

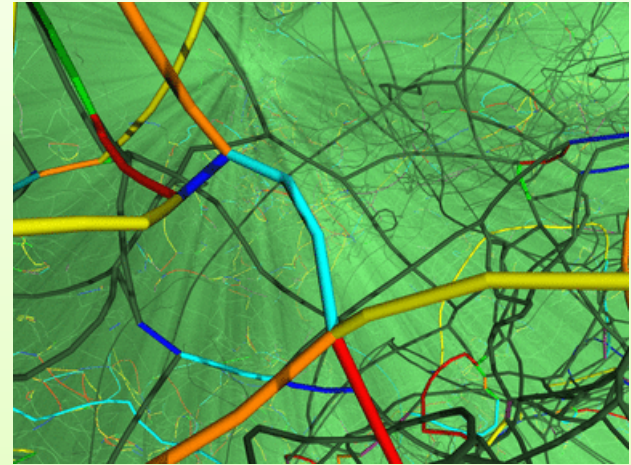
# Coming Soon | A Toolset to Fight Non-deterministic Bugs at Scale

Dong H. Ahn, Chris Chambreau, Ignacio Laguna, Greg Lee, Kento Sato, and Martin Schulz (LLNL)  
Simone Atzeni, Ganesh Gopalakrishnan and Zvonimir Ramarić (Univ. of Utah)  
Joachim Protze (RWTH Aachen University)



# Not all bugs are created in equal!

- LLNL's multiphysics code for laser simulation with over 1M LOC – Non-deterministic crashes w/ a threaded solver
- Above certain optimization levels and certain scales (8K MPI processes)
- **Cost scientists months ...**



- ParaDiS – Non-deterministic fatal conditions when scaled to tens of thousands of processes
- **Impacted science for 1 year...**

# Pruner aims at providing a reproducibility toolset to fight the harmful effects of non-determinism

```
$ srun -n 128 ./tst  
= 0.24512003
```

```
$ srun -n 128 ./tst  
= 0.24511999
```

```
$ srun -n 128 ./tst  
error: Segfault
```

```
$ srun -n 128 ./tst  
= 0.24512000
```



- Reproducibility
  - The ability to repeat program executions with the same results or the same behavior
- Many sources of non-determinism hamper this ability
- Greater challenges: non-determinism in concurrency scheduling
- Project Pruner
  - Provide a scalable reproducibility tool set

# Coming soon, pruning toolset to detect and control targeted sources of non-determinism

- **Archer** - Scalable and accurate OpenMP data race detector
- **ReMPI** - Scalable record and replay for MPI
- **NINJA** – a novel noise injection tool to expose unintended message races and significant complement ReMPI
- **FLiT** – floating point litmus tester from Utah which allows to test and quantify floating-point non-determinism induced by different compilers and compiler flags

In the process of assembling these components into a GitHub organization at <https://github.com/PRUNER> with a Q1/2017 GA release goal



# References

- Kento Sato, Dong H. Ahn, Ignacio Laguna, Gregory L. Lee, Martin Schulz, Chris Chembreau Noise Injection Techniques for Reproducing Subtle and Unintended Message Races, to appear in the Proceeding of the 22nd ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, Austin, Texas, Feb, 2017
- Simone Atzeni, Ganesh Gopalakrishnan, Zvonimir Rakamaric, Dong H. Ahn, Ignacio Laguna, Martin Schulz, Greg L. Lee, Joachim Protze, Matthias S. Müller, ARCHER: Effectively Spotting Data Races in Large OpenMP Applications, In the Proceedings of the 30th IEEE International Parallel and Distributed Processing Symposium (IPDPS 2016), Chicago May, 2016
- Kento Sato, Dong H. Ahn, Ignacio Laguna, Gregory L. Lee, Martin Schulz, Clock Delta Compression for Scalable Order-Replay of Non-Deterministic Parallel Applications, In the Proceedings of IEEE/ACM Supercomputing Conference (SC|15), Austin, Texas, Nov, 2015
- Joachim Protze, Simone Atzeni, Dong H. Ahn, Martin Schulz, Ganesh Gopalakrishnan, Matthias S. Müller, Ignacio Laguna, Zvonimir Rakamaric, and Greg L. Lee, Towards Providing Low-Overhead Data Race Detection for Large OpenMP Applications, In the proceedings of the LLVM Compiler Infrastructure in HPC (a workshop held in conjunction with SC14), Nov, 2014





**Lawrence Livermore  
National Laboratory**