

# Parallelization of Radio module using Gyges

---

A. Augusto Alves Jr, with Nikos Karastathis

Presented at CORSIKA development meeting - KIT, Karlsruhe

October 13, 2022



## Recap: Amdahl's law

- Predicts the expected speedup from parallelism:

Validity of the Single Processor Approach to Achieving  
Large-Scale Computing Capabilities

Amdahl, Gene M.

AFIPS Conference Proceedings (30): 483-485 (1967)

doi:10.1145/1465482.1465560

- It is expressed as

$$S(n) = \frac{1}{(1 - p) + \frac{p}{n}}$$

where:  $S(n)$  is the speedup in function of the number of cores/threads.  $n$  is number of cores/threads and  $p$  is the fraction of code that is parallelizable.

## CORSIKA, Radio and Amdahl's law

- Radio module calculates the signal corresponding to each particle/track for each antenna of the detector.
- And it often runs as one of the final operations in the sequence. Keep in mind that currently the all algorithms in the sequence run sequentially. Same applies for the processing of the particles in the stack.
- We parallelized the calculation of the signal over the detector. It means, in a per particle/track basis, the antennas response are processed in parallel.
- The expected overall speed-up (CORSIKA wise) then depends hugely of the detector size, i.e. number of antennas in the detector. I minor, but still significantly, it also depends on the calculations in the antenna itself.
- With a large detector, the importance Radio module operations grows and tends to dominate the sequence. In such, situations the speed-up is larger. This is what Amdahl's law predicts.

# Gyges

`Gyges` is a lightweight C++20 header-only library to manage thread pooling.

- With `Gyges`, thread creation and destruction costs can be paid just once in the program lifetime.
- Threads from the pool pick-up tasks as they became available. If there is no task, the threads just go sleeping.
- Tasks can be submitted from multiple threads. The submitter gets a `std::future` for monitoring the task in-place.
- Task assignment and running can be stopped at any time acting over an `std::stop_token`.
- A `gyges::gang` can also be created or put in a “hold-on” state. The processing of the tasks will be postponed until it is put on “unhold” status. The threads are not keep busy-waiting, they are put to sleep until “unhold” command is sent.
- Two implementations of `gyges::for_each`. One of them able to recycle an already existing `gyges::gang`.

## Some implementations hints, not details

---

- `corsika::RadioProcess` instances own the `gyges::gang`. Size of it can be specified at construction time :

---

```
1 auto propagator = make_simple_radio_propagator(enviroment);  
2 auto coreas     = make_radio_process_CoREAS(detector, propagator, nthreads);  
3 auto zhs       = make_radio_process_ZHS(detector, propagator, nthreads);
```

---

- No need to specify template parameters with the new interface.
- The implementation will define a task per thread (not per antenna). Each task will iterate over the same number of antennas, so that each thread will perform about the same amount of work.



# What is being profiled?

---


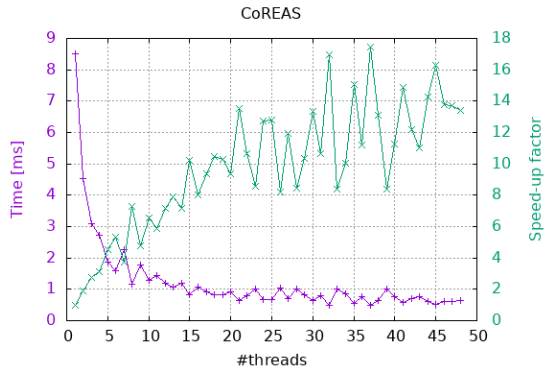
```
1 ...
2 auto propagator = make_simple_radio_propagator(enviroment);
3 auto coreas     = make_radio_process_CoREAS(detector, propagator, nthreads);
4 auto zhs       = make_radio_process_ZHS(detector, propagator, nthreads);
5 ...
6 //start chronometer
7 coreas.doContinuous(particle, base, true);
8 //stop chronometer
9 ...
10 //start chronometer
11 zhs.doContinuous(particle, base, true);
12 //stop chronometer
```

---

# Performance: detector with 1k antennas

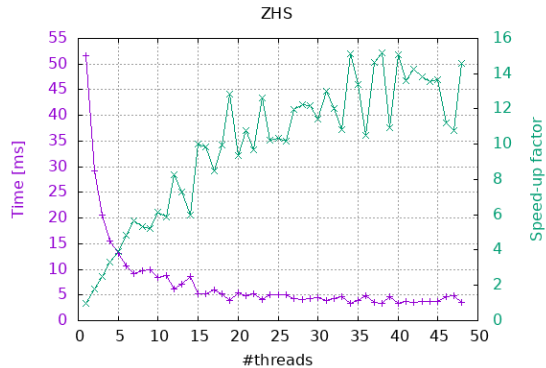
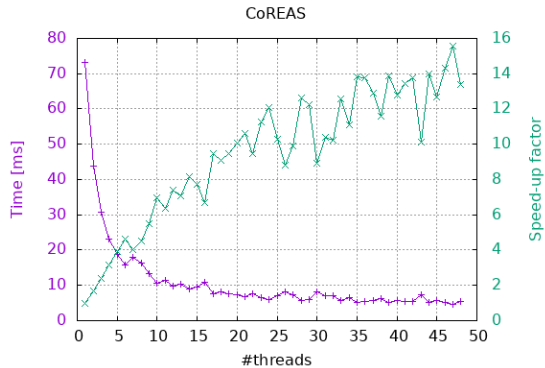


# Performance: detector with 10k antennas

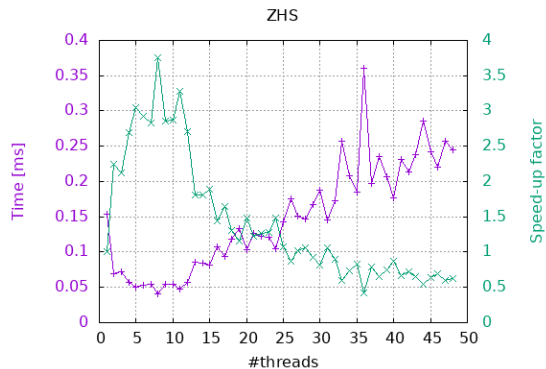





# Performance: detector with 100k antennas



# Performance: detector with 200 antennas



- Parallelized Radio module behave as expected, in terms of performance and physics (See Nikos' slides).
- Scaling behavior for small detectors probably improves more complex propagators.
- One should also pay attention to the absolute value of the timing for a given number of threads, when passing from 200 to 10,000 antennas.

## A performance issue found during this work

---

- In initial stages of this study we found out the opposite to the expected behavior: time always increasing with the number of threads.
- Investigating further we found out that the culprit was the `std::shared_ptr<>` managing the the coordinate system.
- This pointer is shared among all geometry objects. By distributing this pointer among different threads we was distributing a lock.
- We substituted it by an `corsika::dumb_ptr` and the glitch went away and we even got some performance gain in single thread application.

**Backup Slides**

# Gyges example

---

```
1 #include <future>
2 #include <iostream>
3 #include <random>
4 #include <vector>
5 #include <gyges/gang.hpp>
6
7
8 int main(int argv, char** argc)
9 {
10     //number of random numbers to accumulate per task
11     unsigned max_nr = 1000000000;
12
13     // it will create a gang with the number
14     // of cores supported by the hardware.
15     gyges::gang thread_pool{};
16
17     std::cout << "The gang has #" << thread_pool.size() << " workers\n";
18
19     //tasks will accumulate max_nr of random numbers
20     //and set the result in the corresponding position of a vector
21
22     std::vector<double> results(thread_pool.size(), 0.0);
23     std::vector<std::future<void>> monitors;
```

# Gyges example

---

```
1 for(std::size_t i=0; i< thread_pool.size() ; ++i)
2 {
3     //used to obtain a seed for the random number engine
4     std::random_device rd;
5     auto seed = rd();
6     //where to place the result
7     auto result_iterator = results.begin() + i;
8
9     //lambda function getting the necessary parameters to perform the task.
10    auto Task = [ result_iterator, max_nr, seed ](std::stop_token t) {
11
12        double partial_result = 0;
13        std::mt19937 generator( seed );
14        std::uniform_real_distribution<double> distribution(0.0, 1.0);
15
16        for( unsigned nr = 0; nr< max_nr; ++nr)
17            partial_result+=distribution(generator);
18        //set results
19        *(result_iterator) = partial_result;
20    };
21    // task submission
22    auto future = thread_pool.submit_task( Task );
23    monitors.push_back( std::move(future) );
24
25 }//close for loop
```

## Gyges example

---


```
1     //check the tasks and print the result
2     for(std::size_t i=0; i< monitors.size(); ++i ){
3         monitors[i].get();
4         std::cout << "Task #" << i << " completed. Result: "<<     results[i] << std::endl;
5     }
6
7     //stop the gang or let it get destroyed exiting scope
8     thread_pool.stop();
9
10    return 0;
11 }
```

---



# Gyges example

```
1 [augalves@LabHome Gyges_Proj] $ ./examples/use_gangs
2 The gang has #8 workers
3 Task #0 completed. Result: 4.99999e+08
4 Task #1 completed. Result: 4.99998e+08
5 Task #2 completed. Result: 4.99998e+08
6 Task #3 completed. Result: 4.99975e+08
7 Task #4 completed. Result: 4.99992e+08
8 Task #5 completed. Result: 5.00011e+08
9 Task #6 completed. Result: 5.00009e+08
10 Task #7 completed. Result: 4.99997e+08
11 [augalves@LabHome Gyges_Proj] $
```



Basically  $6 \times 10^9$  calls to RNG plus the accumulation operation performed in about 10s.

Profiling...

Thanks