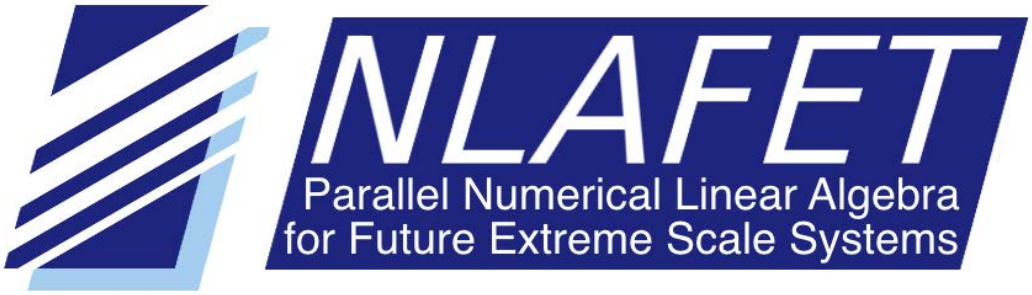

H2020-FETHPC-2014: GA 671633

D7.7

Dissemination report

Period M19-M42

April 2019

Document information

Scheduled delivery	2019-04-30
Actual delivery	2019-04-29
Version	1.0
Responsible partner	UMU

Dissemination level

PU – Public

Revision history

Date	Editor	Status	Version	Changes
2017-04-25	Lennart Edblom	Draft	0.1	First draft version
2017-04-29	Lennart Edblom	Final	1.0	Final revision after internal reviews

Authors

Lennart Edblom and Bo Kågström, UMU

Internal reviewers

Sven Hammarling, UNIMAN, Iain Duff, STFC, Laura Grigori, INRIA

Contributors

Team members from all partners have contributed to the content.

Copyright

Acknowledgements

This project has received funding from the *European Union's Horizon 2020 research and innovation programme* under the grant agreement number 671633.

Table of Contents

1. Introduction

The Description of Action document (GA Annex 1) states for deliverable D7.7:

"*D7.7: Second dissemination report*
Final dissemination report, covering M19-M42."

This deliverable is in the context of tasks 7.1-7.4 of WP7: *Dissemination and community outreach*.

The goal and main impact of NLAFET is to develop and deploy new algorithms, tools and library software that leading-edge applications need in order to attain high performance and thereby be able to utilize future extreme-scale systems effectively. To realize the expected impacts and to make NLAFET well known and recognized, we engage in relevant dissemination and outreach activities. Our stakeholders and audience include the world-wide scientific community, hardware and software vendors, application experts, and the general computing community.

In the following, we describe our efforts for achieving the objectives presented in the M3-deliverable D7.1 *Dissemination and community outreach plan*, and present the results of the main activities during the period May 1, 2017 until April 30, 2019 (M19-M42).

2. Dissemination and outreach activities

This main section of the document is structured with respect to the four WP7 tasks: *Dissemination* (Task 7.1), *Website* (Task 7.2), *Open source activities* (7.3), and *Community outreach* (Task 7.4).

2.1. Dissemination (Task 7.1)

The dissemination activities reported here include information about public deliverables of the project, scientific publications, presentations at and organization of conferences, workshops, and minisymposia.

Public NLAFET deliverables

All public deliverables are published on the website although not all are yet approved by the European Commission (EC), and available online at http://www.nlafet.eu/public-deliverables.

During the first 18 months (M1-M18) of the project 16 deliverables were approved by EC and published. These are listed in deliverable report D7.4 and available online.

During the second reporting period up to now (M19-M42), in total 23 deliverable reports have been submitted and published on the website. The listing of the deliverables below is ordered by submission date (e.g., M24 means end of month 24).

Month 24:

D3.2 Algorithm design for symmetrically structured factorizations (M24 - STFC)
D3.5 Software for for Highly Unsymmetric Factorizations (M24 - STFC)
D3.6 Algorithm Design for Hybrid Methods (M24 – STFC)

D4.3 Prototype software, phase 2 (M24 - INRIA)
D6.2 Novel methods for static and dynamic scheduling (M24 - UNIMAN)
D6.6 Algorithm-based fault tolerance techniques (M24 - UNIMAN)

Month 30:

D2.2 Scalability and tunability of factorization algorithms (M30 – UMU)
D2.6 Prototype software for eigenvalue problem solvers (M30 - UMU)
D4.4 Performance evaluation (M30 – INRIA)
D5.2 Software integration (M30 – INRIA)

Month 33, 39, and 40:

D7.6 Batched BLAS Specification (M33 – UNIMAN)
D2.4 Batched BLAS 2018 Reference Implementation (M36 – UNIMAN)
D4.5 Integration (M36 - INRIA)
D3.3 Software for Symmetrically Structured Factorizations (M39 - STFC)
D3.7 Software for Hybrid Methods (M40 – STFC)

The final eight (8) deliverables will be available as soon as they are submitted to the EC portal:

D2.7 Eigenvalue solvers for non-symmetric problems (M42 – UMU)
D2.9 Novel SVD algorithms (M42 - UNIMAN)
D5.3 Validation and evaluation (M42 - INRIA)
D6.3 Evaluation of software prototypes (M42 – UMU)
D6.5 Evaluation of auto-tuning techniques (M42 - UMU)
D6.7 Prototypes for tiled one-sided factorizations with algorithm-based fault tolerance. (M42 - UNIMAN)
D7.7 Second dissemination report (M42 - UMU)
D7.8 Release of the NLAFET library (M42 - UMU)

Scientific publications - NLAFET Working Notes

In addition to the deliverable documents, progress and results of the NLAFET project are regularly published in the NLAFET Working Notes series. The first 13 Working Notes are listed in deliverable D7.4, the dissemination report for M1-M18. During May 2017 to April 2019 the following reports have been published. All reports are available online at http://www.nlafet.eu/working-notes/ .

WN14: Iain Duff and Florent Lopez**:** Experiments with sparse Cholesky using a parametrized task graph implementation**.** *NLAFET Working Note WN-14*, June, 2017. Also as Technical Report RAL-TR-2017-006, Science & Technology Facilities Council, UK.

WN15: Maksims Abalenkovs, Jack Dongarra, Mark Gates, Azzam Haidar, Jakub Kurzak, Piotr Luszczek, Mawussi Zounon, Samuel Relton, Jakub Sistek, David Stevens, Ichitaro Yamazaki, Asim YarKhan*:* PLASMA 17 Functionality Report: Parallel BLAS and Norms, Linear Systems and Least Squares, Mixed Precision and Matrix Inversion. *NLAFET Working Note WN-15,* June, 2017. Also published as LAPACK Working Notes 293 (http://www.netlib.org/lapack/lawnspdf/lawn293.pdf).

WN16: Maksims Abalenkovs, Negin Bagherpour, Jack Dongarra, Mark Gates, Azzam Haidar, Jakub Kurzak, Piotr Luszczek, Samuel Relton, Jakub Sistek, David Stevens, PanruoWu, Ichitaro Yamazaki, Asim YarKhan, Mawussi Zounon: PLASMA 17 Performance Report. Linear Systems and Least Squares; Haswell, Knights Landing, POWER8. *NLAFET Working Note WN-16,* June, 2017. Also published as LAPACK Working Notes 292 (http://www.netlib.org/lapack/lawnspdf/lawn292.pdf).

WN17: Iain Duff, Florent Lopez and Stojce Nakov: Sparse direct solution on parallel computers. *NLAFET Working Note WN-17*, October, 2017. Also published as Technical Report RAL-TR-2017-010, Science & Technology Facilities Council, UK.

WN18: Mahmoud Eljammaly, Lars Karlsson and Bo Kågström: An Auto-Tuning Framework for a NUMA-Aware Hessenberg Reduction Algorithm. *NLAFET Working Note WN-18*, October, 2017. Also as Report UMINF 17.19, Dept. of Computing Science, Umeå University, SE-901 87 Umeå, Sweden.

WN19: Jan Papez, Laura Grigori, Radoslav Stompor: Solving linear equations with messenger-field and conjugate gradients techniques – an application to CMB data analysis. *NLAFET Working Note WN-19*, March 2018. Also as INRIA Research Report 9157, Paris, France.

WN20: Sébastien Cayrols, Iain Duff and Florent Lopez: Parallelization of the solve phase in a task-based Cholesky solver using a sequential task flow model. *NLAFET Working Note WN-20*, October, 2018. (Submitted for publication). Also published as Technical Report RAL-TR-2018-008, Science & Technology Facilities Council, UK.

WN21: Iain Duff, Jonathan Hogg and Florent Lopez: A new sparse symmetric indefinite solver using A Posteriori Threshold Pivoting. *NLAFET Working Note WN-21*, December, 2018. (Submitted for publication). Also published as Technical Report RAL-TR-2018-012, Science & Technology Facilities Council, UK.

WN22: Timothy Davis, Iain Duff and Stojce Nakov: Design and implementation of a parallel Markowitz threshold algorithm. (Submitted for publication). *NLAFET Working Note WN-22*, February, 2019. Also published as Technical Report RAL-TR-2019-003, Science & Technology Facilities Council, UK.

Scientific publications - journals and proceedings

In total 24 peer-reviewed scientific publications with NLAFET researchers as (co)authors, presenting results from NLAFET work, have been published during the reporting period. These are:

Björn Adlerborn, Lars Karlsson, and Bo Kågström. *Distributed One-Stage Hessenberg-Triangular Reduction with Wavefront Scheduling.* SIAM Journal of Scientific Computing, 40 (2):C157-C180, 2018. (https://doi.org/10.1137/16M1103890)

Alan Ayala, Xavier Claeys, Laura Grigori: *Affine low-rank approximations*. Journal of Scientific Computing, 79:1135-1160, 2019. https://doi.org/10.1007/s10915-018-0885-5

Zvonimir Bujanović, Lars Karlsson, Daniel Kressner: *A Householder-Based Algorithm for Hessenberg-Triangular Reduction.* SIAM Journal on Matrix Analysis and Applications, SIAM Publications 2018, Vol. 39, (3) : 1270-1294. (DOI: 10.1137/17M1153637)

Jack Dongarra, Mark Gates, Azzam Haidar, Jakub Kurzak, Piotr Luszczek, Panruo Wu, Ichitaro Yamazaki, Asim YarKhan, Maksims Abalenkovs, Neigin Bagherpour, Sven Hammarling, Jakub Sistek, David Stevens, Mawussi Zounon, Samuel Relton: *PLASMA: Parallel Linear Algebra Software for Multicore Using OpenMP.* ACM Transactions on Mathematical Software. Accepted July 24, 2018.

Jack Dongarra, Sven Hammarling, Nick Higham, Samuel Relton, Pedro Valero-Lara, and Mawussi Zounon: *The Design and Performance of Batched BLAS on Modern High-Performance Computing Systems*. ICCS 17, ETH Zurich, Procedia Computer Science, Volume 108, 2017, Pages 495-504, DOI: https://doi.org/10.1016/j.procs.2017.05.138

Jack Dongarra, Sven Hammarling, Nicholas J. Higham, Samuel D. Relton, Pedro Valero-Lara and Mawussi Zounon: *Optimized Batched Linear Algebra for Modern Architectures*. Euro-Par 2017: Parallel Processing, F.F. Rivera, T.F. Pena, and J.C. Cabaleiro, editors, volume 10417 of Lecture Notes in Computer Science, Springer-Verlag, Cham, 2017, pages 511–522. DOI: 10.1007/978-3-319-64203-1_37.

Jack Dongarra, Sven Hammarling, Nicholas J. Higham, Samuel D. Relton, and Mawussi Zounon: *Creating a Standardised Set of Batched BLAS Routines*. Proceedings of the Fourth Workshop on Sustainable Software for Science: Practice and Experiences (WSSSPE4, 2016), Gabrielle Allen, Jeffrey Carver et al, volume 1686, CEUR Workshop Proceedings.

Iain S. Duff, Florent Lopez and Stojce Nakov: *Sparse direct solution on parallel computers.* In Numerical Analysis and Optimization: NAOIV 2017, editors M. Al-Baali and L. Grandinetti and A. Purnama. Springer Proceedings in Mathematics & Statistics 235, June 2018, pp 67-98. ISBN 978-3-319-90026-1. DOI: https://doi.org/10.1007/978-3-319-90026-1_4

Iain S. Duff, Jonathan Hogg and Florent Lopez: *Experiments with sparse Cholesky using a sequential task-flow implementation*. Numerical Algebra, Control & Optimization, 2018, 8 (2) : 237-260. June 2018. DOI: 10.3934/naco.2018014

Iain S. Duff, Florent Lopez: *Experiments with Sparse Cholesky Using a Parametrized Task Graph Implementation*. In: Wyrzykowski R., Dongarra J., Deelman E., Karczewski K. (eds) Parallel Processing and Applied Mathematics. PPAM 2017. Lecture Notes in Computer Science, vol 10777,pp 197-206. Springer 2018. ISBN 978-3-319-78024-5 DOI: https://doi.org/10.1007/978-3-319-78024-5_18.

Mahmoud Eljammaly, Lars Karlsson, Bo Kågström: *On the Tunability of a New Hessenberg Reduction Algorithm Using Parallel Cache Assignment*. In: Wyrzykowski R., Dongarra J., Deelman E., Karczewski K. (eds) Parallel Processing and Applied Mathematics. PPAM 2017. Lecture Notes in Computer Science, vol 10777, pp 579-589. Springer, Cham. (https://doi.org/10.1007/978-3-319-78024-5_50)

Mahmoud Eljammaly, Lars Karlsson, Bo Kågström: *An auto-tuning framework for a NUMA-aware Hessenberg reduction algorithm*. ICPE '18 Companion of the 2018 ACM/SPEC International Conference on Performance Engineering, ACM Digital Library 2018 : 5-8. (DOI: 10.1145/3185768.3186304)

Azzam Haidar, Ahmad Abdelfattah, Mawussi Zounon, Stanimire Tomov, Jack Dongarra: *A Guide for Achieving High Performance With Very Small Matrices On GPU: A case Study of Batched LU and Cholesky Factorizations*. IEEE Transactions on Parallel and Distributed Systems, Vol. 29, No. 5, May 2018, DOI: 10.1109/TPDS.2017.2783929

Azzam Haidar, Ahmad Abdelfattah, Mawussi Zounon, Srikara Pranesh, Panruo Wu, Stanimire Tomov, and Jack Dongarra: *The Design of Fast and Energy-Efficient Linear Solvers: On The potential Of Half Precision Arithmetic And Iterative Refinement Techniques.* Computational Science - ICCS 2018 - 18th International Conference, DOI: 10.1007/978-3-319-93698-7.

Azzam Haidar, Stanimire Tomov, Jack Dongarra, Nick Higham: *Harnessing GPU's Tensor Cores Fast FP16 Arithmetic to Speedup Mixed-Precision Iterative Refinement Solvers.* The International Conference for High Performance Computing, Networking, Storage, and Analysis (SC18).

Laura Grigori, Sébastien Cayrols, and James W. Demmel: *Low rank approximation of a sparse matrix based on LU factorization with column and row tournament pivoting*. SIAM Journal on Scientific Computing, 40(2):181-209, 2018. https://doi.org/10.1137/16M1074527

Carl Christian Kjelgaard Mikkelsen, Lars Karlsson: *Blocked Algorithms for Robust Solution of Triangular Linear Systems*. In: Wyrzykowski R., Dongarra J., Deelman E., Karczewski K. (eds) Parallel Processing and Applied Mathematics. PPAM 2017. Lecture Notes in Computer Science, vol 10777, pp 68-78. Springer Cham. (https://doi.org/10.1007/978-3-319-78024-5_7)

Carl Christian Kjelgaard Mikkelsen, Angelika Beatrix Schwarz, Lars Karlsson: *Parallel robust solution of triangular linear systems*, Concurrency and Computation - Practice and Experience (DOI: 10.1002/cpe.5064)

Weifeng Liu, Ang Li, Jonathan Hogg, Iain S. Duff, Brian Vinter: *Fast synchronization-free algorithms for parallel sparse triangular solves*. Concurrency and Computation: Practice and Experience, 2017, vol 29, no 21. John Wiley & Sons. DOI: https://doi.org/10.1002/cpe.4244

Mirko Myllykoski: *A Task-Based Algorithm for Reordering the Eigenvalues of a Matrix in Real Schur Form*. In: Wyrzykowski R., Dongarra J., Deelman E., Karczewski K. (eds) Parallel Processing and Applied Mathematics. PPAM 2017. Lecture Notes in Computer Science, vol 10777, pp 207-216. Springer, Cham, (https://doi.org/10.1007/978-3-319-78024-5_19)

Mirko Myllykoski, Tuomo Rossi, Jari Toivanen: *On solving separable block tridiagonal linear systems using a GPU implementation of radix-4 PSCR method*. Journal of Parallel and Distributed Computing, Elsevier 2018, Vol. 115: 56-66. (DOI: 10.1016/j.jpdc.2018.01.004)

Jan Papez, Laura Grigori and Radek Stompor: *Solving linear equations with messenger-field and conjugate gradient techniques: An application to CMB data analysis*. Astronomy & Astrophysics Journal, vol 620:A59, 2018. http://arxiv.org/abs/1803.03462

Angelika Schwarz, Lars Karlsson: *Scalable eigenvector computation for the nonsymmetric eigenvalue problem.* Parallel Computing, Vol 85, pp131-140. Elsevier July 2019. DOI: https://doi.org/10.1016/j.parco.2019.04.001

Ichitaro Yamazaki, Jakub Kurzak, Panruo Wu, Mawussi Zounon and Jack Dongarra: *Symmetric Indefinite Linear Solver Using OpenMP Task on Multicore Architectures*. IEEE Transactions on Parallel and Distributed Systems, vol. 29, no. 8, pp 1879-1892, Aug. 2018. doi: 10.1109/TPDS.2018.2808964

Scientific publications - other technical reports

Other technical reports relevant for the NLAFET project include:

A. Ayala, X. Claeys, and L. Grigori: *Linear-time CUR approximation of BEM matrices*. Research Report RR-9208, INRIA, October 2018. (Now submitted)

L. Grigori and O. Tissot: *Reducing the communication and computational costs of enlarged Krylov subspaces conjugate gradient.* Technical Report 9023, INRIA, 2017. (Now submitted)

L. Grigori and O. Tissot: *Scalable Linear Solvers based on Enlarged Krylov subspaces with Dynamic Reduction of Search Directions*. Research Report RR-9190, INRIA, July 2018. (Now submitted)

Papers that are submitted or in preparation

H. Al Daas, L. Grigori, P. Jolivet, P. H. Tournier, *A robust multilevel additive Schwarz preconditioner.* (In preparation).

Sébastien Cayrols, Florent Lopez, Iain Duff et.al: *Solving linear least squares problems using block Cimmino.* (In preparation).

J. Demmel, L. Grigori, and A. Rusciano, *An improved analysis and unified perspective on deterministic and randomized low rank matrix approximations.* (In preparation).

Iain Duff and Florent Lopez: *Sparse LU factorization with a posteriori threshold pivoting*. (In preparation).

Iain Duff, Stojce Nakov: *Solving highly unsymmetric systems using distributed and shared memory parallelism.* (In preparation).

Laura Grigori, Jan Papez, Radek Stompor, and Olivier Tissot: *Solving CMB parametric component separation problem using subspace recycling*. (In preparation), 2019.

Ani Miraçi, Jan Papez, and Martin Vohralík: *A multilevel algebraic error estimator and the corresponding iterative solver with p-robust behavior.* (submitted for publication), March 2019.

Presentations etc related to NLAFET, describing NLAFET work and results

The work in and results from NLAFET have been presented at 24 different events, in several of them with more than one presentation. When there are several authors of a presentation in the listings below, the *speaker* is marked in italics.

- *Nordic e-Infrastructure Conference, NeIC2017, Umeå, May 31-June 1, 2017.*

  Plenary talk: NLAFET: Parallel Numerical Linear Algebra for Future Extreme Scale Systems. Bo Kågström, Umeå University.

- *Householder Symposium XX, Virginia Tech, Blacksburg, Virginia, June 18-23, 2017.*

  Plenary talk: Low rank approximation of a sparse matrix based on LU factorization with column and row tournament pivoting, Laura Grigori, INRIA, France.

  Plenary talk: NLAFET: parallel numerical linear algebra for future extreme scale systems. Bo Kågström, Umeå University.

  Presentation: New Developments in the Solution of Large Sparse Unsymmetric Systems. Iain Duff, STFC, UK.

  Poster presentation: How Fast Direct Solvers Can Benefit from GPU-acceleration. Mirko Myllykoski, Umeå University.

- *The Platform for Advanced Scientific Computing (PASC) Conference, Lugano, Switzerland, June 26-28, 2017.*

  Minisymposium presentation (MS45): Enlarged Krylov subspace methods for reducing communication: Laura Grigori, *Olivier Tissot*, INRIA, France

- *Power-Aware Computing Workshop - PACO2017. July 5-8, 2017, MPI Magdeburg, Germany.*

  Keynote speaker: A Look at Energy Saving on the Intel Knights Landing for Linear Algebra Computations. Jack Dongarra, University of Manchester.

  Keynote speaker: Communication avoiding algorithms for linear algebra kernels. Laura Grigori, INRIA.

- *46th International Conference on Parallel Processing, August 14-17, 2017, Bristol, UK*

  Keynote speaker: An Overview of Communication Avoiding Algorithms for Dense and Sparse Linear Algebra. Laura Grigori, INRIA

- *23rd International European Conference on Parallel and Distributed Computing (Euro-Par 2017), Santiago de Compostela, Spain, August 28-September 1, 2017.*

  Presentation: Optimized Batched Linear Algebra for Modern Architectures. Jack Dongarra, Sven Hammarling, Nicholas Higham, Samuel Relton and *Mawussi Zounon*, University of Manchester, UK.

- *Parallel Processing and Applied Mathematics (PPAM) 2017, September 10-13, 2017, Lublin, Poland.*

  Keynote speaker: An Overview of High Performance and a Look at Energy Saving on the Intel Knights Landing for Linear Algebra Computations, *Jack Dongarra, Univ. of Manchester.*

Presentations:

- On the Tunability of a New Hessenberg Reduction Algorithm Using Parallel Cache Assignment – *Mahmoud Eljammaly, Lars Karlsson and Bo Kågström, Umeå University*
- Blocked Algorithms for Robust Solution of Triangular Linear Systems – *Carl Christian Kjelgaard Mikkelsen and Lars Karlsson, Umeå University*
- A Task-Based Algorithm for Reordering the Eigenvalues of a Matrix in Real Schur Form – *Mirko Myllykoski, Umeå University*
- Experiments with sparse Cholesky using a parametrized task graph implementation – *Florent Lopez and Iain Duff, Rutherford Appleton Laboratory, STFC.*

- *Swedish e-Science Academy 2017, 11-12 October 2017, Umeå University, Sweden*

  Coorganizer, Welcome presentation: Bo Kågström, Umeå University.

  Poster presentations:

  - A Task-Based Algorithm for Reordering the Eigenvalues of a Matrix in Real Schur Form*, Mirko Myllykoski, Umeå University*
  - Blocked and Robust solution of triangular systems*, Carl Christian K. Mikkelsen and Lars Karlsson, Umeå University*
  - An Auto-tuning Framework for a NUMA-Aware Hessenberg Reduction Algorithm*, Mahmoud Eljammaly, Lars Karlsson and Bo Kågström, Umeå University*

- *SuperComputing 2017, Denver, November 12-17, 2017*

  Birds of a Feather Session organization: Batched, Reproducible and Reduced Precision BLAS. Organized by Jack Dongarra, Mawussi Zounon et al.

  BoF presentations:

  - Batched BLAS in applications. *Mawussi Zounon*, University of Manchester, UK.
  - BLAS Standardization Process. *Sven Hammarling*, University of Manchester, UK.

  Presentation: Linear Algebra Libraries for High-Performance Computing; Scientific Computing with Multicore and Accelerators. Jakub Kurzak, Jack Dongarra, Michael Heroux, James Demmel.

- *SIAM Conference on Parallel Processing for Scientific Computing, March 7-10, 2018, Waseda University, Tokyo, Japan.*

  Minisymposium MS43 organization: Parallel Numerical Linear Algebra for Future Extreme-Scale Systems – Part I of II: Iain Duff, STFC, UK & CERFACS, Toulouse, France and Bo Kågström, Umeå University, Sweden

  Minisymposium MS43 presentations:

  - Dense Linear Systems for Extreme Scale, *Jack J. Dongarra*, University of Tennessee, ORNL, and University of Manchester, USA
  - The Batched BLAS, *Sven J. Hammarling* and Mawussi Zounon, The University of Manchester, UK
  - Task Based Robust Computation of Eigenvectors, *Carl Christian Kjelgaard Mikkelsen,* Lars Karlsson, Mirko Myllykoski, and Angelika Schwarz, Umeå University, Sweden

Minisymposium MS55 organization: Parallel Numerical Linear Algebra for Future Extreme-Scale Systems – Part II of II: Iain Duff, STFC, UK & CERFACS, Toulouse, France and Bo Kågström, Umeå University, Sweden

Minisymposium MS55 presentations:

- Sparse Solvers on Extreme-Scale Systems*, Iain Duff*, STFC, UK and CERFACS, Toulouse, France
- Task-Based Sparse Direct Solver for Symmetric Indefinite Systems, *Florent Lopez*, Rutherford Appleton Laboratory, United Kingdom; Iain Duff, STFC, UK and CERFACS, Toulouse, France
- Robust Algebraic Preconditioners for Large Scale Applications, Laura Grigori, *Simplice Donfack*, and Olivier Tissot, INRIA, France
- Enlarged Conjugate Gradient Method for Reducing Communication, *Olivier Tissot* and Laura Grigori, INRIA, France

Minisymposium MS54 presentation: Iterative Refinement in Three Precisions for Fast and Accurate Solution of Ill-Conditioned Sparse Linear Systems. *Nicholas J. Higham,* University of Manchester, Erin C. Carson, New York University.

Minisymposia MS3 and MS14 organization: On Batched BLAS Standardization, Part I and II. Mawussi Zounon, Negin Bagherpour, University of Manchester, Azzam Haidar, Univ. of Knoxville, Siva Rajamanickam, Sandia National Laboratories.

Poster presentation: Efficient Robust Multi-Shift Triangular Solves, Angelika Schwarz, Mirko Myllykoski, Carl Christian Kjelgaard Mikkelsen, and Lars Karlsson, Umeå University, Sweden

- *15th Copper Mountain Conference on Iterative Methods, Colorado, USA. March 25-30, 2018.*

  Presentation: A Symbiotic Relationship of Direct and Iterative Methods to Target Sparse Solution on Extreme Scale Computers*.* Iain Duff, STFC, UK.

- *International Conference on Performance Engineering (ICPE 2018), April 9-13, 2018, Berlin, Germany.*

  Presentation: An Auto-Tuning Framework for a NUMA-Aware Hessenberg Reduction Algorithm. *Mahmoud Eljammaly,* Lars Karlsson and Bo Kågström, Umeå University.

- *2018 SIAM Conference on Applied Linear Algebra, May 4-8, 2018, Hongkong, China.*

  Invited plenary speaker: Enlarged Krylov Subspace Methods and Robust Preconditioners. Laura Grigori, INRIA

  Poster presentation: Messenger-field and conjugate gradients in Cosmic Microwave Background data analysis. Laura Grigori, *Jan Papez,* Radek Stompor, INRIA.

- *ICIAM Workshop on Applied and Industrial Mathematics 2018, Philadephia, USA, May 10-11 2018.*

  Presentation: The direct solution of sparse linear systems on extreme-scale computers. Iain Duff, STFC.

- *8th SIAM Workshop on Combinatorial Scientific Computing (CSC18). Bergen, Norway. June 6-8, 2018.*

  Presentation: Design and implementation of a parallel threshold Markowitz algorithm. *Iain Duff*, Stojce Nakov, STFC, UK.

- *International Supercomputing (ICS) 2018, June 24-28, 2018, Frankfurt, Germany*

  Birds of a Feather session organization and presentation: Leveraging mixed precision and emerging low precision operators in HPC application*. Jack Dongarra, Univ of Manchester and Univ. of Tennessee, Eric Petit, Intel, Marc Casas, BSC.*

  Birds of a Feather session organization and presentation: Batched BLAS Standardization*. Jack Dongarra, Univ of Manchester and Univ. of Tenessee, Mawussi Zounon, Sven Hammarling, University of Manchester, Piotr Luszczek, Innovative Computing Laboratoty, Knoxville, Pedro Valero-Lara, BSC.*

  Presentation in the Focus Session Future Accelerated Math Library Design: Numerical Linear Algebra for Future Extreme Scale Systems. Jack Dongarra, Univ. of Tennessee, Oak Ridge National Lab, and Univ. of Manchester

- *10th International Workshop on Parallel Matrix Algorithms and Applications (PMAA), ETH Zürich, Switzerland, June 27-29, 2018.*

  Presentation: Towards Distributed Tasking in the PLASMA Numerical Library. *Mawussi Zounon, University of Manchester,* George Bosilca, Jack Dongarra, Reazul Hoque.

  Presentation: Task-Based Sparse Direct solver for Symmetric Indefinite Systems. *Florent Lopez,* Iain Duff, STFC, UK.

  Presentation: Recycling Krylov method for the solution of sequence of linear systems. *Hussam Al Daas,* Laura Grigori, Pascal Hénon, Philippe Ricoux, Olivier Tissot, INRIA, France.

- *2018 SIAM Annual Meeting, July 9-13, 2018, Portland, Oregon, USA*

  Minisymposia MS150 – Low Precision Arithmetic for Dense Numerical Linear Algebra. Organized by Pierre Blanchard and Nicholas J. Higham, University of Manchester.

- *International Symposium on Computational Science at Scale, September 5-7, 2018, Erlangen, Germany*

  Invited keynote speaker: Scalable and Robust Linear Solvers. Laura Grigori, INRIA

- *Sparse Days Meeting @ CERFACS, Toulouse, September 27-28, 2018.*

  Presentation: Preconditioned Linear Solvers in CMB data analysis. Laura Grigori, *Jan Papez,* Radek Stompor, INRIA, France.

  Presentation: Task-based sparse direct solvers for symmetric systems in the NLAFET Project. Florent Lopez, STFC, UK.

- *Swedish e-Science Academy 2018, 16-17 October 2018, Uppsala, Sweden*

  Poster presentation: Task-based Structured QR Factorization using Parallel Critical Path*, Mahmoud Eljammaly,* Lars Karlsson and Bo Kågström, Umeå University.

- *PRACE-CoEs-FETHPC-EXDCI Workshop, Bruhl, Germany, October 30-31, 2018.*

  Presentation: NLAFET Research Objectives, Dissemination and Outreach Activities. Bo Kågström, Umeå University.

  Presentation: The CERFACS model and industry participation in the FETHPC Project NLAFET. Iain Duff, STFC, UK.

- *SuperComputing 2018, Dallas, Texas, November 12-15, 2018*

  Poster presentations:

  - Sparse linear system solvers. *Sébastien Cayrols*, STFC, UK
  - Parallel Numerical Linear Algebra for Future Extreme Scale Systems. *Sébastien Cayrols*, STFC, UK

- *2019 SIAM Conference on Computational Science and Engineering, February 25 - March 1, 2019, Spokane, USA*

  Invited plenary speaker: Communication Avoiding: The Past Decade and the New Challenges. Laura Grigori, INRIA.

  Minisymposium organization MS313: Parallel Numerical Linear Algebra for Future Extreme-scale Systems: Iain Duff, STFC.

  Minisymposium MS313 presentations:

  - Distributed Tasking in the PLASMA Numerical Library. *Mawussi Zounon,* Jakub Sistek, Jack Dongarra, University of Manchester, UK.
  - Solution of Sparse Unsymmetric Systems. Iain Duff, Florent Lopez and *Stojce Nakov*, STFC, UK.
  - The Scalability of Block Iterative Methods. *Sébastien Cayrols,* Iain Duff, STFC, UK
  - Preconditioned Linear Solvers in Cosmic Microwave Background Data Analysis. Laura Grigori, *Jan Papez,* Radek Stompor, INRIA, France.

  Minisymposium MS117 presentation: Standardisation of Batched BLAS. *Sven Hammarling,* University of Manchester, UK.

  Minisymposium MS286 presentations:

  - Exploiting Half Precision Arithmetic in Solving Ax=b. *Nicholas J. Higham,* Srikara Pranesh, Mawussi Zounon, University of Manchester.
  - Experiments with Mixed Precision Algorithms in Linear Algebra*. Jack Dongarra*, University of Manchester.

  Minisymposium MS231 presentation: Task-based Sparse Direct Solver for Symmetric Indefinite Systems. *Florent Lopez,* Iain Duff, STFC, UK.

  Minisymposia MS186 and MS 219 organization: Advances in Analyzing Floating-point Errors in Computational Science, part I and II. Nicholas J. Higham, Pierre Blanchard, University of Manchester.

- *Numerical algorithms for high-performance computational science meeting*, funded by the Royal Society Hooke. Scientific meeting, April 2019, London. ***Event organizers*** J. Dongarra, L. Grigori, and N. Higham.

  Poster presentation: On communication avoiding methods in CMB data analysis. T. Cimic, L. Grigori, J. Papez, R. Stompor.

Upcoming presentations

- *Sparse Days Meeting @ CERFACS, Toulouse, July 11-12, 2019.*

  Presentation: Title: TBA, *Florent Lopez,* STFC, UK.

- *International Congress on Industrial and Applied Mathematics (ICIAM), Valencia, Spain, July 15-19, 2019.*

  Presentation: Recent advances in the direct solution of sparse equations on parallel computers. Iain Duff, STFC, UK.

- *The 7th international conference on Numerical Algebra and Scientific Computing (NASC), Nanjing, China, October 19-23.*

  Presentation: Title: TBA, Iain Duff, STFC, UK

2.2. Website (Task 7.2)

The NLAFET website is the main dissemination channel for promotion of the project and functions as the public interface of NLAFET. Hence, it contains basic information about the project, the partners involved, how to contact the coordinator, the main focus of the research, and the publication results of the project thus far. It serves as a focal point for everyone interested in the objectives, results, impact, and progress of the project.

The website is available at www.nlafet.eu and has seven main sections:

- HOME
- ABOUT; an overview of the background and main aims of the project
- RESEARCH; brief description of the four research-focused work packages WP2-4 and WP6
- USE CASES / APPLICATIONS; describes the work in WP5
- PUBLICATIONS; listings of deliverables and publications
- SOFTWARE
- NEWS

Below, we display a screenshot of the home page of the NLAFET public website. The top banner shows the northern lights above part of the Umeå University campus.

Figure 1: The NLAFET web site

In line with our expectations, the NLAFET website has received considerably more attention during the second period of the project. From June 2017 to end of April 2019, there have been over 158 000 visits from 52 736 unique IP addresses. By comparison, the site had visits from around 11 000 ip addresses during the first 18 months of the project. So far, NLAFET deliverables have been downloaded 8 165 times to 1 420 different IP addresses, which is almost a tenfold increase compared to the first project period. We expect these numbers to increase further when the ten final deliverables submitted during 2019 soon will be publicly available.

The site is described in more detail in deliverable D7.2: *Collaborative infrastructure,* that can be downloaded from the site.

2.3. Open source activities (Task 7.3)

A GitHub platform has been established for the software developed and deployed in the NLAFET project. The platform can be found using the link https://github.com/NLAFET/.

GitHub is a web-based application for managing source codes that relies on the version control system Git (https://git-scm.com/). In the NLAFET project the basic structure is to create a repository for each piece of software that we are developing and deploying. However, many of the repositories include several pieces of software that together form a package of routines for a subset of fundamental linear algebra operations considered in the NLAFET project. Typically, such a package is organized in folders or even subfolders to reflect different types of functionality, tools and solvers.

The current public NLAFET library release includes 13 repositories which cover the topics in the list below. Besides source code each repository include documentation that describes the installation, usage and testing of the main software components.:

- Dense matrix factorizations and solvers.
  - BBLAS: Final batched BLAS reference implementation.
  - plasma: Snapshot of PLASMA library routines for dense factorizations and solves.
  - PlaStar: Some StarPU implementations of dense factorizations and solves.
- Solvers and tools for standard and generalized dense eigenvalue problems.
  - SEVP-PDHSEQR-Alg953: Nonsymmetric standard eigenvalue problems—distributed memory package for computing a real standard Schur form $S = Q^T AQ$.
  - GEVP-PDHGEQZ: Nonsymmetric generalized eigenvalue problems—distributed memory package for computing a real generalized Schur form $(S, T) = Q^T (A, B)Z$.
  - StarNEig: Task-based library including a full suite of software for solving nonsymmetric standard ($Ax = \lambda x$) and generalized ($Ax = \lambda Bx$) eigenvalue problems, respectively.
- Sparse direct factorizations and solvers.
  - SpLLT: Sparse LLT solver for $A = A^T$, where A is positive definite.
  - SyLVER: Symmetrically structured factorizations.
  - ParSHUM: Sparse LU solver for highly unsymmetric matrices.
  - BC: Hybrid solver based on block Cimmino for square and rectangular systems.

- Communication optimal algorithms for iterative methods.
  - preAlps: Preconditioned iterative methods and enlarged Krylov methods.
- Cross-cutting tools.
  - PCP-runtime: Parallelizing the critical path.
  - ABFT-factor: Tiled factorizations with ABFT.

A description of the content and organization of the repositories can be found in Deliverable D7.8, *Release of the* NLAFET *library.* Together, these repositories constitute the NLAFET library. In several comparisons with existing state-of-the-art library software, NLAFET shows outstanding results in terms of performance, scalability and accuracy. Such comparisons and results concerning all NLAFET repositories are presented in several NLAFET deliverable reports and NLAFET Working Notes, all available via the NLAFET web-site.

Altogether, the software components of the NLAFET library release for solving fundamental and important numerical linear algebra problems provide novel task-based algorithms using various programming environments (MPI, OpenMP, ParSEC and StarPU) and contribute to the development of parallel numerical linear algebra for future extreme scale systems. For the future, we plan to update and integrate new progress concerning software components in the NLAFET library.

2.4. Community outreach (Task 7.4)

The outreach activities reported here include standardization efforts, other Horizon 2020 and international collaborations, and finally honours and awards appointed to members of the NLAFET Team during the reporting period.

Batched BLAS standardization

As part of our aim to standardise an API for the proposed Batched BLAS (BBLAS) we have continued to participate in relevant events, primarily at conferences related to high performance computing. These have included Birds-of-a-Feather sessions and minisymposia dedicated to the BBLAS, as well as NLAFET minisymposia. Naturally, the aim has been to involve the community in the proposal, to give and receive feedback and to achieve a consensus. The conferences have been international conferences in the USA, Japan and Europe.

NLAFET deliverable D7.6 describes the BBLAS specification and a reference repository is available on the NLAFET website. Deliverable D7.6 is currently being turned into a paper suitable for submission to the ACM Transactions on Mathematical Software, which is the journal that published the specifications of the BLAS themselves..

Other dissemination or outreach activities

Laura Grigori, INRIA, is co-organizing the Gene Golub Summer school on High Performance Data Analytics, June 17-28, 2019 in Aussois, France. Co-organizerers are M. Knepley, O. Schenk, R. Vuduc. The school is sponsored by SIAM through an endowment from the estate of Gene Golub. The series of lectures will develop background on methods and algorithms for data analytics, approximation algorithms to deal with large volumes of data, languages

and tools for implementing those algorithms on large scale computers, and data-driven applications from scientific computing and machine learning. The material presented in some of the lectures is related to the work in NLAFET on low rank matrix approximation. For more information see: https://project.inria.fr/siamsummerschool/

Collaboration with other projects, programmes, working groups, initiatives etc

On May 15-19, 2017, EXDCI (European eXtreme Data and Computing Initiative) and some FETHPC projects arranged workshops, included in the HPC Summit Week in Barcelona. The overall aim was to find synergies among all stakeholders in the European HPC Ecosystem. Bo Kågström and Lennart Edblom participated and represented the NLAFET project.

One year later, May 28-June 1, 2018, the European HPC Summit Week was arranged in Ljubljana, Slovenia. Similar to previous years the event offered a wide variety of workshops covering a number of application areas where supercomputers are key, as well as HPC technologies and infrastructures. The event is a great opportunity to network with all relevant European HPC stakeholders, from technology suppliers and HPC infrastructures to scientific and industrial HPC users in Europe. NLAFET was represented by Lennart Edblom.

On October 30-31, the PRACE-CoEs-FETHPC-EXDCI workshop was arranged in Brühl, Germany. The goals of this workshop were:

- To better support the interaction and collaboration between PRACE, CoEs and FET-HPC projects
- To better collaborate at the dissemination level between PRACE, CoEs, FET-HPC projects and EXDCI
- To coordinate the training efforts with the CoEs - To better align the operational services for the HPC ecosystem
- To give an opportunity to become familiar with the technologies coming out of the European Exascale projects.In this sense, the workshop should act as a venue for match-making between the PRACE + CoEs + FET-HPC projects + EXDCI experts.

Aspects of the NLAFET project results was presented by Bo Kågström, Iain Duff and Lennart Edblom.

Iain Duff is a member of the Project Steering Board representing CERFACS for the Centre of Excellence EoCoE-II that is managed by the Maison de la Simulation in Paris. He is particularly involved in Work Package 3 that is concerned with large scale simulation codes and their implementation on extreme scale computers. The codes are from energy applications in the areas of meteorology, fusion, materials, wind, and water. He is collaborating with colleagues in EoCoE on the use of solvers in these packages. In particular, he is working with researchers in Toulouse on sparse direct methods and block iterative methods.

The cooperation with the ExaHype FETHPC project has included exchange visits with UMU and UNIMAN, and organization of BoF and Focus sessions at ICS 2018; one topic was on Future Accelerated Math Library Design.

Close relationship has been maintained with INTERTWinE, another FETHPC project in which UNIMAN is a partner. INTERTWinE focuses on interoperability of programming models suitable for exascale. A particular focus for interaction with INTERTWinE was the porting of the PLASMA library from its own runtime system QUARK to the OpenMP runtime system, on

which both projects participated.

NLAFET team members have also met and had discussions with the three Horizon 2020 FETHPC projects that focus on auto-tuning, AllScale, ANTAREX, and READEX, especially at EXDCI events.

ETP4HPC has compiled a series of handbooks, describing European Exascale projects and their international collaboration, that has been presented and distributed at the SC conferences and other occasions. NLAFET has contributed with information about the project.

NLAFET has also been active in the preparatory work for the Extreme Scale Demonstrator (ESD) call, actively looking for cooperation opportunities. However, so far no concrete results have been reached; in addition, the ESD call has been delayed.

NLAFET has contributed with answers to the EXDCI questionnaire about FET-HPC projects.

NLAFET has also collaborated with eSSENCE, the strategic Collaborative Research Programme in e-Science between three Swedish universities with a strong tradition of excellent e-Science research: Uppsala University, Lund University and Umeå University. eSSENCE is financed via a governmental initiative to promote high-quality research in areas of strategic societal and industrial importance. eSSENCE research focuses on four main areas: Materials Science, Human Function and Environment, Life Science, and Generic e-Science Methods and Tools, where the latter main area includes topics of collaboration with the NLAFET UMU team.

Members of the NLAFET INRIA team participate in an interdisciplinary project B3DCMB ("Big Bang from Big Data") funded by the French National Research Agency, which aims at addressing some major challenges in CMB data analysis.

Awards and honours

Jack Dongarra, UNIMAN has been elected as a Foreign Fellow of the Royal Society of the United Kingdom.

Laura Grigori, INRIA, is one of the PIs for the project *Extreme-scale Mathematically-based Computational Chemistry* (EMC2) project, that was awarded an ERC Synergy Grant in the 2018 call for projects.