

REPRODUCING PARALLEL STOCHASTIC SIMULATIONS

Enabling parallel and sequential results comparison
before scaling on top supercomputers

HILL David – ISIMA/LIMOS UMR CNRS 6158

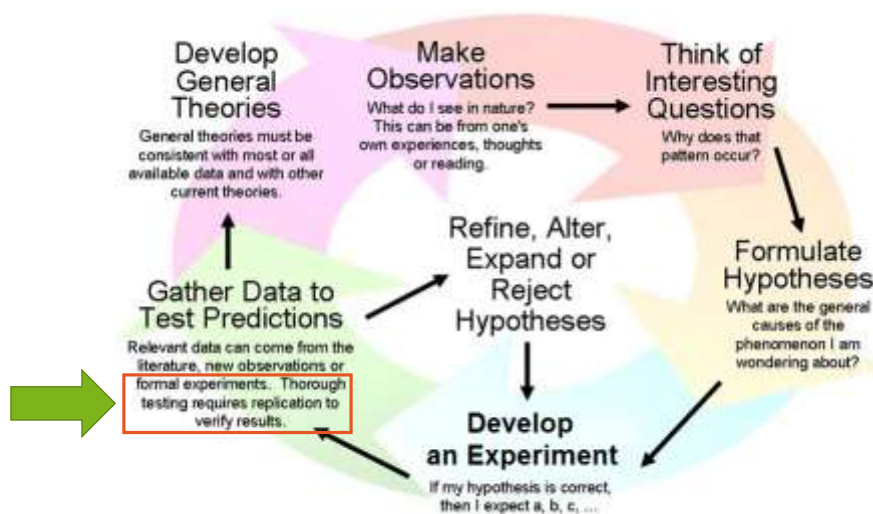


1ST LET'S GO BACK TO PHILOSOPHY OF SCIENCE

Just because we (computer scientists) give our own definitions,
which adds confusion - It's getting cold... isn't it?

THE SCIENTIFIC METHOD

https://i0.wp.com/peegel.info/wp-content/uploads/2017/10/scientific_method.png



Traditionally we have 2 main branches of the scientific method:

- 1 – Deductive branch
Mathematics and formal logic
- 2 – Empirical branch
Statistical analysis of controlled experiments

Hope for a 3rd & 4th branches

- 3 – Large Scale Simulation
- 4 – Data intensive & data driven computer Science

But we do not meet the standards of Branch 1 & 2...

CRITERIA OF THE SCIENTIFIC METHOD

The major criteria of the scientific method are **intangible principles** :

- **Refutability** : a scientific affirmation is said to be rebuttable if it is possible to record an observation or conduct an experiment which, if it were positive, would contradict this statement.
- **Non-contradiction** : is the law that prohibits affirming and denying the same term or proposition
- **Reproducibility** : Science works by drawing "laws" or "principles" from reproducible observations whose main property is to be true as long as no observation has proved otherwise.

REPRODUCIBILITY...

- Many of us know the important work of **Karl Popper** (philosopher of sciences) in modeling and simulation. Karl Popper is generally regarded as one of the greatest philosophers of science of the 20th century.
- The criterion of **reproducibility** is one of the conditions on which Popper distinguishes between the scientific or **pseudo-scientific** character of a study.
- Scientific conclusions **can only be drawn from a well observed and described “event”, which has appeared several times**, observed by different people and/or studies.
- This criterion **eliminates random effects** that distort the results as well as **errors** in judgment or **manipulations** by scientists.

YOU KNOW THAT MANY DOMAINS ARE IMPACTED 'REPRODUCIBILITY CRISIS...'



nature International weekly journal of science

Home | News & Comment | Research | Careers & Jobs | Current Issue | Archive |
 Archive > Specials and supplements archive > Challenges in irreproducible research


SPECIAL

CHALLENGES IN IRREPRODUCIBLE RESEARCH

Science moves forward by collaboration – when researchers verify others' results. Science advances faster when people waste less time pursuing false leads. No research paper can ever be considered to be the final word, but there are too many that do not stand up to further study.

There is growing alarm about results that cannot be reproduced. Explanations include increased levels of scrutiny, complexity of experiments and statistics, and pressures on researchers. Journals, scientists, institutions and funders all have a part in tackling reproducibility. Nature has taken substantive steps to improve the transparency and robustness in what we publish, and to promote awareness within the scientific community. We hope that the articles contained in this collection will help.

Recent articles | Editorial | Features | News and analysis | Comment |
 Perspectives and reviews




WHAT ABOUT COMPUTER SCIENCE ? (PSEUDO-SCIENCE TOO?)

Remember the recent repeatability issues found in a 2015 survey :

A study of 601
'A' ranked papers
(journal or conference papers)
published by the distinguished ACM
(Association for Computing Machinery)
showed that a only bit more than 30% of the results could be reproduced

<http://reproducibility.cs.arizona.edu/>



REPRODUCIBILITY & CORROBORATION

- There is a growing alarm of results that have been published but that cannot be reproduced.
- Science advances faster when people waste less time pursuing false leads.
- Science moves forward by **corroboration** when researchers verify each other's data.
- A study of **top scientific research** in UK (REF) showed that only 11% of medical studies where reproducible. (First page of "The Guardian")



REPRODUCIBLE COMPUTER
SCIENCE IS GOOD

BUT REPEATED COMPUTER
SCIENCE CAN BE BETTER

LET'S HAVE A LOOK AT SOME DEFINITIONS... WHAT DO WE SEE FROM COLLEAGUES

- In Fomel and Claerbout 2009:

- ✓ Reproducibility **often means replication** depending on computer scientists

- In Drummond 2009¹: 1: <http://www.site.uottawa.ca/ICML09WS/papers/w2.pdf>


- ✓ “Reproducibility **requires changes; replicability avoids them**”

- In Demmel and Nguyen 2013 (COMPUTER ORIENTED – means replication)

- ✓ “Reproducibility, i.e. **getting bitwise identical results from run to run**”

- In Revol and Théveny 2013 (COMPUTER ORIENTED TOO – means replication too)

- ✓ “What is called **numerical reproducibility** is the problem of getting the same result when the scientific computation is run several times, either on the same machine or on different machines, with different numbers of processing units, types, execution environments, computational loads, etc.”



THIS IS THE
MOST
COMMON
SCIENTIFIC
SENSE

WHY DO WE NEED REPEATABILITY ?



- If you don't have **repeatability**, **how would we debug our stochastic simulations ?**
How do we repeat/reproduce the events observed in simulations ?
(confirmation of Higgs discovery, etc...)
- In Digital Computer Science we are used to **deterministic computing** and we expect « **repeatability** » of computer experiments. **Computer debugging and program setup is based in repeatability!**
- **Even when we use pseudo-random numbers** for stochastic models, **we are running deterministic experiments since pseudo-random number generators have been carefully designed to be repeatable** (though some computer scientist often use the "reproducible" term...).
- In the context of a Biological or Physical experiment, **repeatability** measures the variation in measurements taken by a single instrument or person under the same conditions, while **reproducibility** measures whether an entire study or experiment can be reproduced in its entirety – by the same research theme or by another team.

Floating point...

- Round off errors
- Order of floating point operations (dynamic execution / out of order)
- ...

Hardware

- Number of processors, Networking Interconnect, devices and latency
- Difference between architectures (regular processors, vs accelerators,...) – Hybrid computing.
- Processor implementation or design bugs
- Soft errors
- ...

Software

- Operating systems, compilers,
- Libraries, dependencies and software stack versions
- Parallelization techniques
- Virtual machines and containers (rare in HPC > bare metal)
- ...

IN ADDITION
TO POSSIBLE
INDIVIDUAL ERRORS
AND MISCONDUCTS...

HERE ARE SOME
TECHNICAL REASONS
FOR **NUMERICAL
REPEATABILITY
FAILURES**

See Intel – 2014 https://www.mpcdf.mpg.de/services/computing/software/languages-1/FP_accuracy_reproducibility.pdf

- **Out-of-order execution** is also known as **dynamic execution**. Most **modern high-performance microprocessors** optimize the execution of instructions based on the availability of input data to avoid delays. By default C & C++ are not impacted, unfortunately Fortran is impacted...


- **Original order of instructions – not always respected !**
 - The micro-processor avoids having parts of its internal computing units being idle by processing the next instructions which are able to run immediately and “independently”.
 - It is the equivalent of the software dynamic recompilation which enables improving instruction scheduling.
 - It may impact floating point operations
- floating point arithmetic is not associative** (for + & *)
- ex: $a + (b + c) \neq (a + b) + c$



Numerically non-expert programmers are legion

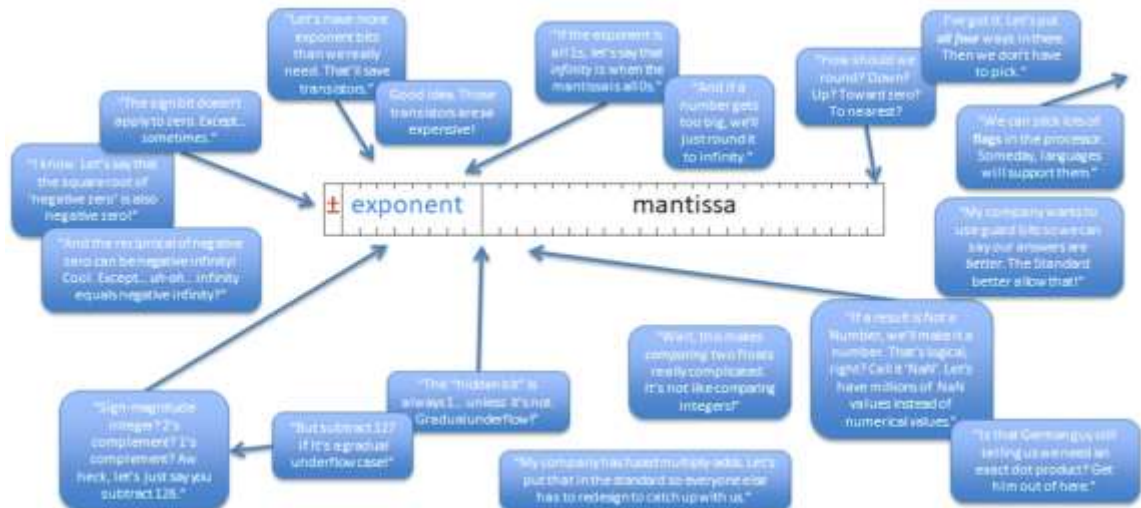
Error-analysis can be very unobvious

Competent error-analysts are extremely rare...



FLOATS WERE DESIGNED BY A COMMITTEE IN THE 1980'S SEE GUSTAFSON'S CONFERENCE – BEATING FLOAT'S AT THEIR OWN GAME

<http://www.johngustafson.net/pdfs/BeatingFloatingPoint.pdf> (Posits, Unums...)



EXAMPLE OF MICROPROCESSOR DESIGN ERRORS AND MISS-BEHAVIORS > HYPER-THREADING, MELTDOWN, SPECTRE,...

[WARNING] Intel Skylake/Kaby Lake processors: broken hyper ...

<https://lists.debian.org/debian-devel/2017/06/msg00308.html> - Traduire cette page

25 juin 2017 - TL;DR: unfixed Skylake and Kaby Lake processors could, in some situations, dangerously misbehave when hyper-threading is enabled. Disable hyper-threading immediately in BIOS/UEFI to work around the problem. Read this advisory for instructions about an Intel-provided fix. SO, WHAT IS THIS ALL ...

Users of systems with Intel Skylake processors may have two choices:

1. If your processor model (listed in `/proc/cpuinfo`) is 78 or 94, and the stepping is 3, install the non-free "intel-microcode" package with base version 3.20170511.1, and reboot the system. THIS IS THE RECOMMENDED SOLUTION FOR THESE SYSTEMS, AS IT FIXES OTHER PROCESSOR ISSUES AS WELL.

Skylake and Kaby Lake CPUs have broken hyper-threading - Fudzilla

<https://www.fudzilla.com/.../43964-skylake-and-kaby-lake-cpus-ha...> - Traduire cette page


26 juin 2017 - During April and May, Intel started updating processor documentation with a new errata note and it turned out that the reason was that Skylake and Kaby Lake silicon has a microcode bug it did not want any one to find out about. The errata is described in detail on the Debian mailing list, and affects Skylake ...



RELIABILITY & HPC...SILENT & SOFT ERRORS...


1. Change the system state by 'external forces'

- Alpha particles
- Cosmic rays (High Energy Particles from space)
- Thermal neutrons
- Variation in voltage, temperature, etc.



2. They are at the origin of ECC...to avoid bits flips in memory cells

- There is also a **rising of soft errors in arithmetic units !!!**
- **The more we size down the more this problem increases.**
- **Chip manufacturers spend money and silicon space to avoid this kind of errors:**
 - Samsung, GlobalFoundries, and IBM introduced the world's first **5nm** chip with GAAFET transistors, GAA (gate-all-around) stacked nano-sheet transistors.




3. Soft errors are difficult to detect **and reproduce** Using spare time of Supercomputers to check ?

17

RUN TO RUN REPEATABILITY ERRORS

SEE THE WORK OF FRANCOIS THOMAS – OPTIMIZATION OF WEATHER APPLICATIONS
ON POWER AND X86 ARCHITECTURES (TOULOUSE CERFACS) – **BITWISE REPEATABILITY** IS THE TARGET

My post-Mayascale prediction : « The weather forecast for Dec, 21st, 2018 could be different from the weather forecast for Dec, 21st, 2018, itself different from the weather forecast for Dec, 21st, 2018. »





WE DON'T HAVE EASY SOLUTIONS – BUT TOOLS ARE COMING...

Evolutions of containers like Singularity for HPC - efficient binary containers (ready for ARM processors...)

Embedded publishing :

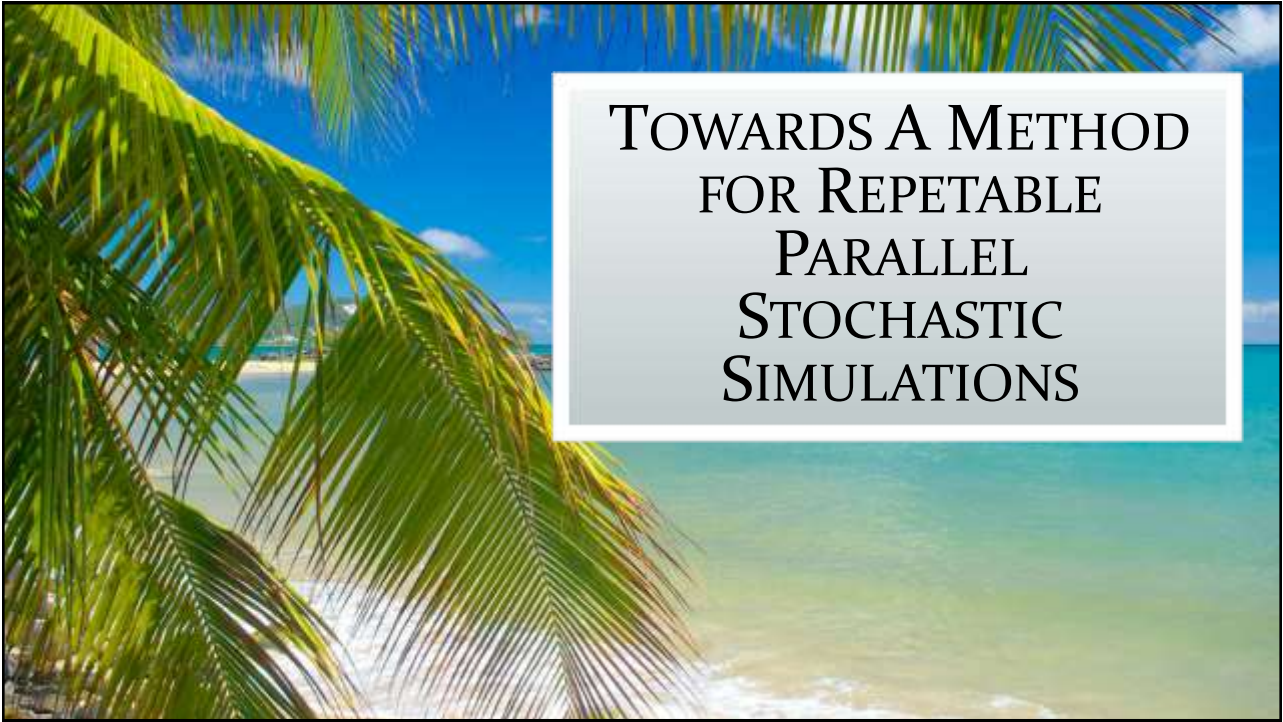
Sweave, knitR, ReScience, SHARE, Verifiable Computational Research, SOLE, Collage Authoring Environment.

Dissemination Platforms:

IPOL, ResearchCompendia.org, Madagascar, MLOSS.org, CoRR (NIST), RunMyCode.org, nanoHUB.org, thedatahub.org, Open Science Framework, Scientific Open Data,...

Workflow Tracking & Resarch Environment :

Sumatra, CoRR (NIST), CDE, Kepler, Chameleon, Galaxy, Tavera, Pegasus, Jupyter notebook, GenePattern,



TOWARDS A METHOD FOR REPETABLE PARALLEL STOCHASTIC SIMULATIONS

Parallel Stochastic Simulations... Various requirements...

Most Parallel Monte Carlo Simulations are often easy to parallelize.

- Particularly when they fit with the **independent bag-of-work** paradigm.
- Such stochastic simulations can easily tolerate a loss of jobs, if hopefully enough jobs finish for the final statistics...
- **Requirements:**
 - Fine Generator, Fine Parallelization technique and “independent” Parallel random streams.
 - Random statuses should be small and fast to checkpoint at Exascale (Original MT – 6Kb status – MRG32K3a 6 integers)
- Should fit with **different distributed computing platforms / HPC nodes**
 - Using regular processors
 - Using hardware accelerators : GP-GPUs, Intel IGP/GPU Xe, Phi, (and FPGAs ?)

21

EVEN IF WE HAVE NO DEPENDENCY BETWEEN ELEMENTARY COMPUTING REPEATABILITY OF PARALLEL SIMULATION IS NOT GRANTED

A system being of collection of interacting “objects” (dictionary definition)
– a simulation will make all those objects evolve during the simulation time with a precise modeling goal.

- **To obtain repeatability – think parallel when you design your sequential code** : Assign an « independent » random stream and initialization status for the pseudo-random number generator of each stochastic object of the simulation.
- An object could also encapsulate a random variate used at some points of the simulation. Every random variate could also have their own random stream with the same approach.
- **This O.O. approach, applied to stochastic objects, is the key to have a reference sequential program that we will be able to compare to a parallel version.**

[Hill 1996] : HILL D., “Object-oriented Analysis and Simulation”, Addison-Wesley, 1996, 291 p.

1. Check with some top PRNGs used with different execution context (hardware,

-

Errors found: Different Compilers (2 cases) With Identical Hardware (2 cases) Operating

[illegible]


Table 5: Results for TaryNIT, via PRNG via Core 17-46MHz using Windows 7 with NISTN

Expected results CHU644OUT.TXT	Results obtained with MnGW.gpy
1.1281366994756	1.1281366994756
1.36520337675698	1.36520337675698
1.21819878626463	1.21819878626464

Table 7: Results for TinyMT_32 PRNG with open64-i386 on virtual machines of Ubuntu-13.04 and 14.04

Expected results CHECK32 OUT. EXT	Results obtained with Ubuntu 13 on Virtual Box	Results obtained of Ubuntu 14 on Virtual Box
0.6455914	0.6455913	0.6455913
0.9415597	0.9415598	0.9415598
0.9034473	0.9034472	0.9034474
0.9348063	0.9348064	0.9348064
0.7581965	0.7581964	0.7581964

REMEMBER THE POTENTIAL IMPACT OF THE GENERATOR QUALITY FOR SENSITIVE APPLICATIONS...



Example of two results of a local TEST simulation – PDE Harmonic solution computed with Brownian movements.

- On the left the image is obtained with the current Linux rand (which is already far better than the old std UNIX rand on 15bits).
- On the right – same simulation with the 2002 version of Matsumoto Mersenne Twister
Then, the right solution is obtained : ellipsoid with a circular section.

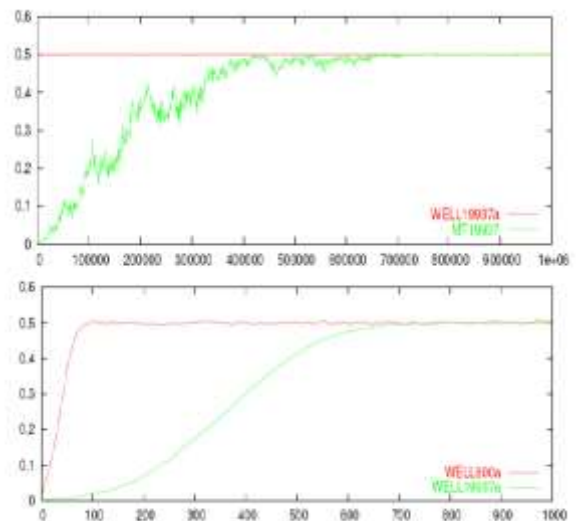
THERE IS NO PERFECT GENERATOR - THE INITIALIZATION MATTERS TOO... EX: FIRST MERSENNE TWISTER : A KNOWN DEFAULT...

Ex: very long recovery of **zero-excess initial state** for MT19237 before 2002.
(More than 700 000 drawing to recover).

Figure on the right is an extract from:
(Panneton et al., 2006)

Panneton F., L'Ecuyer P. and Matsumoto M. Improved Long-Period Generators Based on Linear Recurrences Modulo 2 [Article] // ACM Transactions on Mathematical Software. - 2006. - 1 : Vol. 32. - pp. 1-16.

Number of prints to find proportions of 1 and 0 equilibrate after an initialization with a large majority of 0
Tested with WELL and Mersenne Twister 19937
(ordinate: proportion of 1 and on the x-axis the number of prints made).



SOME TOP PRNGS (PSEUDO RANDOM NUMBER GENERATORS)

Green PRNG are said 'crush' resistant (TestU01 software) and can be recommended:

- **MRG** (Multiple Recursive Generator) – **slow but top API for reproducing parallel simulations**

$$x_i = (a_1 * x_{i-1} + a_2 * x_{i-2} + \dots + a_k * x_{i-k} + c) \bmod m - \text{with } k > 1$$

Ex: **MRG32k3a** & **MRG32kp** – by L'Ecuyer and Panneton
- **MLFG** (Multiple Lagged Fibonacci Generator) – Non linear
 by Michael Mascagni MLFG 6331_64
- **Mersenne Twisters** – by Matsumoto, Nishimura, Saito (**MT**, **SFMT**, **MTGP**, **TinyMT**...)
- **WELLS generators** by – Panneton, L'Ecuyer and Matsumoto, L'Ecuyer
- **Phylox and Threefry** – by Salmon et al. presented at SC'11 with crypto background and a parameterization technique to distribute different. In his master's thesis, Liang Li (Prof. Mascagni's student couldn't reproduce these tests. I had the same problem with Philox4x32-10.

See the following reference for advices including hardware accelerators.

HILL D. PASSERAT-PALMBACH J. MAZEL C., TRAORE, M.K., "Distribution of Random Streams for Simulation Practitioners", Concurrency and Computation: Practice and Experience, June 2013, Vol. 25, Issue 10, pp. 1427-1442. ²⁷

QUICK SURVEY OF RANDOM STREAMS PARALLELIZATION

(1) USING THE SAME GENERATOR

- The **Central Server** (CS) technique (avoid for HPC reproducibility)
- The **Leap Frog** (LF) technique. Means partitioning a sequence $\{x_i, i=0, 1, \dots\}$ into 'n' sub-sequences, the j^{th} sub-sequence is $\{x_{kn+j-1}, k=0, 1, \dots\}$ - like a deck of cards dealt to card players.
- The **Sequence Splitting** (SS) – or blocking or **regular/fixed spacing** technique. Means 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 **Cycle Division** or **Jump ahead** approach. Analytical computing of the generator state in advance after a huge number of cycles (generations). Jump Ahead technique (can be used for both Leap Frog or Sequence splitting)
- The **Indexed Sequences** (IS) - or random spacing. Means that the generator is initialized with 'n' different seeds/statuses

QUICK SURVEY OF RANDOM STREAMS PARALLELIZATION

(2) USING DIFFERENT GENERATORS:

Parameterization:

The same type of generator is used with different parameters for each processor meaning that we produce different generators

- A paper describes an implementation of parallel random number sequences by varying a set of different parameters instead of splitting a single random sequence (Chi and Cao 2010).
- For Mersenne Twister a **dynamic creation technique** is available.
- **Phylox & Threfry (2011)** are counter based generator, thus they **propose parametric statuses formed only by a single key** that can be set at runtime according to a unique identifier for each thread.
 - They can be used on CPU or GPU with good performance in terms of throughput and memory footprint (a GPU version of these PRNGs is supplied directly by their authors, either for CUDA or for OpenCL).

29

A METHOD FOR REPEATABLE PARALLEL STOCHASTIC SIMULATIONS

Remember that a stochastic program is « deterministic » if we use (initialize and parallelize) correctly the pseudo-random number.

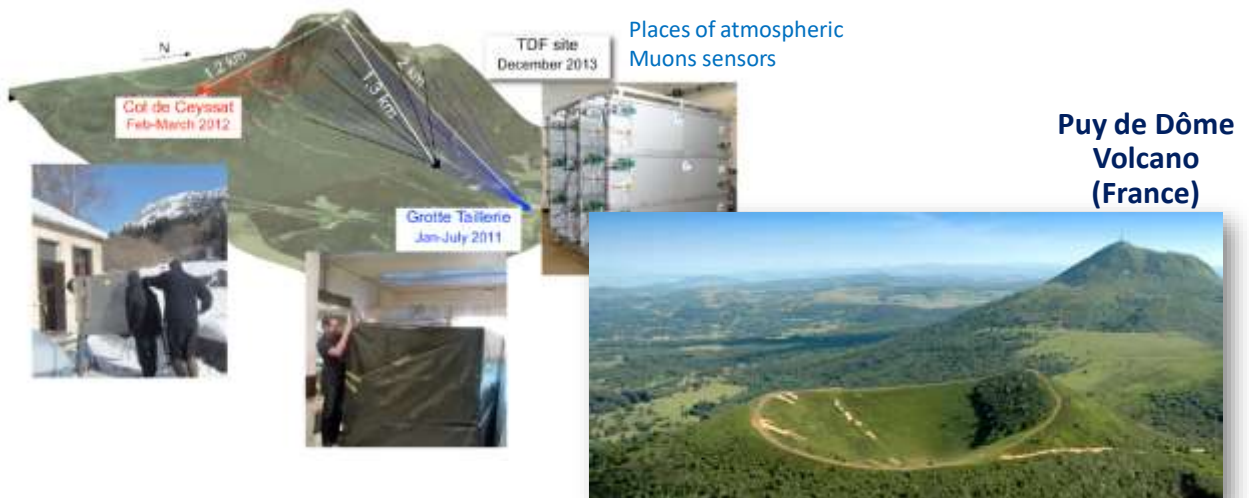
1. **An object oriented approach has to be chosen for every stochastic objects which has its own random stream.**
2. Select **a modern and statistically sound generators** according to the most stringent testing battery (TestU01);
3. Select **a fine parallelization technique adapted to the selected generator,**
4. The simulation must first be designed as a sequential program which would emulate parallelism: this sequential execution – **with compiler flags set on 'repeatability'** – will be the reference to compare parallel and sequential execution at small scales on the same node.
5. Externalize, sort or give IDs to the results for reduction in order to keep the execution order or use compensated algorithms

[Hill 2015] :

Hill D., "Parallel Random Numbers, Simulation and reproducibility". IEEE/AIP - Computing in Science and Engineering, vol. 17, no 4, 2015, pp. 66-71.


TEST APPLICATION FOR PARALLEL MONTE CARLO SIMULATION OF MUONIC TOMOGRAPHY

'ORIGINAL' TEST/REFERENCE APPLICATION WITH PHYSICISTS
MUONIC TOMOGRAPHY – UP TO A BILLIONS OF THREADS...



PRINCIPLE OF MUONIC TOMOGRAPHY

Atmospheric muons go through matter. Depending on their energy and of the matter they traverse it is possible to reconstruct the inner image of a large edifice with multiple sensors (figure by Samuel Béné)





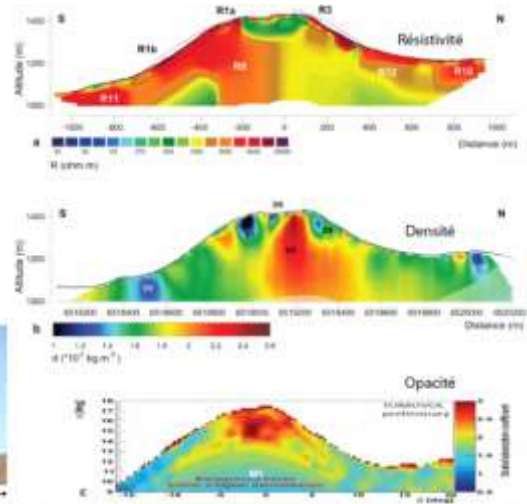
The muon is an elementary particle similar to the electron, with a negative charge and a spin of $1/2$, but with a much greater mass. It is classified as a lepton. The muon is not believed to have any sub-structure—that is, it is not thought to be composed of any simpler particles (as is the case of other leptons).

TOMUVOL PROJECT

<http://www.obs.univ-bpclermont.fr/tomuvol/presentation.php>



LMV (Laboratoire Magmas et Volcans) and LPC (Laboratoire de Physique Corpusculaire) made a joint venture with computer scientists for this TOMUVOL project (TOMographie MUonique des Volcans)



TARGET NODES WITH REGULAR XEON & INTEL XEON PHI

XEON PHI – STILL ON TOP CEA MACHINE LIKE JOLIOT – CURIE > 9 PF

Parallel stochastic simulation of muonic tomography – Aim finish computing **in less than 24h**

- Parallel programming model using p-threads <https://github.com/HeisSpiter/HPCsim>
- Each Muon is a stochastic object
- Multiple streams using MRG32k3a
- A billion threads handled by a single node
- Compiling flags set to maximum reproducibility – Sequential results obtained after 5 weeks – 3 months for a single Phi core (results below are with all CPU/Phi cores).

1st we did a round of sequential optimization with the code given by our physicists colleagues
16X on a single CPU core

Performance of a billion event simulation when parallelized on 1 Phi, 1 CPU, 2 CPUs

	Intel Xeon Phi 7120P	Intel Xeon E5-2650v2	2x Intel Xeon E5-2650v2
Time	48 h 49 min	36 h 32 min	18 h 17 min
Speedup	1	1.34	2.67

SCHWEITZER, P., MAZEL, C., FEHR, F., CÂRLOGANU, C., HILL D., "Proper parallel Monte Carlo for computed tomography of volcanoes", Proceedings of the 2013 International Conference on High Performance Computing & Simulation, ACM/IEEE/IFIP, Helsinki July 1st-5th, 2013, pp. 519-526.

REPRODUCIBILITY BETWEEN PHI & REGULAR XEON

FIRST ATTEMPTS

- First try with simple compilations of simulation to study the validity of the results
Intel C compiler with the "-O2 -g -Wall -Wextra" - (no -fast-math no aggressive -O3)
- For Xeon Phi, we added the "-mmic" option. (no -fast-math no aggressive -O3).
- We evaluate the deviation in the results when the compilation is left free (limited to 1000 muons events – muon reaching the detector)
- Very important differences in final muon energy have been noticed (up to 0.18 GeV)
- Also noticed important differences for the final position (up to 0.3 m).
- If the initial energy of the particle is between 5 GeV and 10 TeV, its final energy is between 0.15 GeV and 5 TeV (or even zero, if it does not even reach the detector). A difference of 0.18 GeV is therefore not acceptable.
- The detector has plans whose size is one meter by one meter. An inaccuracy of 0.3 m on the end position means a 30% inaccuracy on one dimension of the plane!

MORE CAREFUL ATTENTION TO COMPILER FLAGS

After different tries with Intel Compiler flags we retained the following:

- ✓ `"-fp-model precise -fp-model source -fimf-precision=high -no-fma"`
for the compilation on the Xeon Phi – (no `-fast-math` no aggressive `-O3`)
- ✓ `"-fp-model precise -fp-model source -fimf-precision=high"`
for the compilation on the Xeon CPU – (again no `-fast-math` no aggressive `-O3`)

With this set of flags, the results on the two architectures are reproducible (the same order).

Both of them have the same sign and the same exponent (even if some exceptions would be theoretically possible, they would be very rare and haven't been observed).

The only bits that can differ between these results are the least significant bits of the significand.

For a given exponent e , and a result $r1 = m \times 2e$, the closest value greater than $r1$ is $r2 = (m + \epsilon_d) \times 2e$, where ϵ_d is the value of the least significant bit of the significand: $\epsilon_d = 2^{-52} \approx 2.22 \cdot 10^{-16}$.

BIT FOR BIT REPRODUCIBILITY STUDY (x86 VS K10M)

As announced by Intel we cannot expect bit for bit reproducibility when working with such different architectures - in our case (x86 & k10m).

- However with the best compiler flags, **we observed bit for bit reproducibility in single precision but not in double precision** (but **with the best compiler flags we found for reproducibility**)
- The relative difference between processors (E5 vs Phi) in double precision were analyzed and are shown here.

Relative CPU-Phi differences between the results and number of altered bits

Difference ↓ \ Result →	Position X	Position Z	Direction X	Direction Y	Direction Z
0 bit: bit for bit reproducibility	4922	4934	4896	4975	4913
1 bit: $1.11\text{E-}16 \leq \Delta < 2.22\text{E-}16$	25	21	14	5	18
2 bits: $2.22\text{E-}16 \leq \Delta < 4.44\text{E-}16$	21	18	52	4	31
3 bits: $4.44\text{E-}16 \leq \Delta < 8.88\text{E-}16$	15	12	23	6	12
4 bits: $8.88\text{E-}16 \leq \Delta < 1.78\text{E-}15$	10	7	5	4	10
≥ 5 bits: $1.78\text{E-}15 \leq \Delta < 2.25\text{E-}11$	7	8	10	6	16

Run-to-Run Reproducibility of Floating-Point Calculations for Applications on Intel® Xeon Phi™ Coprocessors (and Intel® Xeon® Processors) – by Martin Cordel - <https://software.intel.com/en-us/articles/run-to-run-reproducibility-of-floating-point-calculations-for-applications-on-intel-xeon>

See also P. Schweitzer thesis & paper : SCHWEITZER P., CIPIERE S., DUFAURE A., PAYNO H., PERROT Y., HILL D. and MAIGNE L., "Performance evaluation of multi-threaded Geant4 simulations using an Intel Xeon Phi cluster", Scientific Programming, Article ID 980752, 10 pages, 2015. doi:10.1155/2015/980752.



CONCLUSION

COMPUTERS CAN BE AMPLIFIERS OF ERRORS...

- **Huge Numerical differences when we do not pay attention to compiler flags**
- **Repeatability achieved only for identical execution platforms.**
- **Comparison possible with sequential results – gain in confidence (with the given method)**
- **Numerical Reproducibility is possible (not repeatability) for Parallel Stochastic applications with independent computing on different architectures.**
- **Can be resilient to soft errors on supercomputers (use statistics – ‘N out of M’).**
- Key elements of a method have been presented to produce numerically reproducible results for parallel stochastic simulations **comparable with a sequential implementation** (before large scaling on coming Exascale systems)
- Numerical replication is important for scientists in many sensitive areas, finance, climate, nuclear safety, medicine...