

Extreme Scale Resilience Home

The Argonne National Laboratory/MCS/[Extreme Scale Resilience](#) group covers fault tolerance and resilience for HPC simulations and data analytics at extreme scale

Lead: [Franck Cappello, ANL](#)

Topics and people

- **Multi-level Checkpoint / Restart:** Bogdan Nicolae, Leonardo Bautista Gomez (Postdoc now at BSC), Franck Cappello
 - Main project: VeloC (ECP)
- **Lossy compression,** Sheng Di, Franck Cappello.
 - Main projects: EZ (ECP), CODAR (ECP)
- **Silent soft errors/data corruptions detectors and compression:** Sheng Di, Franck Cappello
 - Main project: Aletheia (NSF)
- **Failure characterization and prediction:** Sheng Di, Rinku Gupta, Franck Cappello
 - Main project: Catalog (DOE ASCR)
- **Failure modeling and fault tolerance optimizations:** Sheng Di
- **Fault tolerance protocols:** F. Cappello

Main collaborators: Marc Snir (ANL and UIUC), Jon Calhoun (Clemson), Bill Kramer (UIUC), Bogdan Nicolae (IBM Dublin), Thomas Ropars (EPFL), Amina Guermouche (UVSQ), Frederic Vivien (Inria), Yves Robert (LIP), Satoshi Matsuoka (Titech), Mitsuhsisa Sato (U. Tsukuba), Omer Subasi (BSC), Osman Unsal (BSC), Leonardo Bautista Gomez (BSC)

Tools and software

- **SZ (Error Bounded Lossy Compressor for floating point data sets)**
- **Z-checker (An lossy data compression assessment tool)**
- **AID (Adaptive Impact-Driven Detection) library for SDC detection**
- **FTI** (operational prototype): Fault Tolerance Interface for multi-level checkpoint/restart (in memory checkpointing, checkpointing on remote nodes, erasure encoding, etc.)
- **HELO/ELSA** (operational prototypes): System event clustering and Failure predictor
- **MPICH-HFT** (prototype under development): Fault tolerant MPI with hierarchical fault tolerant protocol

Main collaborative activities

- [Illinois-Inria-ANL-BSC-JSC-Riken-UTK Joint Laboratory on Petascale Computing](#)

Recent Publications (from 2013)

- S. Di, H. Guo, E. Pershey, M. Snir, F. Cappello, **Characterizing and Understanding HPC Job Failures over The 2K-day Life of IBM BlueGene/Q System**, IEEE/IFIP 49th International Conference on Dependable Systems and Networks (*IEEE DSN19*), Portland, USA, 2019.
- D. Tao, S. Di, X. Liang, Z. Chen, F. Cappello, **Optimizing Lossy Compression Rate-Distortion from Automatic Online Selection between SZ and ZFP**, in IEEE Transactions on Parallel and Distributed Systems (*IEEE TPDS*), 2019.
- X. Zou, T. Lu, W. Xia, X. Wang, W. Zhang, S. Di, D. Tao, F. Cappello, **Accelerating Relative-error Bounded Lossy Compression for HPC datasets with Precomputation-Based Mechanism**, in Proceedings of the 35th International Conference on Massive Storage Systems and Technology (*MSST19*), 2019.
- S. Di, H. Guo, R. Gupta, E. Pershey, M. Snir, F. Cappello, **Exploring Properties and Correlations of Fatal Events in a Large-Scale HPC System**, in IEEE Transactions on Parallel and Distributed Systems (*IEEE TPDS*), 2018.
- W. He, H. Guo, T. Peterka, S. Di, F. Cappello, HW Shen, **Parallel Partial Reduction for Large-Scale Data Analysis and Visualization**, in The 8th IEEE Symposium on *Large Data Analysis and Visualization (IEEE LDAV)* in conjunction with *IEEE VIS 2018*, Berlin, Germany, October 21, 2018.
- C. Wang, N. Dryden, F. Cappello, and M. Snir. **Neural Network Based Silent Error Detector**, in IEEE CLUSTER 2018, 2018. **[best paper award (in the programming and system software track)]**
- A. Murat Gok, S. Di, Y. Alexeev, D. Tao, V. Mironov, F. Cappello. **PaSTRI: Error-bounded Lossy Compression for Two-Electron Integrals in Quantum Chemistry**, in IEEE CLUSTER 2018, 2018. **[best paper award (in the application, algorithms and libraries track)]**
- X. Liang, S. Di, D. Tao, Z. Chen, and F. Cappello. **Efficient Transformation Scheme for Lossy Data Compression with Point-wise Relative Error Bound**, in IEEE CLUSTER 2018. **[best paper award (in the Data, Storage, and Visualization track)]**
- D. Tao, S. Di, X. Liang, Z. Chen, and F. Cappello. **Fixed-PSNR Lossy Compression for Scientific Data**, in IEEE CLUSTER 2018. (short paper)
- D. Tao, S. Di, X. Liang, Z. Chen and F. Cappello. **Optimization of Fault Tolerance for Iterative Methods with Lossy Checkpointing**, in 27th ACM Symposium on High-Performance Parallel and Distributed Computing (ACM HPDC2018), 2018.
- S. Di, D. Tao, X. Liang, and F. Cappello. **Efficient Lossy Compression for Scientific Data based on Pointwise Relative Error Bound**, in IEEE Transactions on Parallel and Distributed Systems (IEEE TPDS), 2018.
- H. Guo, S. Di, R. Gupta, T. Peterka, F. Cappello, **La VALSE: Scalable Visual Analysis of Logs for Fault Characterization on Supercomputers**, in EG Symposium on Parallel Graphics and Visualization (*ECPGV2018*), 2018.
- D. Tao, S. Di, Z. Chen, and F. Cappello. **In-Depth Exploration of Single-Snapshot Lossy Compression Techniques for N-Body Simulations**, Proceedings of the 2017 IEEE International Conference on Big Data (BigData2017), Boston, MA, USA, December 11 - 14, 2017, short paper.
- A. Murat Gok, D. Tao, S. Di, V. Mironov, Y. Alexeev, F. Cappello. **PaSTRI: A Novel Data Compression Algorithm for Two-Electron Integrals in Quantum Chemistry**, in IEEE/ACM 29th The International Conference for High Performance computing, Networking, Storage and Analysis (SC2017). [poster]
- D. Tao, S. Di, H. Guo, Z. Chen, and F. Cappello. **Z-checker: A Framework for Assessing Lossy Compression of Scientific Data**. in The International Journal of High Performance Computing Applications (IJHPCA), 2017.

- S. Di, F. Cappello. **Optimization of Error-Bounded Lossy Compression for Hard-to-Compress HPC Data**. in IEEE Transactions on Parallel and Distributed Systems (IEEE TPDS), 2017.
- E. Berrocal, L. Bautista-Gomez, S. Di, Z. Lan, and F. Cappello. **Toward General Software Level Silent Data Corruption Detection for Parallel Applications**. in IEEE Transactions on Parallel and Distributed Systems (IEEE TPDS), 2017.
- F. Cappello, R. Gupta, S. Di, E. Constantinescu, T. Peterka, and S. M. Wild. **Understanding and improving the trust in results of numerical simulations and scientific data analytics**. in 10th workshop on resilience in high performance computing (resilience) in Clusters, Clouds and Grids, in the conjunction with 23rd International European Conference on Parallel and Distributed Computing (Euro-Par), 2017.
- I. T. Foster, M. Ainsworth, B. Allen, J. Bessac, F. Cappello, J. Youl Choi, E. M. Constantinescu, P. E. Davis, S. Di, et al.. **Computing Just What You Need: Online Data Analysis and Reduction at Extreme Scales**. in 23rd International European Conference on Parallel and Distributed Computing (*Euro-Par 2017*), 2017. pp. 3-19.
- D. Tao, S. Di, Z. Chen, and F. Cappello. **Exploration of Pattern-Matching Techniques for Lossy Compression on Cosmology Simulation Data Sets**. Proceedings of the 1st International Workshop on Data Reduction for Big Scientific Data (DRBSD1) in Conjunction with ISC'17, Frankfurt, Germany, June 22, 2017.
- S. Di, Y. Robert, F. Vivien, and F. Cappello. **Toward an Optimal Online Checkpoint Solution under a Two-Level HPC Checkpoint Model**, in IEEE Transactions on Parallel and Distributed Computing (IEEE TPDS), 2017.
- S. Di, R. Gupta, E. Pershey, M. Snir, F. Cappello. **LogAider: A tool for mining potential correlations in HPC Log Events**. in IEEE/ACM 17th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (ACM CCGrid2017), Spain, 2017.
- D. Tao, S. Di, F. Cappello. **A Novel Algorithm for Significantly Improving Lossy Compression of Scientific Data Sets**, " in International Parallel and Distributed Processing Symposium (IEEE/ACM IPDPS 2017), Orlando, Florida, 2017.
- Pierre-Louis Guhur, Hong Zhong, Tom Peterka, Emil Constantinescu and Franck Cappello, **Lightweight and Accurate Silent Data Corruption Detection in Ordinary Differential Equation Solvers**, EuroPar 2016
- E. Berrocal, L. Bautista Gomez, S. Di, Z. Lan and F. Cappello, **Exploring Partial Replication to Improve Lightweight Silent Data Corruption Detection for HPC Applications**, EuroPar 2016
- O. Subasi, S. Di, L. Bautista-Gomez, P. Balaprakash, O. Unsal, J. Labarta, A. Cristal and F. Cappello. **Spatial Support Vector Regression to Detect Silent Errors in the Exascale Era**, IEEE/ACM CCGRID'2016
- S. Di, F. Cappello, **Adaptive Impact-Driven Detection of Silent Data Corruption for HPC Applications**, IEEE Transactions on Parallel and Distributed Computing, to appear, 2016
- S. Di, F. Cappello, **Fast Error-bounded Lossy HPC Data Compression with SZ**, IEEE IPDPS 2016
- L. Bautista-Gomez, A. Gainaru, S. Pernau, D. Tiwari, S. Gupta, F. Cappello, C. Engelmann, M. Snir, **Reducing Waste in Large Scale Systems through Introspective Analysis**, IEEE IPDPS 2016
- T. Martsinkevich, T. Ropars, F. Cappello, **Addressing the last roadblock for message logging in HPC: alleviating the memory requirement using dedicated resources**, Euro-Par 2015 workshop on Resilience - Resiliency in High Performance Computing with Clouds, Grids, and Clusters, 2015
- L. Bautista Gomez and F. Cappello, **Detecting Silent Data Corruption for Extreme-Scale MPI Applications**, EuroMPI 2015
- T. Martsinkevich, O. Subasi, O. Unsal, F. Cappello and J. Labarta, **Fault-tolerant protocol for hybrid task-parallel message-passing applications**, FTS 2015 workshop at IEEE Cluster 2015
- L. Bautista-Gomez and F. Cappello, **Detecting and Correcting Data Corruption in Stencil Applications through Multivariate Interpolation**, FTS 2015 workshop at IEEE Cluster 2015
- L. Bautista-Gomez and F. Cappello, **Exploiting Spatial Smoothness in HPC Applications to Detect Silent Data Corruption**, IEEE HPCC 2015
- E. Berrocal, L. Bautista-Gomez, S. Di, Z. Lan, F. Cappello, **Lightweight Silent Data Corruption Detection Based on Runtime Data Analysis for HPC Applications**, short paper, ACM HPDC 2015
- S. Di, E. Berrocal, F. Cappello, **An Efficient Silent Data Corruption Detection Method with Error-feedback Control and Even Sampling for HPC Applications**, IEEE CCGRID 2015
- S. Di, E. Berrocal, K. Heisey, L. Bautista-Gomez, R. Gupta, F. Cappello, **Towards Effective Detection of Silent Data Errors for HPC Applications**, Poster, IEEE/ACM SC14
- L. Bautista Gomez, P. Balaprakash, S. Bouguerra, S. Wild, F. Cappello and P. Hovland, **Energy-Performance Tradeoffs in Multilevel Checkpoint Strategies**, Poster, IEEE Cluster 2014
- S. Di, F. Cappello, **GloudSim: Google Trace based Cloud Simulator with Virtual Machines**, in Journal of Software: Practice and Experience (*Wiley SPE*), 2014.
- S. Bouguerra, A. Gainaru, F. Cappello, **Failure prediction: what to do with unpredicted failures?**, to appear in *International Journal of High Performance Computing Applications*.
- F. Cappello, A. Geist, B. Gropp, B. Kramer, M. Snir, **Toward Exascale Resilience: 2014 update**, International Journal on Supercomputing Frontiers and Innovations, Vol 1, Num 1, 2014, <http://superfri.org/superfri/article/view/14>
- S. Di, L. Bautista-Gomez, F. Cappello, **Optimization of Multi-level Checkpoint Model with Uncertain Execution Scales**, to appear in IEEE /ACM SC14
- M. Snir et al. **Addressing failures in exascale computing**, to appear in *International Journal of High Performance Computing Applications*, 2014
- S. Di, D. Kondo, and F. Cappello, **Characterizing and Modeling Cloud Applications/Jobs on a Google Data Center**, To appear in Journal of Supercomputing, 2014.
- L. Bautista-Gomez, Franck Cappello, et. al, **GP GPUs: How to Combine High Computational Power with High Reliability (Embedded Tutorial)**, Design, Automation & Test in Europe, DATE'14
- S. Di, S. Bouguerra, L. Bautista Gomez, F. Cappello, **Optimization of Multi-level Checkpoint Model for Large Scale HPC Applications**, IEEE IPDPS 2014
- S. Di, C.-L. Wang, F. Cappello, **Adaptive Algorithm for Minimizing Cloud Task Length with Prediction Errors**, IEEE transaction on Cloud Computing.
- L. Bautista Gomez and F. Cappello, **Detecting Silent Data Corruption through Data Dynamic Monitoring for Scientific Applications**, Poster, to appear in Proceedings of ACM PPoPP 2014
- G. Bosilca, A. Bouteiller, E. Brunet, F. Cappello, J. Dongarra, A. Guermouche, T. Herault, Y. Robert, F. Vivien, D. Zaidouni, **Unified Model for Assessing Checkpointing Protocols**, To appear in Concurrency and Computation: Practice and Experience, Wiley, 2013
- L. Bautista Gomez, F. Cappello, **Improving Floating Point Compression through Binary Masks**, Proceedings of IEEE BigData 2013
- T. Ropars, T. Martsinkevich, A. Guermouche, A. Schiper, F. Cappello, **SPBC: Leveraging the Characteristics of MPI HPC Applications for Scalable Checkpointing**, Proceedings of IEEE/ACM SC13
- S. Di, Y. Robert, F. Vivien, D. Kondo, C. L. Wang, F. Cappello, **Optimization of Cloud Task Processing with Checkpoint-Restart Mechanism**, Proceedings of IEEE/ACM SC13
- A. Bouteiller, F. Cappello, J. Dongarra, A. Guermouche, T. Herault and Y. Robert, **Multi-criteria checkpointing strategies: optimizing response-time versus resource utilization**, Proceedings of EuroPar 2013
- S. Di, D. Kondo, F. Cappello, **Characterizing Cloud Applications on a Google Data Center**, short paper, Proceedings of ICPP2013
- A. Gainaru, F. Cappello, M. Snir, B. Kramer, **Failure prediction for HPC systems and applications: current situation and open issues**, International Journal of High Performance Computing Applications, SAGE, 2013

- *B. Nicolae, F. Cappello, AI-Ckpt: Leveraging Memory Access Patterns for Adaptive Asynchronous Incremental Checkpointing*, Proceeding of ACM HPDC 2013
- *B. Nicolae, F. Cappello, BlobCR: Virtual Disk Based Checkpoint-Restart for HPC Applications on IaaS Clouds*, To appear in Journal of Parallel and Distributed Computing, 2013
- *M. El Mehdi Diour, O. Gluck, L. Lefevre, F. Cappello, ECOFIT: A Framework to Estimate Energy Consumption of Fault Tolerance protocols during HPC executions*, Proceedings of IEEE CCGRID 2013
- *M. S. Bouguerra, A. Gainaru, F. Cappello, L. Bautista Gomez, N. Maruyama and S. Matsuoka, Improving the computing efficiency of HPC systems using a combination of proactive and preventive checkpointing*, Proceedings of IEEE IPDPS 2013
- *M. El Mehdi Diouri, O. Gluck, L. Lefevre, F. Cappello, Towards an energy estimator of fault tolerance protocols*, Poster, in Proceedings of ACM PPoPP 2013