Go-HEP is written in pure-Go, the whole software suite and its dependencies are installable on all Go supported platforms (Linux, macOS, Windows, RPi3, ...) with:

```
$> go get go-hep.org/x/hep/...
```

The ellipsis (...) at the end instructs the `go` tool to compile and install all the libraries and applications that are part of the `go-hep.org/x/hep` repository.

The libraries provided by Go-HEP can be broadly organized around two categories:

- libraries that provide physics and statistical tools (Lorentz vectors, histograms and n-tuples, fits, jet clustering, fast detector simulation, plots, ...)
- libraries that provide low level interoperability with C++ libraries, to allow Go-HEP users to integrate with existing analyses pipelines.

Indeed, analyses in HENP are mainly written in C++ and Python. Even though Go has a native foreign function interface called `cgo` that allows Go code to use C code and vice versa, the libraries under `go-hep.org/x/hep` do not call any C++ library (via a C shim library.) This decision allows to retain the quick edit-compile-run development cycle of Go and the easy deployment and cross-compilation of Go applications. Instead, the strategy for interoperating with C++ is to integrate with *e.g.* ROOT [3] at the data file level. Go-HEP provides read/write access for LCIO, LHEF, HepMC, SLHA and YODA files. Go-HEP provides – at the moment – only read access to ROOT files via its `go-hep.org/x/hep/rootio` package.

In the following, we will describe two components of Go-HEP, `fads` and `hep/rootio`, that enable physics analyses.

4. `fads`

4.1. `fads-app`

`fads` is a "FAst Detector Simulation toolkit". This library is built on top of Go-HEP's control framework, `fwk`. The control framework exposes a traditional API, with a `Task` type that can be started, process some event data and then stopped. Each event is processed by a goroutine. The framework runs these tasks in their own goroutine and lets the Go runtime deal with work stealing. Data between tasks are exchanged via an event store service which itself utilizes `channels` to ensure there are no data races. Besides an event store service, `fwk` also provides services for message logging, histogramming and n-tupling. N-tuple and histograms (1D and 2D) are provided by the `hep/hbook` and `hep/hbook/ntup` packages and are persisted using `hep/rio`, a binary file format that heavily draws inspiration from `SIO`, the binary file format the LCIO community uses.

Data dependencies between tasks are described at the job configuration level: each task declares what are its inputs (the type of input and a name that identifies it) and its outputs. `fwk` ensures there are no cycles between the tasks and connects the goroutines together via channels. `fwk` enables task-level parallelism and event-level parallelism via the concurrency building blocks of Go. For sub-task parallelism, users are required – by construction – to use the same building block, so everything is consistent and the Go runtime has the complete picture.

`fads` is itself a transliteration of Delphes [4] (v3.0.12) but with `fwk` underneath instead of ROOT. To assess the performances of Go and `fads` in a realistic settings, the whole ATLAS data card provided with Delphes has been implemented and packaged with `fads` under `hep/fads/cmd/fads-app`. This application is composed of:

- a HepMC file reader task,
- a particle propagator, a calorimeter simulator,
- energy rescalers, momentum smearers,
- electron, photon and muon isolation tasks,
- b-tagging and tau-tagging tasks, and a jet-finder task.

The jet finder task is based on a naive re-implementation of the C++ FastJet [6] library: only the N^3 "dumb" clustering has been implemented so far. A small part of the directed acyclic graph of the resulting `fads-app` application can be found in figure 1.

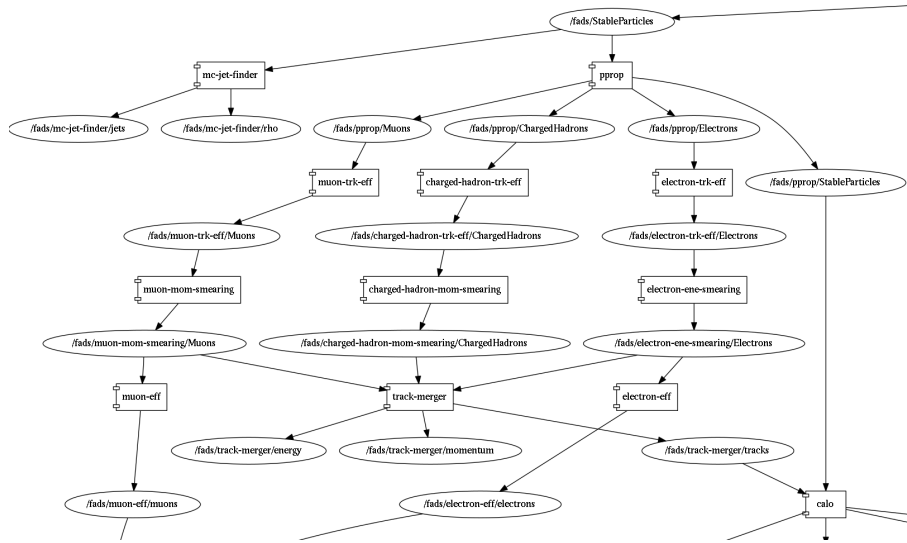


Figure 1. Part of the directed acyclic graph of data dependencies between tasks composing the `fads-app`, a Go transliteration of the Delphes’ ATLAS data card example. Rectangles are tasks, ellipses are data collections.

Delphes was compiled with `gcc-4.8` and the N^3 clustering strategy hard-coded, `fads` was compiled with `Go-1.9`. Timings were obtained on a Linux Xeon CPU E5-4620@2.20GHz server with 64 cores and 128Gb RAM with a 10000 HepMC events input file.

The results for various numbers of threads are shown in figure 2. As Delphes is not thread-safe, we only show the data for one Delphes process and then draw a flat line for visual aid. The graph on the left shows a smaller memory footprint for `fads` that only matches that of Delphes’ when 30 threads are used. The graph on the right shows that `fads` achieves a better event processing rate even when using only one thread, or when using only one thread but with the sequential event loop instead of the concurrent event loop (data point for $n_{thread} = 0$.)

`fads` achieves better performances than Delphes-3.0.12 on this data set. The output simulated data for `fads` was matched bit-by-bit with that of Delphes up to the calorimetry stage². `fads` also achieves these performances without the need to merge output files.

4.2. `fads-rivet-mc-generic`

`fads` exports another application, `fads-rivet-mc-generic`, a reimplement of the MC_GENERIC analysis from the Rivet [5] toolkit.

² We could not manage to match data further down the sequence because of Delphes’ usage of PRNG to seed other PRNGs in the calorimeter. Control histograms agreed at the statistical level.




Figure 2. Memory usage (left) and event processing rate (right) for `fads` (solid red curve) and `Delphes` (green dash line).

	MaxRSS (<i>Mb</i>)	Real (<i>s</i>)	CPU (<i>s</i>)
<code>Rivet</code>	27	13.3	13.3
<code>fads</code>	23	5.7	5.7

Table 1. Runtime performances of `Rivet` and `fads` to read a file with Z events.

`fads` shows better performances than the C++ application. As `Rivet` does not use ROOT for this application, the memory usage of `Rivet` is closer to that of `fads`. However, the event processing rate of `fads` with only one thread is twice as good as the one of `Rivet`. This is because the job steering and the event loop of `Rivet` are written in Python.

5. `hep/rootio`

The `hep/rootio` package is a pure-Go package that:

- decodes and understands the structure of `TFiles`, `TKeys`, `TDirectory` and `TStreamerInfos`,
- decodes and deserializes `TH1x`, `TH2x`, `TLeaf`, `TBranch` and `TTrees`.

At the moment, `hep/rootio` only concerns itself with reading ROOT files, although writing ROOT files is on the roadmap. Nonetheless, `hep/rootio` can already be useful and provides the following commands:

- `cmd/root-ls` lists the content of ROOT files,
- `cmd/root-print` extracts histograms from ROOT files and saves them as PDF or PNG,
- `cmd/root-dump` prints the contents of `TTrees`, event by event,
- `cmd/root-diff` prints the differences between the contents of two `TTrees`,
- `root-cnv-npy` converts simple `TTrees` to NumPy array data files,
- `cmd/root-srv` allows to interactively plot histograms from a file and branches from a tree. `root-srv` being pure-Go, it can be hosted on Google AppEngine.

`rootio` can read flat TTrees with C++ builtins, static and dynamic arrays of C++ builtins. `rootio` can also read TTrees with user defined classes containing `std::vector<T>`³, another user defined class, `std::string` or `TString`, arrays of C++ builtins and C++ builtins.

To assess the performances of `hep/rootio` with regard to the original C++ implementation, we created two input files of 10^6 entries and 100 branches of `float64`. One file was compressed with the default settings of ROOT-6.10 (686Mb) and the other had no compression (764Mb).

fCompression=0	VMem (Mb)	MaxRSS (Mb)	Real (s)	CPU (s)
ROOT	517	258	6.7	6.6
hep/rootio	43	42	12.9	12.9
fCompression=1	VMem (Mb)	MaxRSS (Mb)	Real (s)	CPU (s)
ROOT	529	292	12.6	12.0
hep/rootio	83	82	35.8	35.8

Table 2. Runtime performances of C++/ROOT and `hep/rootio` to read a file with 10^6 events, using no compression (up) and with default compression (down).

Table 2 shows the results obtained running the two programs with the compressed file and the uncompressed file as inputs. While the memory footprint of the Go program is almost an order of magnitude lower than its C++ counterpart, it is twice as slow in the non-compressed case and almost thrice in the compressed case. The degraded performances in the compressed case have been tracked back to the decompression package from the standard library which could probably be further optimized. On the otherhand, the degraded performances in the non-compressed case come from the still young implementation of `hep/rootio` which could be also further optimized, *e.g.* preemptively loading multiple consecutive entries.

6. Conclusions

Go improves on C++/Python and addresses their deficiencies with regard to code distribution, code installation, compilation, development and runtime speeds. Go also provides builtin facilities to tackle concurrency efficiently and easily. Go-HEP provides some building blocks that are already competitive with battle-tested C++ programs, both in terms of CPU, memory usage and cores' utilization. Further improvements are still necessary in the ROOT I/O area, both in terms of performances and features, so that Go-HEP can be part of physics analyses pipelines.

References

- [1] The Go programming language <https://golang.org>
- [2] S Binet, E Busato, D Blyth, M Ughetto, P Waller and B Wieck, Go-HEP: libraries for High Energy Physics analyses in Go, The Journal of Open Source Software, 2017/09, <https://doi.org/10.21105/joss.00372>
- [3] R Brun and F Rademakers, ROOT - An Object Oriented Data Analysis Framework, Proceedings AIHENP'96 Workshop, Lausanne, Sep. 1996, Nucl. Inst. & Meth. in Phys. Res. A 389 (1997) 81-86
- [4] The DELPHES 3 collaboration, DELPHES 3: a modular framework for fast simulation of a generic collider experiment, [software] [http://dx.doi.org/10.1007/JHEP02\(2014\)057](http://dx.doi.org/10.1007/JHEP02(2014)057)
- [5] A Buckley et al, Rivet user manual, [software] <https://doi.org/10.1016/j.cpc.2013.05.021>
- [6] M Cacciari, G Salam, G Soyez, FastJet manual, [software] <https://dx.doi.org/10.1140/epjc/s10052-012-1896-2>

³ where T is a C++ builtin or a `std::string` or a `TString`