

Possible issues to optimize stochastic simulation time with parallel sequences and unrolling techniques



David R.C. Hill and Romain Reuillon

romain.reuillon@isima.fr

LIMOS – UMR CNRS 6158

Université Blaise Pascal

ISIMA, Campus des Cézeaux BP 10125

63177 Aubière Cedex

FRANCE



Parallelism in Monte Carlo Applications

- Computationally intensive but naturally parallel
- Appropriate for independent *bag-of-work* paradigm
- Fits the distributed computing paradigm
- Requirements :

Independence of underlying random number streams

SPRNG

(Scalable Parallel Random Number Generation) library

« *Random number generators, particularly for parallel computers, should not be trusted.* »

Paul Coddington

- It is necessary that RNG can be generated in parallel (ie each process may have an autonomous access to a sub-sequence issued from a common global sequence (for correlation problems))
- If such an autonomy is not guaranteed, the potential parallelism of the application is affected (if for instance processes access to a central RNG, even if this generator is also run in parallel)
- The main problem is to find partitioning techniques which preserve the good properties required to guarantee not only the efficiency of the simulation but mainly the credibility of the results (L'Ecuyer proposed an interesting approach in 2001)

A decorative graphic at the top of the slide consists of two groups of three circles. The first group on the left has a solid light purple circle on the left, a white circle with a light purple outline in the middle, and a solid light purple circle on the right. The second group on the right has a solid light purple circle on the left, a white circle with a light purple outline in the middle, and a solid light purple circle on the right. The word "Requirements" is written in a black serif font, centered over the first group of circles.

Requirements

- *prop. 1* : numbers are uniformly generated
- *prop. 2* : the sequence is uncorrelated
- *prop. 3* : the sequence is reproducible
- *prop. 4* : the generator is portable on any computer
- *prop. 5* : the sequence can be changed by adjusting a seed or a status
- *prop. 6* : the period is as large as possible
- *prop. 7* : the generator satisfy any randomness test
- *prop. 8* : a quick generation is obtained
- *prop. 9* : the generator uses a limited computer memory



Requirements for Parallel RNGs


- *prop. 10* : it is easy to split the numbers into many independent sub-sequences that are allocated to different LPs, without the need of communication or synchronization; each sub-sequence is a good sequential RNG
- *prop. 11* : there is no correlation between the sub-sequences on different LPs (crossed-correlation)



Parallelizing RNGs : partitioning methods


- The LF (Leap Frog) technique:
partitioning a sequence $\{x_i, i=0, 1, \dots\}$ into 'n' sub-sequences, the j^{th} sub-sequence is $\{x_{kn+j-1}, k=0, 1, \dots\}$
- The SS (Sequence Splitting or Cycle Splitting) technique:
 - Partitioning a sequence $\{x_i, i=0, 1, \dots\}$ into 'n' sub-sequences
The j^{th} sub-sequence is $\{x_{k+(j-1)m}, k=0, \dots, m-1\}$
where m is the length of each sub-sequence
 - The user deterministically chooses widely separated seeds in same generator.
- The IS (Independent Sequences or Cycle Splitting) technique:
 - Using randomly generated seeds
 - Using different parameters for the same kind of generator (example Matsumoto and Nishimura MT) => need theoretical studies of mathematical properties at first.

Leap Frog Technique



"Long-Range" correlation problem:
auto-correlation

Sequence splitting technique



"Long-Range" correlation problem :
Crossed-correlation

Independent sequences technique



x_0 x_1 x_2 x_3 x_4 x_5 x_6 x_7 x_8 x_9 x_{10} x_{11} x_{12} x_{13} x_{14} x_{15}



x'_0 x'_1 x'_2 x'_3 x'_4 x'_5 x'_6 x'_7 x'_8 x'_9 x'_{10} x'_{11} x'_{12} x'_{13} x'_{14} x'_{15}





y_0 y_1 y_2 y_3 y_4 y_5 y_6 y_7 y_8 y_9 y_{10} y_{11} y_{12} y_{13} y_{14} y_{15}



y'_0 y'_1 y'_2 y'_3 y'_4 y'_5 y'_6 y'_7 y'_8 y'_9 y'_{10} y'_{11} y'_{12} y'_{13} y'_{14} y'_{15}

Hybrid technique



Hybridizing LF and SS:

partitioning a sequence into 'n' sub-sequences,

the jth sub-sequence is $\{x(kn+j-1)m+s, k=0, \dots$ and $s=0, \dots, m-1\}$

where m is the length of contiguous sub-blocks.

Reducing "Long-Range" correlations




Parallelizing PRNGs : tests

- Theoretical tests:
the random sequences should have the properties of a random sample drawn from the uniform distribution.
(Series of Knuth, DieHard tests, L'Ecuyer TestU01)
- Empirical tests:
As stated by Paul Coddington in 1996 : *“It is strongly recommended that all simulations be done with two or more different generators, and the result compared to check whether the random number generator is introducing a bias ”*

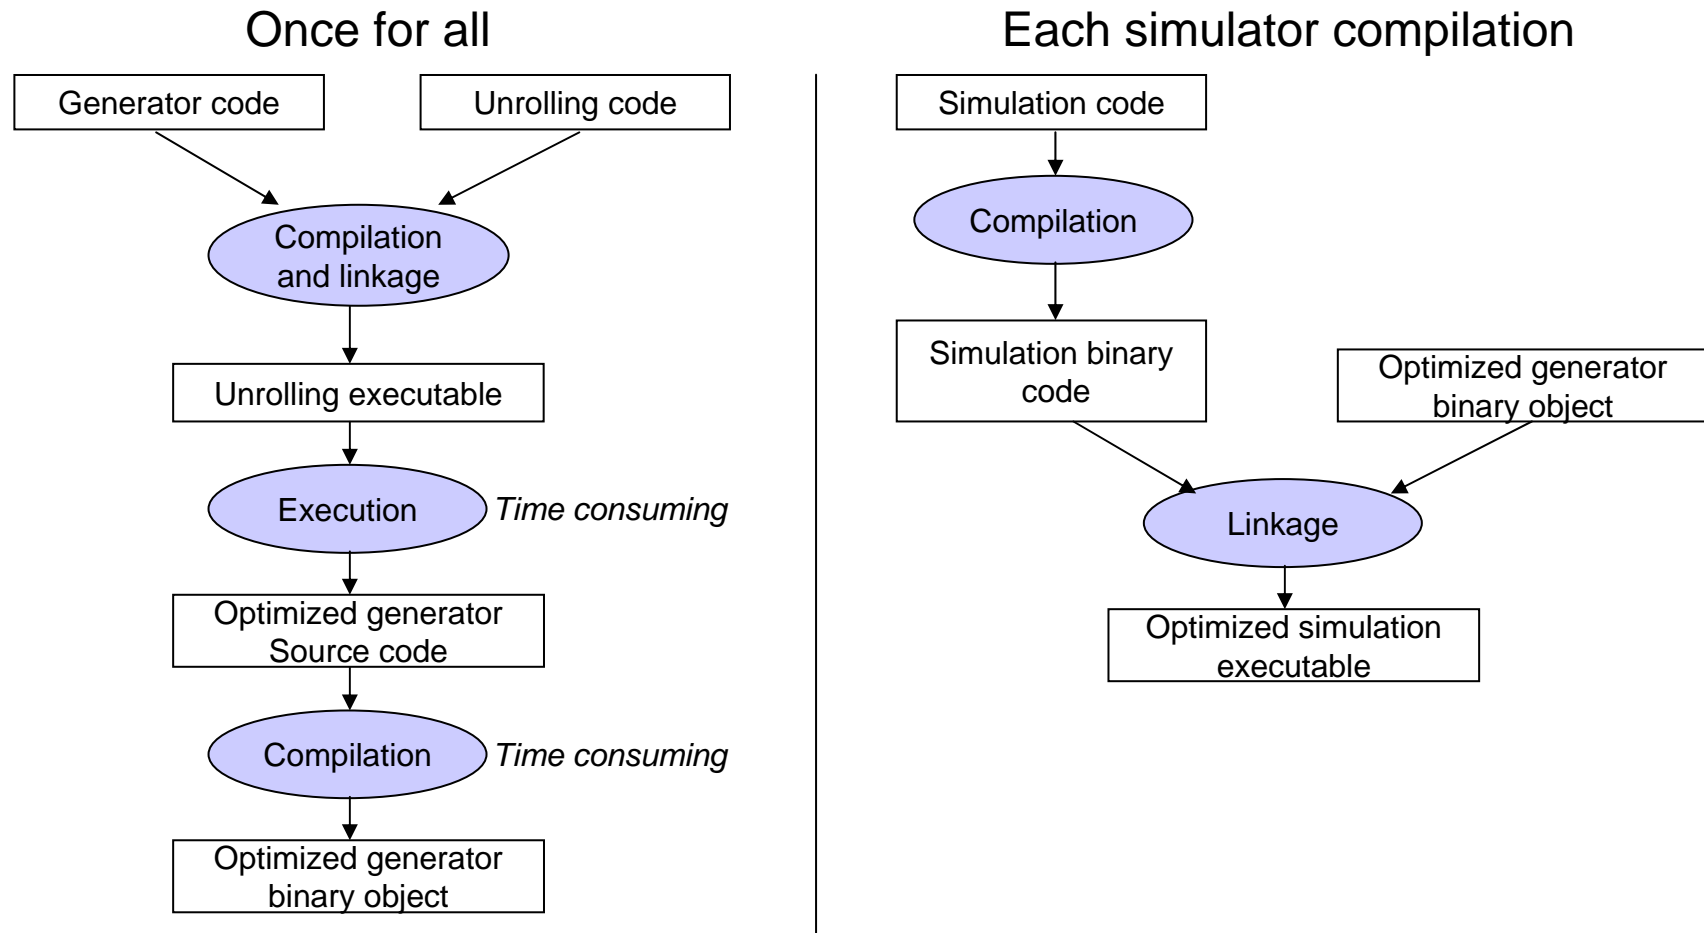
The classical approach to test PRNG is to test separately each sub-sequence, and then to test the entire sequence generated

Random number generation in Geant4 : up to 20% of running time

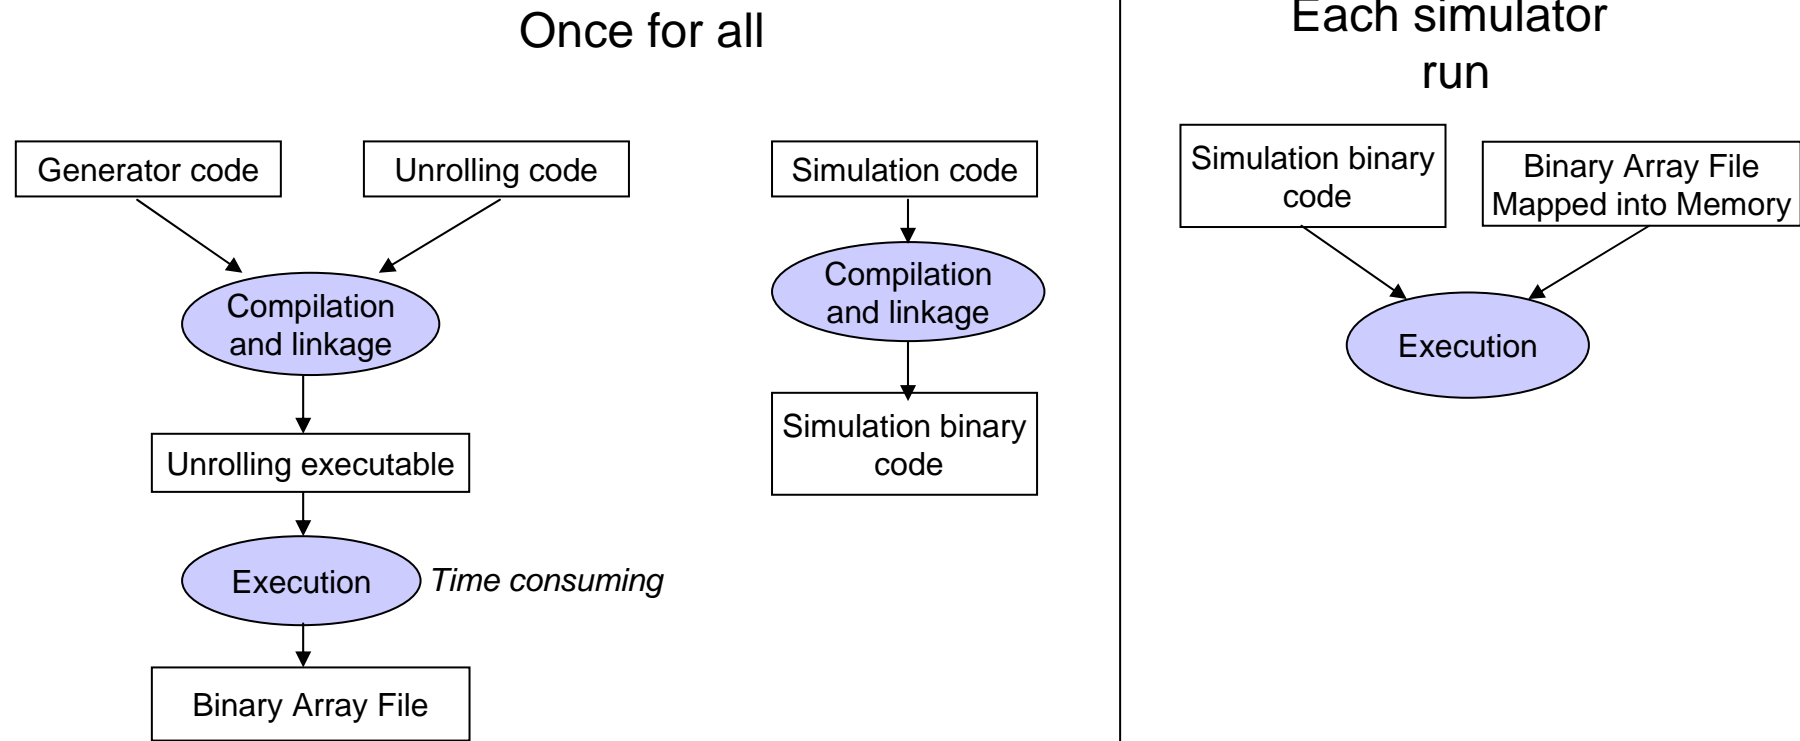
Occupation CPU de CLHEP Pour le Noyau



A speedup technique faster than getting RN from hardware cards : metaprogramming by « unrolling » : URNG

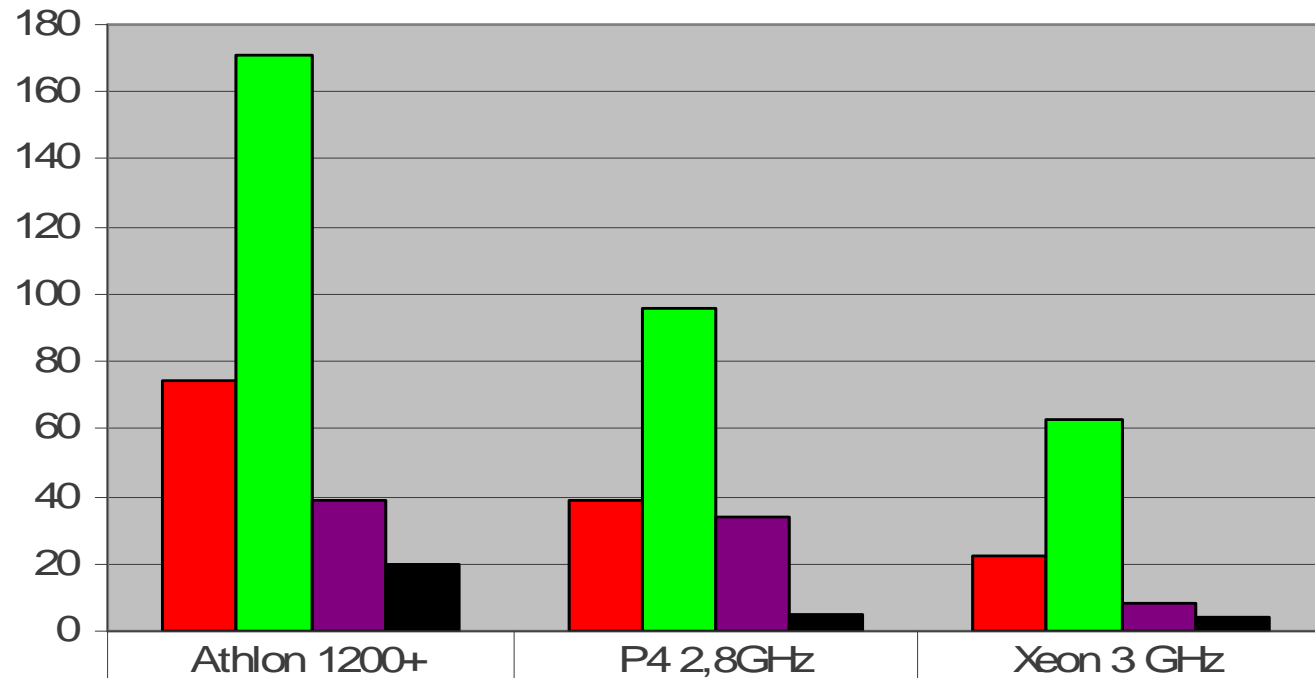


Metaprogramming useless for files larger than RAM size => Unrolling with « Memory Mapping »



Comparative tests under Linux on 3 common generators and the unrolled version

Mean Generation Time (in ms) for 1 million drawings



■ Sobol GSL	74	39	22
■ Shuffled NR	171	96	63
■ Quick and Dirty	39	34	8
■ Unrolled	20	5	4



Conclusion

- Acceleration of Monte Carlo methods
 - Variance Reduction
 - Quasi Monte Carlo
 - Optimising number generation
 - Parallel Computing
- The best option to reduce Geant4 computing time is parallel computing
 - Be careful with random number generation !!!!