



DALI

Digits Architectures Logiciels Informatique



LIRMM

Optimizing Cherenkov photons generation and propagation in Corsika

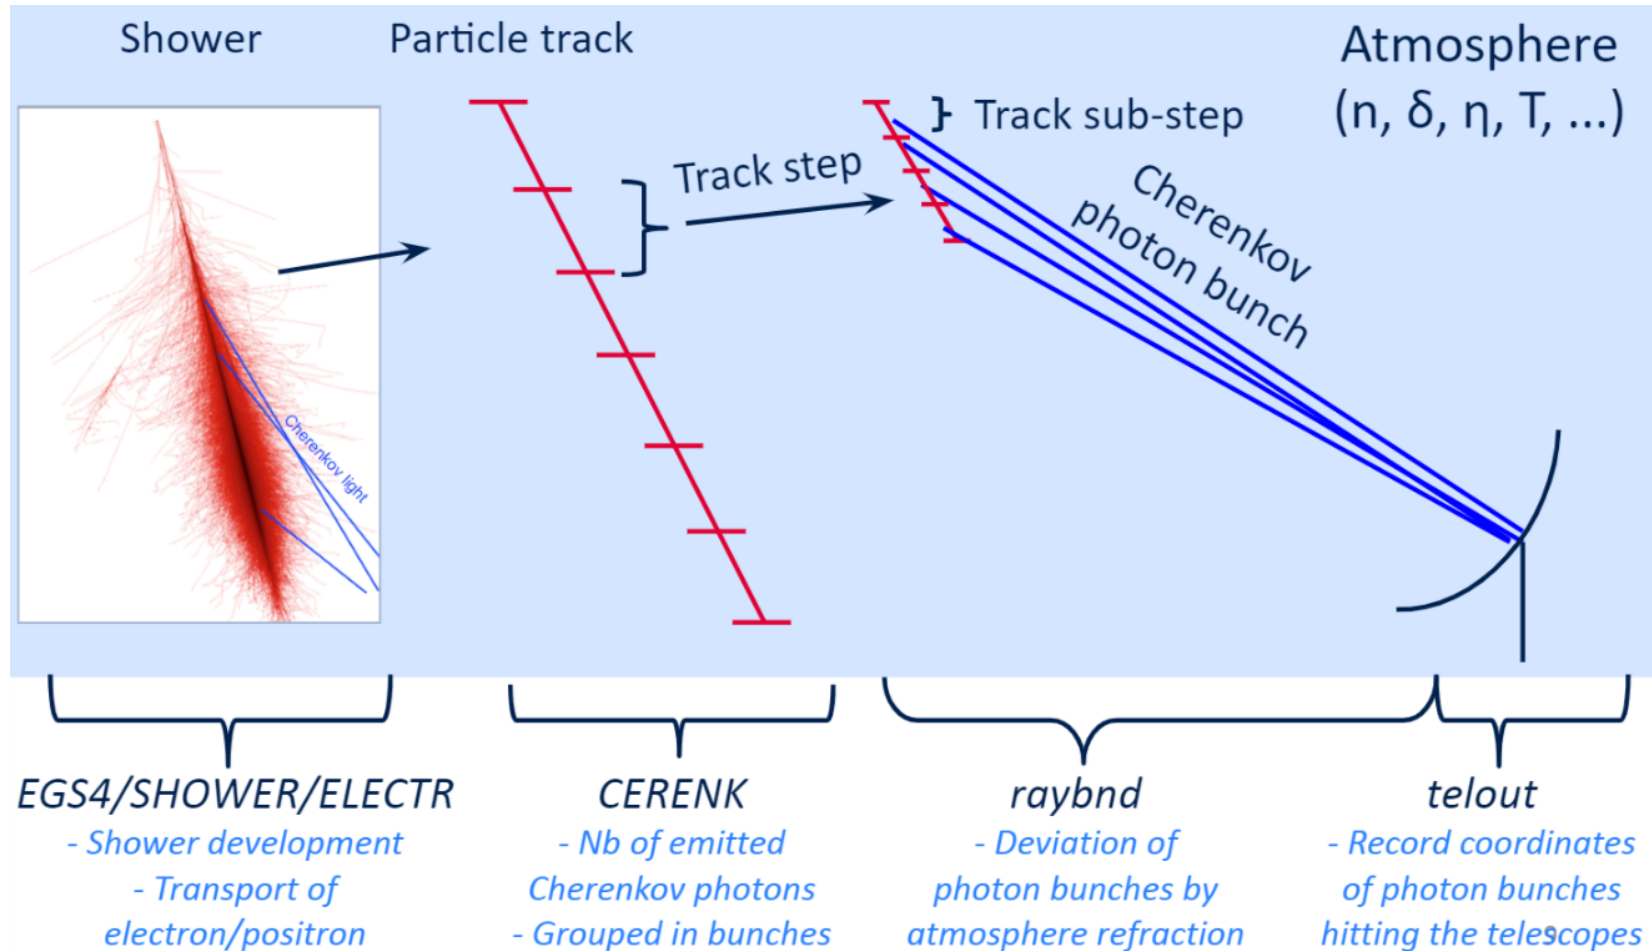
ADNANE KHATTABI



Content

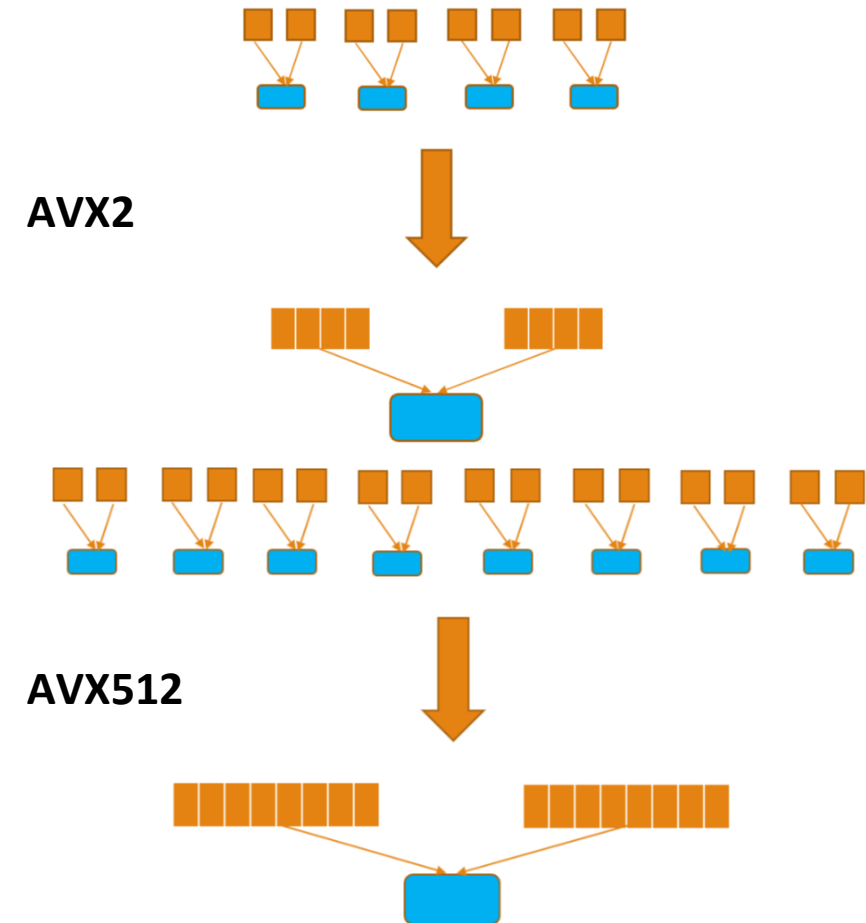
1. Cherenkov photons production and propagation.
2. Corsika profiling.
3. Optimization process.
 1. Raybnd
 2. Cerenk
 3. Cerlde
4. Experiments and performance.
5. Conclusion

Cherenkov photons production and propagation



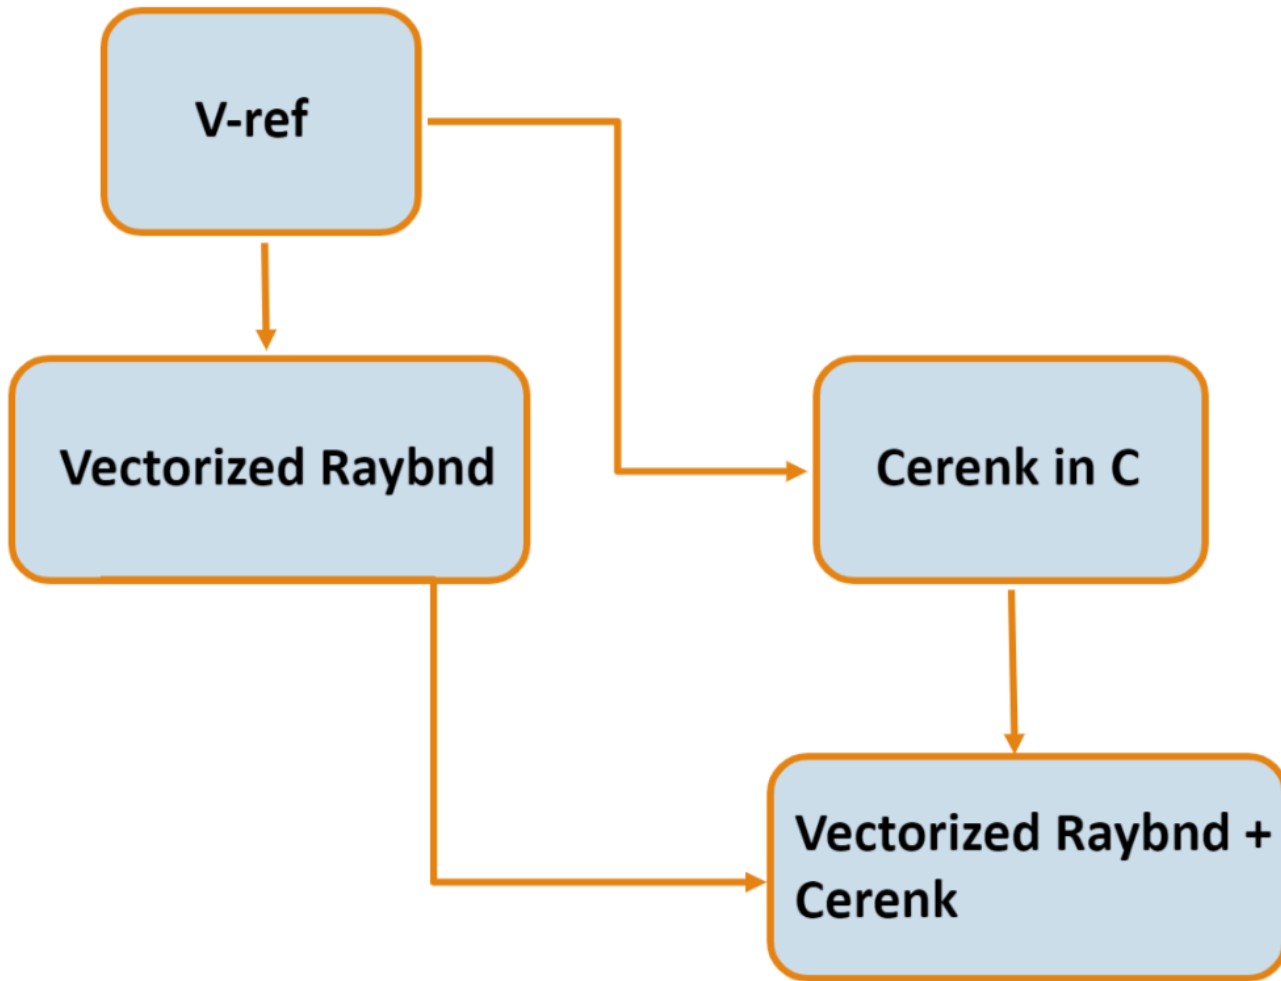
Vectorization

- Having to compute a number of scalar operations on same type variables.
- Storing a number of variables in a vector and using a vectorized operation equivalent to the scalar one.



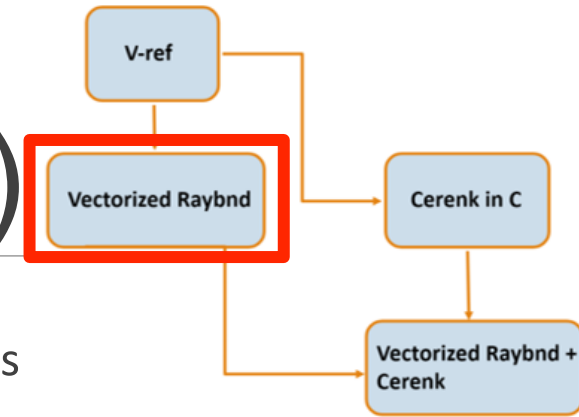
Vectorization

- Intel Intrinsics to manually vectorize the code .
- External libraries that implement vectorized functions such as math libraries.
 - SIMD vector libm
 - <https://gitlab.com/cquirin/vector-libm>
 - <https://hal.archives-ouvertes.fr/hal-01511131/document>
- Automatic vectorization of the code by the compiler.



Optimization process

Optimization process (Raybnd)



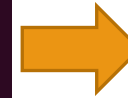
- Previous work on the vectorization Raybnd by computing multiple photon bunches once.
 - <https://gite.lirmm.fr/cta-optimization-group/cta-optimization-project/wikis/uploads/d9fe1652712cbfffe6a28c5d8c3a5da9/CorsikaOpt.pdf>
- The vectorization of mathematical functions (exp, cos ...) using a dedicated library
 - <https://gitlab.com/cquirin/vector-libm>
- Unrolling Cerenk Fortran loop to use vectorized Raybnd.
- Factorizing the code to avoid redundant function calls.

Optimization process (Raybnd)

- Advancing the optimization of Raybnd by:

- Restructuring the tests.
- Isolating computations to facilitate detection of vectorizable loops by compiler.

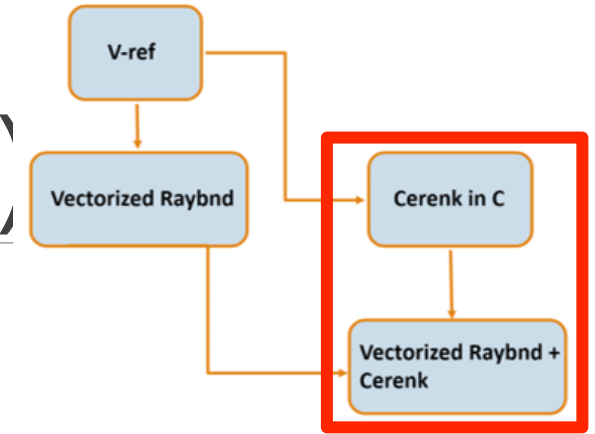
```
*u *= sin_t_obs/sin_t_em;  
*v *= sin_t_obs/sin_t_em;  
if ( (*w) >= 0. )  
    *w = sqrt(1.-sin_t_obs*sin_t_obs);  
else  
    *w = -sqrt(1.-sin_t_obs*sin_t_obs);  
  
*dx += hori_off * (*u)/sin_t_obs;  
*dy += hori_off * (*v)/sin_t_obs;  
  
*dt += travel_time;
```



```
for(int i=0; i< VECTOR_SIZE; i++){  
    u[i] *= sin_t_obs[i]/sin_t_em[i];  
}  
  
for(int i=0; i< VECTOR_SIZE; i++){  
    v[i] *= sin_t_obs[i]/sin_t_em[i];  
}  
  
for(int i=0; i< VECTOR_SIZE; i++){  
    if(w[i] >= 0)  
        w[i] = sqrt(1.-sin_t_obs[i]*sin_t_obs[i]);  
    else  
        w[i] = -sqrt(1.-sin_t_obs[i]*sin_t_obs[i]);  
}  
  
for(int i=0; i< VECTOR_SIZE; i++){  
    dx[i] += hori_off[i] * (u[i])/sin_t_obs[i];  
}  
  
for(int i=0; i< VECTOR_SIZE; i++){  
    dy[i] += hori_off[i] * (v[i])/sin_t_obs[i];  
}  
  
for(int i=0; i< VECTOR_SIZE; i++){  
    dt[i] += travel_time[i];  
}
```

- Validation of vectorization by checking assembler code.

Optimization process (Cerenk)



- Optimization opportunity in the vectorization of unrolled loop in Cerenk.
- Translating the function from Fortran to C for easier application of the transformations applied to Raybnd.
- The Translation to C requires:
 - Defining the data structures mainly the Fortran common blocks, in C.
 - Making sure conditional expressions are equivalent between the two languages.
 - Linking the C version of Cerenk to the required functions still in Fortran.

Optimization process (Cerlde)

- LONGI option is impossible to activate with vectorized versions with scalar Cerlde.

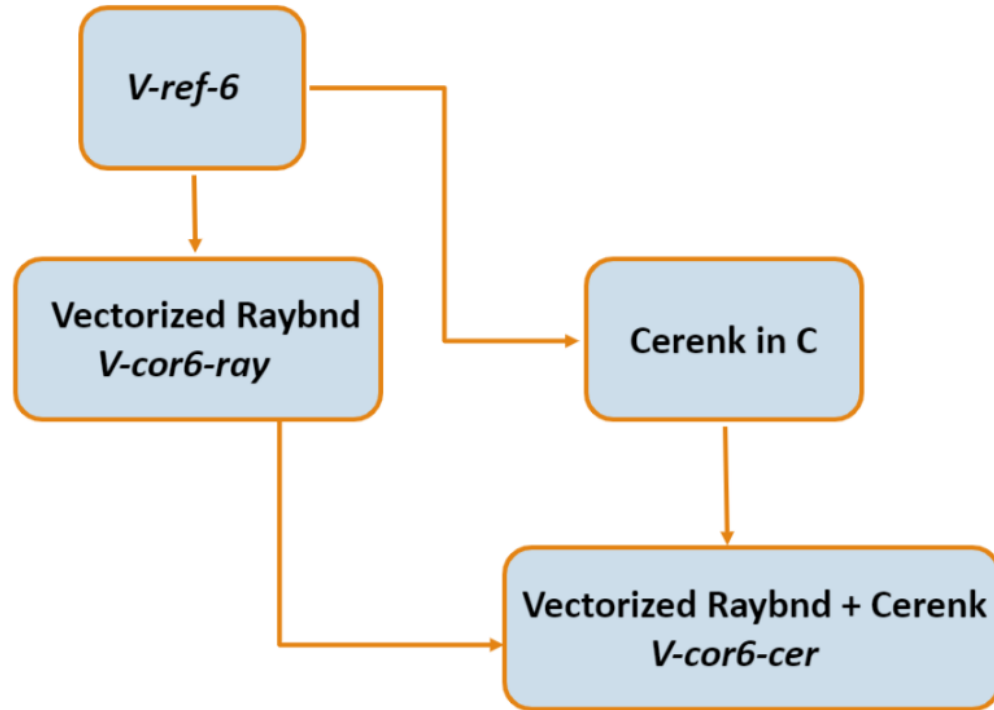


A vectorized Cerlde function is necessary.

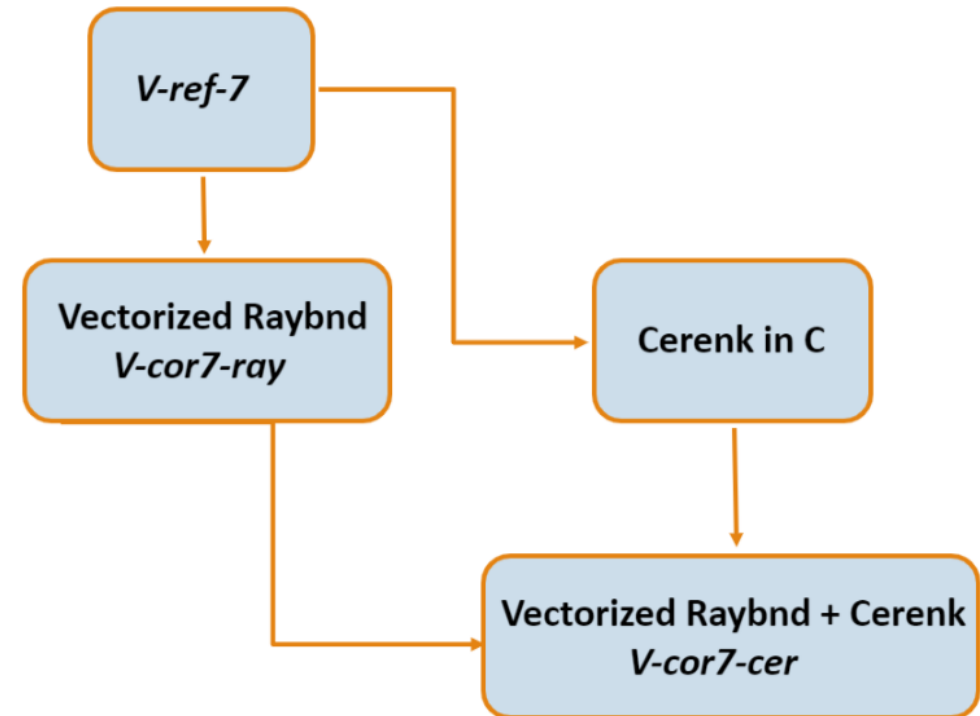
- Inlining Cerlde in the vectorized portion of CERENK while calling the sequential function in the computation of the remaining photon bunches.

Optimization process

Corsika v6 / IACT 1.51



Corsika v7 / IACT 1.59

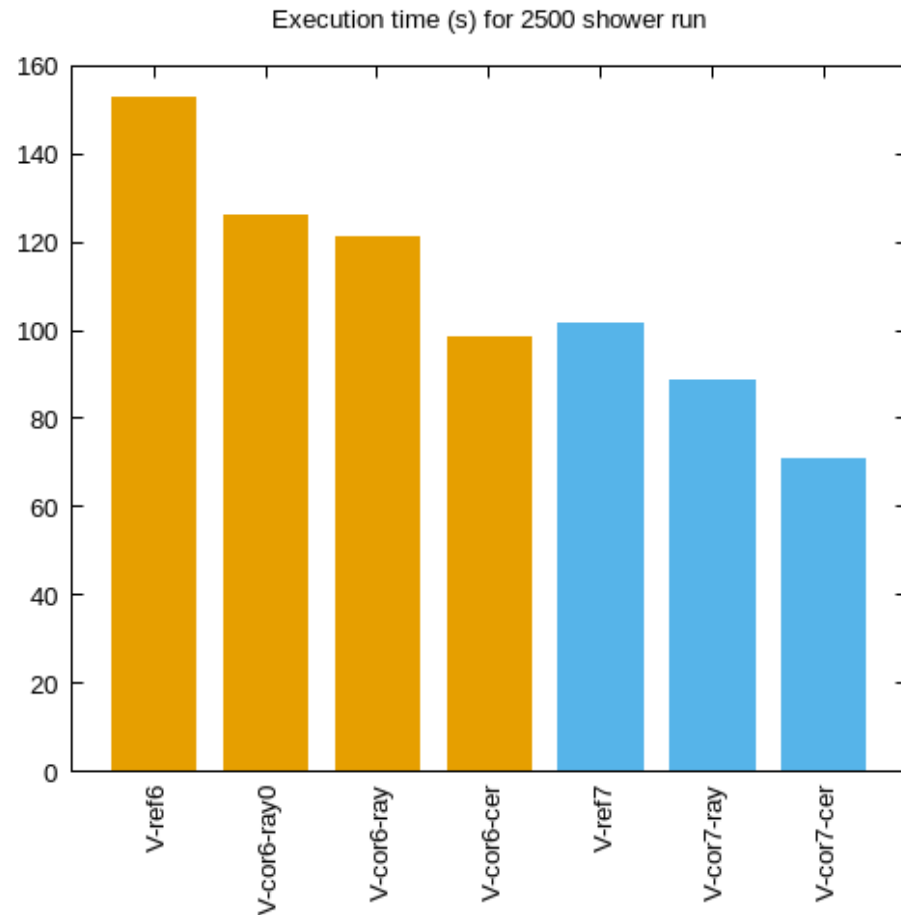


Experimental setup

- Dedicated server running running CentOS Linux release 7.4 - 64 bits.
- Compiler: **gcc 8.2.1**
- Compilation flags: **-O3 -mavx2**
- Running conditions:
 - gamma, prod3 Paranal baseline 20 deg. zenith angle (LONGI disabled)
 - Using keep-seeds option for random number generation to obtain reproducible runs
- Run validation using python script based on <https://github.com/fact-project/pyeventio>

```
Architecture:          x86_64
CPU op-mode(s):        32-bit, 64-bit
Byte Order:            Little Endian
CPU(s):                16
On-line CPU(s) list:   0-15
Thread(s) per core:    2
Core(s) per socket:    4
Socket(s):             2
NUMA node(s):         2
Vendor ID:             GenuineIntel
CPU family:            6
Model:                85
Model name:            Intel(R) Xeon(R) Gold 5122 CPU @ 3.60GHz
Stepping:              4
CPU MHz:               3600.000
BogoMIPS:              7200.00
Virtualization:        VT-x
L1d cache:             32K
L1i cache:             32K
L2 cache:              1024K
L3 cache:              16896K
```

Experiments and Performance



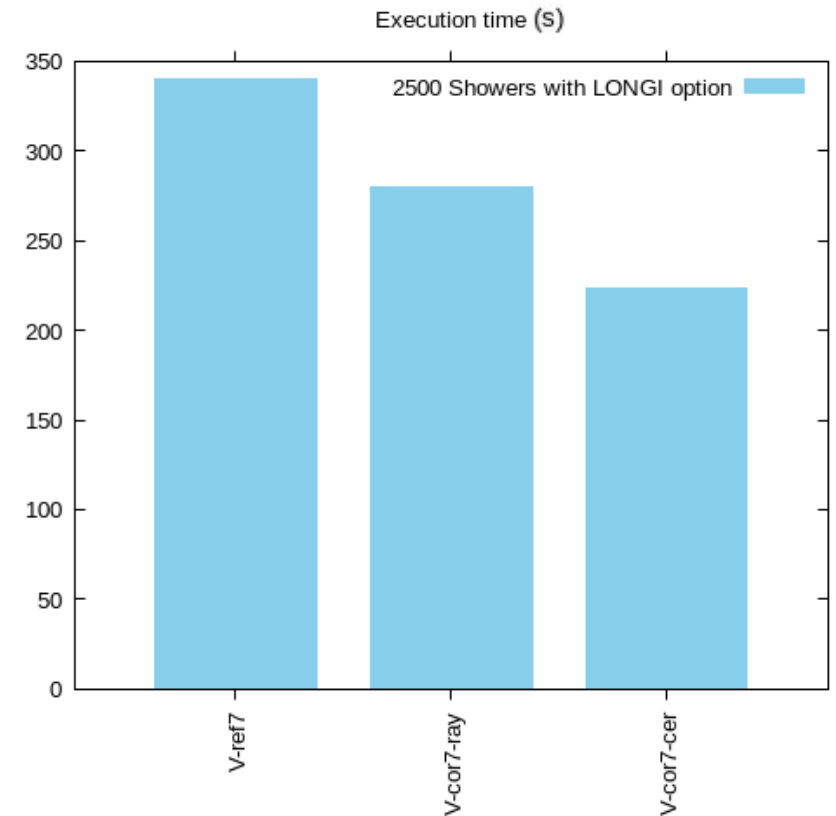
	V-cor6-ray0	V-cor6-ray	V-cor6-cer
Speed-up	1.21	1.26	1.55

	V-cor7-ray	V-cor7-cer
Speed-up	1.14	1.44

- Less CPU time spent in Raybnd in Corsika v7 : [New interpolation scheme](#).

Experiments and Performance

- The vectorization of CERLDE → A gain in performance.
 - Speedup of V-cor7-ray + LONGI option : **1.21**
 - Speedup of V-cor7-cer + LONGI option : **1.52**
- All previous experiments have vector length set to 4
 - Speedup of V-cor7-cer with -mavx512 flag (vector length = 8) : **1.60**



Conclusions

- The optimization of Corsika via vectorization:
 - Corsika v6 / IACT 1.51 a speedup of up to **1.55**
 - Corsika v7 / IACT 1.59 a speedup of up to **1.44** and **1.60** with vectors of 8 doubles.
- Corsika v7 / IACT 1.59 with LONGI option speedup : **1.52**
- Vectorized versions have identical results as the Ref versions so no accuracy is sacrificed.
- Optimization plans based on memory access and precision reduction are currently being investigated.
- The vectorized versions are in the process of being packaged for distribution.

