

# *8th Workshop of the* INRIA-Illinois Joint Laboratory for PetaScale Computing

November 19-21 2012

Argonne National Laboratory

## Google “joint-lab” → first link

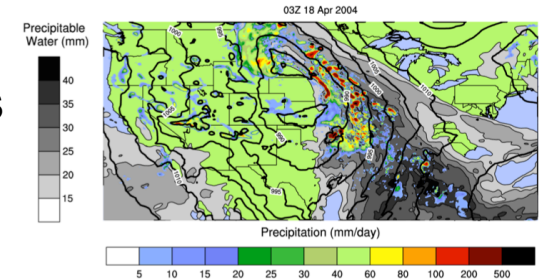
- Between INRIA and UIUC
- In the context of Blue Waters
- Adding soon ANL as partner
- 2 co-executive directors: Marc Snir (ANL, UIUC) and Franck Cappello (INRIA)
- 3 types of activities: workshops, visits, missions + *Collaborations: FR, EU, America, ASIA*
- Productions: Publications and Software
- Will last until end of 2013.
- Working on an extension – JLESC - to other partners



# INRIA-Illinois-ANL Eco system

## Applications (PRAC):

- Contagion on very large social networks
  - Earthquake system science
  - Formation of the first galaxies
- ETC.



**Users** (Center for  
Atmospheric research,  
Japan Earth quake, etc.)

**Application developers**


**System practitioners.**

**Algorithm research** and  
development (Numerical  
methods, Mapping + load  
balancing, FT, I/O)

**Hardware vendor** (CRAY, IBM)

**System software R&D**  
(Charm++, HDF5, MPICH,  
Integrated System Consol,  
FTI, etc.)

**HPC System**



**Mira**



# Joint-Lab Areas of Collaboration

- **Numerical Algorithms and Libraries**





(Communication avoiding, dynamic scheduling, linear algebra, adjoint methods, etc.);

- **Programming and runtime** (parallel languages, runtime, multi-threading, process mapping, scheduling, communication libraries, etc.)

- **Resilience** (FT-protocols, event log analysis, multi-level checkpointing, resilience performance modeling, energy consumption, etc. );

- **System software** (monitoring, resource-job manager, I/O & file systems, visualization, etc.).

- **Cloud** (HPC Cloud, Fault tolerance, testbed, etc.)







« Application driven »: From Applications to system

## Applications


Tornado Simulation



Earth Quake Simulation




Climate Simulation





## Numerical libraries

3D-FFT




Linear Algebra

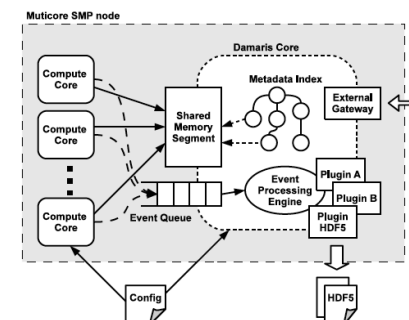


## System Software


New fault Tolerance Approaches



New I/O Approaches



New Archiving Approaches



- The most efficient way to implement collaborations: **Visits!**  
→ The joint-lab all travel and stay costs.

HOME  
ABOUT  
EVENTS  
PEOPLE  
RESEARCH  
► VISITS  
PUBLICATIONS  
SOFTWARE  
EDUCATION  
GRANTS  
JOINT LAB WIKI

## Researcher Visits to the Joint Lab

- [54] Tatiana Martsinkevich, INRIA, from July 2cd 2012 to July 2013
- [53] Mohamed Slim Bouguerra, INRIA, from July 12th to July 11th 2013
- [52] Matthieu Dorier, INRIA, from July 30th to Aug. 4th 2012
- [51] Alexandru Costan, INRIA, from July 23rd to July 27th 2012
- [50] Francois Tessier, INRIA, from May. 6th to May 19th 2012
- [49] Olivier Guittonneau, U. Paris 6, from May. 1st to Sept. 10th 2012
- [48] Maxime Laine, U. Paris 6, from Apr. 7cd to Sept. 10th 2012
- [47] Adrien Remy, INRIA, from Apr. 2cd to Apr. 30th 2012
- [46] Mitsuhsa Sato, University of Tsukuba, from Mar. 22cd to Mar. 23th 2012
- [45] Hideyuki Jitsumoto, University of Tsukuba, from Jan. 20th to Mar. 20th 2012
- [44] Vivek Kale, UIUC, from Jan 18th. to Feb. 4th 2012
- [43] Rodrigo Kassick, UFRGS, from Nov. 20th to Dec. 4th 2011
- [42] Laercio Pilla, INRIA, from Nov. 20th to Dec. 4th 2011

## Publications:

HOME  
ABOUT  
EVENTS  
PEOPLE  
RESEARCH  
VISITS  
► PUBLICATIONS  
SOFTWARE  
EDUCATION  
GRANTS  
JOINT LAB WIKI

### Joint Laboratory Publications

#### A Framework to Estimate Energy Consumption of Fault Tolerance protocols during HPC executions

Poster, to appear in Proceedings of ACM PPoPP 2013  
*M. El Mehdi Diour, O. Gliuck, L. Lefevre, F. Cappello*

#### Fault prediction under the microscope: A closer look into HPC systems

Proceedings of IEEE/ACM SC12  
*A. Gainaru, F. Cappello, Marc Snir, Bill Kramer*

#### Damaris: How to Efficiently Leverage Multicore Parallelism to Achieve Scalable, Jitter-free I/O

Proceedings of IEEE Cluster 2012  
*M. Dorier, G. Antoniu, F. Cappello, M. Snir, L. Orf*

#### Hierarchical Clustering Strategies for Fault Tolerance in Large Scale HPC Systems

Proceedings of IEEE Cluster 2012  
*L. B. Gomez, T. Ropars, N. Maruyama, F. Cappello, S. Matsuoka*

#### A Hierarchical Approach for Load Balancing on Parallel Multi-core Systems

Proceedings of ICPP2012  
*L.L. Pilla, C. Pousa Ribeiro, D. Cordeiro, C. Mei, A. Bhatele, P. O. A. Navaux, F. Broquedis, J.-F. Mehaut, L. V. Kale*

#### Scalable Reed-Solomon-based Reliable Local Storage for HPC Applications on IaaS Clouds

Proceedings of EuroPar 2012.  
*L. Bautista Gomez, B. Nicolae, N. Maruyama, F. Cappello, S. Matsuoka*

#### Energy considerations in Checkpointing and Fault Tolerance protocol

Proceeding of IEEE/IFIP DSN/FTXS 2012.  
*M. El Mehdi Diour, O. Gliuck, L. Lefevre, F. Cappello*

## Software being transferred to production:

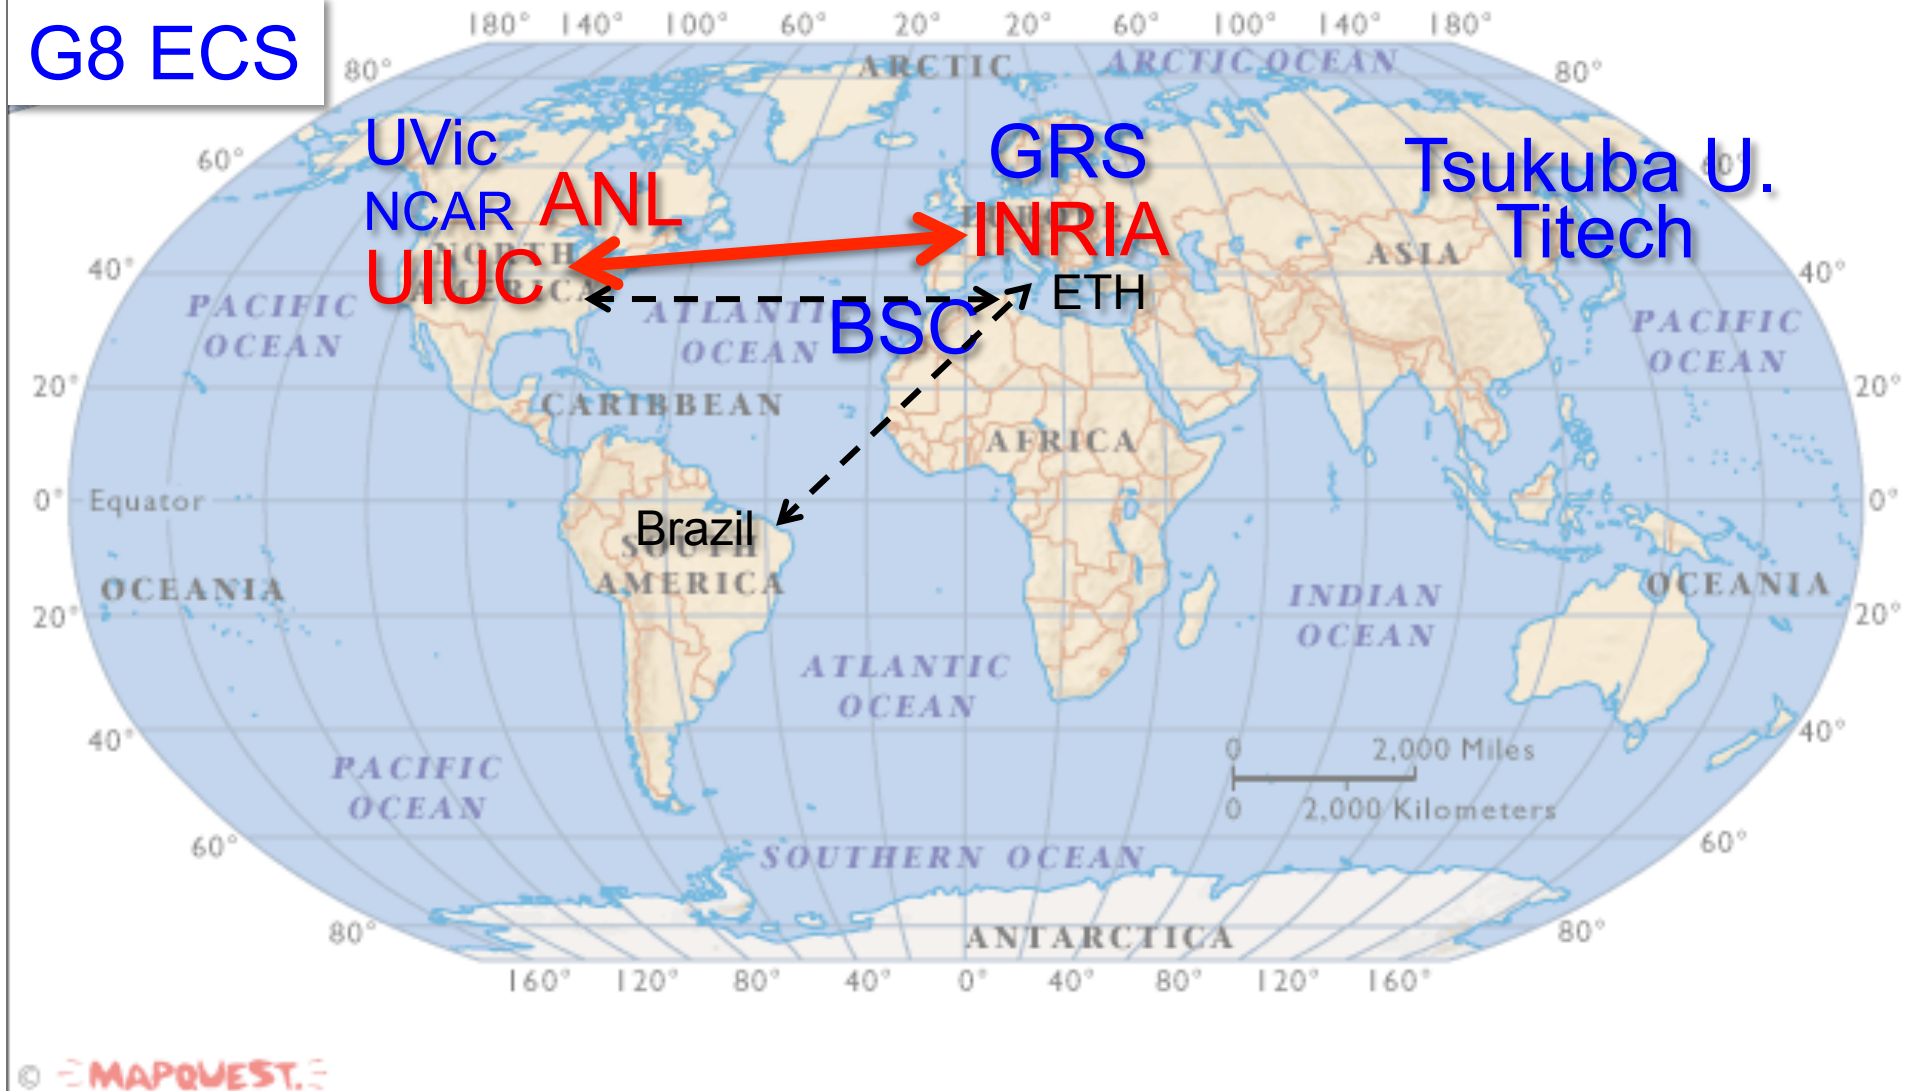
- HELO, ELSA (system event analysis and failure prediction)
- Damaris (I/O and visualization)
- FTI (fault tolerance Interface, multilevel checkpoint/restart)



# International collaborations - opportunities

JLPC

G8 ECS



# 8th Workshop Highlights

- Session 1 **Big Apps, Big DATA - Big I/O**  
I/O, file system, Scientific Data Generator, Visualization
- Session 2 **Programming Models/Runtime**  
Tri-diagonal Solvers, MPI3, Actor model for Multicore Architectures
- Session 3 **Numerical algorithms and Methods**  
Automatic tuning, CBM data analysis, Adjoint codes
- Session 4 **Big Systems**  
OS and Runtime Software for Exascale, Blue Waters
- Session 5 **Cloud**  
Big Data, Contrail European project, Workflow Scheduling,
- Session 6 **Resilience**  
Preventive+Proactive Checkpointing, CIFTS, failure prediction

# Mini-workshops



Smaller groups and focus on specific topic

More like a brainstorming meeting. Try to identify some 2-3 research issues of common interest.

Define how the collaboration will be implemented in the next 6 months. Student visit is the most efficient way.

The research results will be presented at the next workshop

Visits could be to UIUC, ANL or France (any INRIA locations)  
The Joint-lab will support visits.

This is the perfect time to define topic of internships



# Workshop Logistics 1/3

- Monday plenary session in **1416**
- Tuesday and Wednesday sessions in **1416** and **1407**
- We reconvene Wednesday in **1416** before lunch for report and closing
- For any question, please ask **Janet** in the atrium

# Workshop Logistics 2/3

- All lunches will take place in this room
  - Lunch boxes on Wednesday
- Breaks will take place in the atrium
- Dinners (Monday and Tuesday) will take place in different restaurants
- Buses will bring us to restaurants




# Workshop Logistics 3/3

- Presentation: Use your laptop
- We collect ALL presentations
- Presentations will be stored on the join-lab web site.
- Please put you presentation on the USB key, in pdf format, right after your talk.



Joint Laboratory  
for Petascale Computation

# Acknowledgments



Joint Laboratory  
for Petascale Computation

# Enjoy the Workshop!

