NLAFET

Parallel Numerical Linear Algebra for Future Extreme-Scale Systems

H2020-FETHPC-2014: GA 671633

D7.1
Dissemination and Community Outreach Plan

January 2016

Document information

Scheduled delivery	2016-01-31
Actual delivery	2016-01-26
Version	Final
Responsible partner	UMU

Dissemination level

PU - Public

Revision history

Date	Editor	Status	Version	Changes
2016-01-05	Lennart Edblom	Draft	0.1, 0.6-0.7	Initial version in docx format Revisions - comments
2016-01-13	Lars Karlsson	Draft	0.2-0.4	Google Docs with comments
2016-01-22	Bo Kågström	Draft	0.5, 0.6-0.8	Revisions – partner comments
2016-01-26	Bo Kågström	Final	0.9	Final revision

Authors

Bo Kågström, Lennart Edblom, and Lars Karlsson, UMU

Internal reviewers

Jack Dongarra, UNIMAN; Iain Duff, STFC; Laura Grigori, INRIA

Contributors

In addition to the authors and reviewers, the following team members have contributed to the content: Jonathan Hogg, STFC; Carl Christian Mikkelsen, UMU, Nick Higham, UNIMAN.

Copyright

Acknowledgements

The NLAFET 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. Executive summary 4

2. Introduction 4

2.1. Activities and objectives of NLAFET 5

2.2. Expected impacts of NLAFET 6

3. Objectives 7

4. Audiences 7

4.1. Scientific community 7

4.2. Hardware and software vendors 8

4.3 Application experts 8

4.4 General computing community 8

5 Strategy 8

6 Activities and communication channels 9

6.1 Website 9

6.2 Press releases 9

6.3 Promotional material 10

6.4 Public deliverables and prototypes 10

6.5 Scientific conferences and publications 10

6.6 Workshops with researchers and application experts 11

6.7 Open source library software 12

6.7.1 Source code repository 12

6.7.2 Tutorials 12

6.7.3 Email assistance 12

6.8 Educational efforts and opportunities for young researchers 13

6.9 Standardisation contributions 13

6.10 Collaboration with industry 13

6.11 Collaboration with other EU projects 13

1. Executive summary

This document describes the objectives, strategy and activities that will be deployed in our work with dissemination and outreach activities in the NLAFET project.

NLAFET is the acronym for the complete project title “Parallel Numerical Linear Algebra for Future Extreme Scale Systems”, and is used throughout this document.

The NLAFET project is a direct response to the demand for new mathematical and algorithmic approaches that will make it possible to use the peak capabilities of the hardware of today’s and future Extreme-Scale High Performance Computing (HPC) systems. Section 2 starts by briefly describing this background. To accomplish this, NLAFET will explore new algorithms and investigate advanced scheduling strategies and auto-tuning techniques.

The main impact of NLAFET will be to develop, deploy and make available software to the HPC community. To achieve this a focused dissemination effort is needed. Sections 3 to 5 describe the detailed objectives for this effort, the target audiences and the strategy that will be employed.

NLAFET will use many different communication channels and activities. They include the project website, presentations at major conferences, publications in scientific journals, and distributing promotional material. In addition, many activities will revolve around the developed software, like workshops and tutorials. We will also collaborate with industry and other EU projects, as well as contribute to standardisation efforts. Section 6 describes these activities in detail.

2. Introduction

Today’s largest HPC systems already have a serious gap between the peak capabilities of the hardware and the performance realized by high-performance computing applications. The ongoing development of extreme-scale systems involves increased parallelism at several levels, and extended heterogeneity in multiple dimensions—multi-core CPUs and many-core accelerators, shared and distributed memory, PCI Express and other novel interconnects. Therefore extreme-scale systems present new challenges that could widen the gap so much that it prevents the productive use of future Exascale computers. However, there is considerable potential to improve software design in order to better exploit the hardware performance features already present in the Petascale systems of today and the systems that are expected to exist in the next five years. The NLAFET project is an important step to address these challenges and will contribute to close the “application-architecture performance gap” in the area of linear algebra with its multitude of applications.

The ultimate goal of all dissemination and outreach activities is to support and strengthen the objectives of the NLAFET project and its strategic impact. The objectives and strategic impacts are specified in the Description of Action document (Annex 1 of the Grant Agreement). A brief overview follows.

2.1. Activities and objectives of NLAFET

The NLAFET project, through its research and software deliverables, will lead to the introduction of tools that will simplify the transition to the next generation of extreme-scale computer architectures.

The focus is on the development of novel algorithms and libraries for fundamental numerical linear algebra operations that are at the heart of many calculations in the computational sciences. The main aim is to design, prototype, and distribute a new linear algebra software library capable of scaling to achieve a significant level of the peak performance on extreme-scale systems (i.e., systems capable of 100 Pflop/s and beyond).

The main aim will be accomplished through the following set of broad activities:

- **Exploration of new algorithms**
  The anticipated characteristics of extreme-scale systems demand significant algorithmic innovations on several levels. We will search for and design algorithms that expose as much parallelism as possible, exploit heterogeneity, avoid communication bottlenecks, respond to escalating fault rates, and help meet emerging power constraints.
- **Investigation of advanced scheduling strategies**
  We will explore advanced scheduling strategies and runtime systems focusing on the extreme scale and strong scalability in multi/many-core and hybrid environments.
- **Investigation of advanced auto-tuning strategies**
  Optimizing the performance on an extreme-scale system will require auto-tuning at multiple levels. We will design and evaluate novel strategies and software support for both offline and online auto-tuning.

The specific objectives of the activities above are the following:

- **Dense solvers (WP2)**
  Design algorithms and implement scalable parallel solvers for dense linear systems and eigenvalue problems. The aim is to leverage new results on auto-tuning, communication avoidance, heterogeneity, and scheduling to deliver software capable of running efficiently on future extreme-scale systems.
- **Sparse direct solvers (WP3)**
  Design algorithms and implement scalable parallel direct solvers for sparse linear systems. The aim is to leverage new results on auto-tuning, heterogeneity, and scheduling as well as to develop and apply new theory on communication-avoiding algorithms to deliver software capable of running efficiently on future extreme-scale systems.
- **Sparse iterative solvers (WP4)**
  Design algorithms and implement scalable parallel iterative solvers for sparse linear systems. The aim is to develop and apply new theory on communication-avoiding

and energy-aware algorithms to deliver iterative solvers and preconditioners with improved scalability and energy efficiency.

- **Use cases and applications (WP5)**
  Validate and disseminate results obtained in the project by integrating newly developed software solutions into challenging scientific applications. The results will be further disseminated by arranging workshops, giving tutorials, and involving students.
- **Cross-cutting issues (WP6)**
  Deliver a sustainable set of tools, methods, and software solutions for cross-cutting issues such as scheduling, auto-tuning, and algorithm-based fault tolerance.

Each of the objectives in the itemized list above, corresponds to a NLAFET work package (WP number in parentheses). In addition there are two more work packages, **WP1 - Management and coordination** and **WP7 - Dissemination and community outreach**.

This document presents the initial version of the activity plan for WP7. The work is structured in the following four tasks:

- Dissemination (T7.1): Develop and carry out the concrete dissemination activities of the project.
- Website (T7.2): Develop and maintain the project website over time.
- Open source activities (T7.3): Coordinate the open-source software contributions of NLAFET.
- Community outreach (T7.4): This task will manage the outreach activities of NLAFET, in particular efforts towards future standardization.

In this document we describe different dissemination and outreach activities of which some contribute to several tasks above. Also note that the NLAFET website is described in deliverable D7.2, *Collaborative infrastructure.*

2.2. Expected impacts of NLAFET

NLAFET will provide a software library that leading-edge applications need in order to attain high performance on extreme-scale systems, so we expect that the impact of the project will be very broad.

The main impact of the project will be to develop, deploy and make available software to the scientific community to make it competitive on a world scale, and to contribute to standardization efforts in the area. In addition, there will be benefits for computational science education in offering an extreme-scale set of components; for application developers by giving them a single point of contact for registering their requirements; and for vendors by organizing a community to help them assemble the complete software environment that their systems need for success.

Our research will help characterize and evaluate multi/many-core and heterogeneous architectures in ways that will have broader impact by identifying opportunities for exploiting

these architectures effectively in a variety of contexts. We plan to ensure this broader impact by publishing our techniques in a tutorial manner and providing email assistance in understanding and disseminating our available software. Thus, our proposed pioneering work in this area will help train the next generation of software architects to achieve the promise of high performance with low power that the future and emerging extreme-scale systems offer.

To realize the expected impacts, we need to engage in relevant dissemination and outreach activities. The rest of this document outlines our initial plans in this matter. During the project, these measures will be made more concrete and partly revised as well as being complemented with additional activities.

3. Dissemination objectives

The main objective of our dissemination and community outreach activities is to make the expected impacts of the project become a reality. Specifically, the objectives are:

1. To distribute a new open source numerical linear algebra software library to relevant stakeholders within Horizon 2020 and the HPC communities.
2. To ensure that end users become aware of and start using the library.
3. To become a preferred point of contact for application experts and vendors when it comes to numerical linear algebra library software.
4. To communicate the scientific progress achieved by the NLAFET project. In addition to the obvious scientific merit this would enable us to increase possibilities for future cooperation and to attract talented people to the partner institutions.
5. To raise awareness of the advances made possible by Horizon 2020 and demonstrate the relevance of the project to industries and citizens of Europe.

4. Audiences

Our audience can be classified into four main categories:

- The scientific community,
- the hardware and software vendors,
- the application experts, and
- the general computing community.

4.1. Scientific community

In addition to researchers active in the parallel computing and numerical linear algebra research fields, the scientific community most relevant to the project consists of computational scientists and engineers in various application domains. The validation of algorithms and software developed in NLAFET will be carried out by means of a selection of challenging applications. See the link http://www.nlafet.eu/use-cases/ for more details.

4.2. Hardware and software vendors

Hardware vendors primarily include processor and accelerator manufacturers. Software vendors of interest include companies that develop and distribute large-scale computational software with a numerical linear algebra component. NLAFET is expected to produce results that have impact on the future development of both software and hardware for extreme-scale systems.

4.3 Application experts

The application experts most likely to be interested in our library are those that are already relying upon large-scale matrix computations on supercomputers, representing a broad spectrum of challenging applications in academia and industry. Application specialists, already committed to implement and evaluate results from NLAFET in their own software, represent four complementary real world application domains. See the link http://www.nlafet.eu/use-cases/ for more details.

4.4 General computing community

The general computing community includes local audiences from outside the immediate research constituency. This can be a member of another non-related EU project, but also journalists, students, or technical professionals not directly connected to NLAFET. But the general public is also meant in the broader sense mainly for evoking possible interest in the general aspects of the NLAFET project as an EU research project, but specifically because we are all more and more dependent on the availability of high-performance computations.

5 Strategy

To reach the full set of objectives, we adopt the following strategy.

The NLAFET library will be distributed as open source software and promoted through a website. The software will be made easy to find, download, install, and use. This *meets* the *first objective*.

By targeting the dissemination of results to potential end users, they will become aware of the existence of a useful software library. By actively engaging with end users through workshops and validation case studies, they will get help to get started with using the library. Tutorials distributed through the website further enable end users to access and use the library on their own, and through email and web assistance they can then obtain the support they may need. This *meets* the *second objective*.

Application experts and vendors will over time come to see NLAFET as a preferred point of contact. This requires that the project presents a unified front and has a persistent presence in many different forums, constantly reminding the stakeholders that the project NLAFET is

alive and flourishing. It lets people know how to make contact, and why it is relevant to do so. This *meets* the *third objective*.

The partners in the project will disseminate the scientific results in conferences, workshops and journals. This will show that we are doing interesting and significant work, and will promote interest for future collaboration.

Talented young people can be attracted to choose future studies and do research at the partner institutions both locally by raising awareness through education efforts and local media attention. Altogether, this *meets* the *fourth objective*.

Awareness of the advances made possible by Horizon 2020 funding can be raised through the website and through press releases. In addition, local awareness within the partner institutions and their local communities can be arranged through local education efforts, through internal promotion within the institutions, and through local media attention. This *meets* the *fifth objective*.

The activities and communication channels indicated in our strategy are described in more detail in the sections that follow.

6 Activities and communication channels

The actual activities and channels to be used in the dissemination, communication and outreach of NLAFET, are briefly described in the following subsections.

6.1 Website

The NLAFET website is considered the main communication vehicle to keep in touch with all target audiences. It serves as a focal point for everyone interested in the objectives, results, impact, and progress of the project.

The site will be used to gather all published materials, e.g., information accessible to laypeople, press releases, public deliverables, scientific publications, tutorials, software, etc.

In order to reach the general audience the website will include information about NLAFET in a non-technical language, making it easier to understand what the project is about.

The site is available at http://www.nlafet.eu and is described in more detail in deliverable D7.2: “Collaborative infrastructure” that can be downloaded from the site.

6.2 Press releases

The publication of press releases is an efficient way to reach a broader audience and will be issued by all the partners in the consortium in their respective countries and organizations.

Press releases will be published coinciding with major milestones of NLAFET, thereby drawing attention to the project and raising awareness of its existence, objectives, results, and impacts.

In connection with the NLAFET kick-off meeting at Umeå University, a press release that attracted considerable attention was released. Some resulting publicity in newspapers and magazines is displayed on the link http://www.nlafet.eu/on-the-web/ .

6.3 Promotional material

Promotion materials are a great way of giving out information about the project for personal contact with interested people, e.g., at conferences, fairs or workshops. Initially, we will produce the following materials:

- **Flyer / Factsheet.** The first version will be a fact sheet, containing an overall description of NLAFET, an explanation what makes the project unique, and project details in relation to the EU funding, as well as contact details. With the progress of the project different versions of the flyer, containing new information about the project will be developed.
- **Poster.** A first version of an NLAFET poster will be designed for general dissemination purposes. As opportunities for conference participations arise, new opportunity-driven versions including specific content targeting a concrete event or audience will be designed.
- **Social media.** Today social media offers a broad variety of tools and networks for connecting to people. This might be particularly helpful in building up a topic specific community. But not all channels are suitable for a research project such as NLAFET. It is necessary to choose between the social networks to reach the right people and use our resources in a meaningful way.

  When the project has reached a certain momentum and has results to publish we will become more active in social networks. Primarily, we plan to use Twitter and LinkedIn to promote the project. This will allow us to attract a specific professional audience.

6.4 Public deliverables and prototypes

All public deliverables will be published directly on the website as soon as they are cleared for publication. The repository is available online at http://www.nlafet.eu/public-deliverables.

6.5 Scientific conferences and publications

An important vehicle for raising awareness, establishing new connections, and disseminating results is the publication of scientific articles and contributions to scientific conferences. Scientific publications will be published with “green” and/or “gold” open access options whenever possible to disseminate the work as widely as possible. All scientific publications are further collected and published on the website at http://www.nlafet.eu/publications .

Relevant scientific conferences during 2016 we plan to participate in include

- SIAM Conference on Parallel Processing for Scientific Computing (next meeting in Paris, April 12-15, 2016; Laura Grigori is co-chair, we organize an invited minisymposia, and Iain Duff will give a plenary talk)
- ISC High Performance 2016
- SC2016
- Europar 2016
- PMAA16 – Parallel Matrix Algorithms and Applications
- Sparse Days at CERFACS
- Vecpar

Relevant scientific journals we will attempt to publish in include

- ACM Transactions on Mathematical Software (http://www.nlafet.eu/research/)
- SIAM Journal on Scientific Computing (https://www.siam.org/journals/sisc.php)
- Parallel Computing (http://www.journals.elsevier.com/parallel-computing/)
- International Journal for High Performance Computing Applications
- Concurrency and Computations: Practice and Experience
- Journal of Parallel and Distributed Computing
- Numerical Linear Algebra with Applications
- Numerical Algorithms
- SIAM Journal on Matrix Analysis and Applications)
- Journal of Computational and Applied Mathematics
- BIT Numerical Mathematics

Any publication, at least partially funded by the NLAFET project, will acknowledge NLAFET and the grant under which the project is funded to raise awareness of the project and the impact of EU Horizon 2020 funding.

6.6 Workshops with researchers and application experts

To facilitate the adoption of our results by the scientific community and application experts, we will arrange hands-on workshops. By bringing application experts, numerical mathematicians, library developers, and computer scientists together to discuss challenges, we can show how the progress and results apply to current problems and discover new directions for our research and development.

The goals of the workshops will be threefold:

1. To familiarize technology consumers with advanced concepts and tools emerging from the project and the research community,
2. to solicit feedback on the effectiveness of current and previous generations of numerical linear algebra libraries on extreme-scale systems, and
3. to build a consensus vision on future research directions that will provide the most benefit to end users in the near, medium, and long terms. This vision will, in turn,

inform and influence the research and infrastructure development, both within NLAFET and in the broader research community.

Additionally in **WP5 - Challenging applications - a selection**, we will arrange close collaborations with a few pre-selected application experts that will help us validate and shape our research and in return receive support to integrate our library into their applications. See the link http://www.nlafet.eu/use-cases/ for more details.

6.7 Open source library software

The final goal of the NLAFET project is to deliver a sustainable library containing initial versions of the next generation of extreme-scale solvers for fundamental numerical linear algebra problems. As an open-source project, there is potential to greatly benefit from community contributions. We will support this openness with the following measures.

6.7.1 Source code repository

To maximize the spread and rate of adoption, the code will be made available with a permissive open source licence (a “modified BSD” licence). This kind of licence allows proprietary commercial use and for the software released under the licence to be incorporated into proprietary commercial products. Works based on the material may be released under a proprietary licence as long as the original copyright statements are maintained. This will enable us to drive home our impact by working with hardware vendors (such as Intel and IBM), library developers (such as HSL, LAPACK, NAG), and end-user software vendors (such as The MathWorks, Julia, and R). The source code repository will be publicly available for read-only access and will be linked to from the website, and will be easy to install and use.

6.7.2 Tutorials

Adopting a complicated research-level software library is a risk and a challenge. Step-by-step tutorials for installing, configuring, and using the library can significantly lower the threshold for adoption. Tutorials will be made available through the website and used in local educational efforts and workshops with application experts.

6.7.3 Email assistance

Problems will inevitably occur while trying to adopt our library. Even with tutorials and good documentation, unexpected problems will appear. We will offer a single point of contact through web and email for assistance and support in using our software.

6.8 Educational efforts and opportunities for young researchers

At the partner institutions, the project will have an educational impact by involving both undergraduate and graduate students. This situation presents excellent opportunities for student interaction with postdoctoral and professional research staff, as well as with colleagues in academic, government, and industry research labs. Furthermore, the principal investigators have taught and will continue to teach special topics courses and workshops in their areas of research.

More broadly, we plan to contribute with tutorials at conferences and industry workshops. Some partners have previously taught performance and architecture related tutorials at various international conferences and user group meetings, and we expect performance engineering for emerging hybrid and multi/many-core architectures to be in great demand as a tutorial topic in the future.

6.9 Standardisation contributions

Standardisation activities can be a very effective dissemination path for achieving broad awareness and take-up of project results. Dissemination of NLAFET results through standards bodies in general brings the project higher international recognition and enables collaboration opportunities.

There are however no formal standards bodies in the area of numerical linear algebra. Libraries are dependent upon the formation of community-developed standards. This is essential to ensure the widespread deployment of the software components. We will contribute to standards efforts in this area, in particular regarding a standardisation effort of a batched BLAS interface targeting hybrid architectures.

6.10 Collaboration with industry

Past experience with software we have developed shows that one key to making shared software infrastructures sustainable in that ecosystem is the active engagement and cooperation with its major stakeholders, for example industry. Since the software is made freely available for anyone to use and build upon, including industry, rapid uptake of the results by major industry partners becomes possible.

We also plan to contribute to the work in related European Technological Platforms, in particular the ETP4HPC platform, giving NLAFET still another contact forum with industrial groups and various user groups.

6.11 Collaboration with other EU projects

The consortium members are committed to collaborate with different external working groups including other ICT projects. In particular, we intend to liaise and cooperate with

other projects that were selected for funding in the FETHPC call "Towards exascale high performance computing" and started in autumn 2015. In particular, we want to identify and exploit synergies with such projects with the aim of developing and strengthening the FETHPC initiative.

Our involvement will be based on both the value we gain from different collaborations and the potential value NLAFET technology will bring to other projects and communities. It is important that collaboration is based on mutual value with clear outcomes and scope of work defined at the outset.

So far, initial contacts have been made with the projects ExaHyPE and INTERTWInE.