H2020 FETHPC-1-2014

Enabling Exascale Fluid Dynamics Simulations
Project Number 671571

D4.3 – Exploitation Plan Revision 2

(also contains:
D4.5 – Innovation Management Report Revision,
D4.8 – Dissemination and Communication – Activities and Results – Update 2,
D4.11 – Collaboration Report – Update 2)

WP4: Dissemination and Exploitation

Document Information

Deliverable Number	D4.3, D4.5, D4.8, D4.11
Deliverable Name	Final Report on Dissemination, Exploitation, Collaboration, and Innovation Management
Due Date	30/09/2018 (PM36)
Deliverable Lead	USTUTT
Authors	Lena Bühler, USTUTT
Responsible Author	Lena Bühler, USTUTT, buehler@hlrs.de
Keywords	Dissemination, Exploitation, Innovation management, Collaboration, Standardisation
WP	WP4
Nature	R
Dissemination Level	PU
Final Version Date	30/09/2018
Reviewed by	Erwin Laure, Niclas Jansson, Nick Johnson
MGT Board Approval	30/09/2018

Document History

Partner	Date	Comments	Version
USTUTT	29.08.2018	First draft	0.1
All	31.08.2018	Inputs to tables	0.2
KTH	1.9.2018	Received for document review	0.7
USTUTT	21.9.2018	Update after review	0.9
KTH	30.09.2018	Final version	1.0

List of Tables

List of Figures

Executive Summary

This deliverable is a merge of several deliverables of WP4 (D4.3, D4.5, D4.8, D4.11). It contains updated versions of the Exploitation Plan (D4.2), the plan and report on innovation management (D4.4), the dissemination and communication plan (D4.7) and the collaboration report (D4.10). It summarises the WP4 activities in all these areas with a special focus on the past 18 months.

Some highlights of this deliverable are the various presentations of the project at conferences and events, such as David Moxey's Keynote talk at ICOSAHOM 2018, and the highly engaging usage of social media during the project lifetime.

Further, the document contains a detailed calendar of past events and describes progress achieved with respect to Exploitation results, IPR and Technology Readiness Levels. The document also includes a brief report on the market outlook in HPC technology and CFD software. Finally, this deliverable outlines the intentions of ExaFLOW with regard to standardisation activities.

Table of Contents

1 Introduction

The aim of this document is to summarize all of the activities of WP4 that took place during the project runtime and evaluate how the goals of exploiting project results were met, how effective the dissemination process was, as well as assessing the collaborations with other projects and organisations to secure the sustainability and continuity of the research in scaling CFD towards exascale.

This document underwent continuous updates until the end of the project and was therefore considered a living document. The aim of this, apart from more accurate planning, was to ensure that changes in the ambitions of ExaFLOW and the engagement opportunities found may be documented in a coherent manner throughout the project.

The engagement described consists of a number of routes, such as through collaboration with the EC and other EU projects, and designed to foster co-operative working and results sharing. In addition to the formal groups we recognise that there will be a number of opportunities for bi-lateral interactions between ExaFLOW and other activities. This document reports on these engagements.

Alongside the collaboration with other H2020-projects we acknowledged and pursued collaboration with other projects, funded by other-than-the-EC funding schemes, i.e. local and national governments and international sources. Again this document reports on these engagements.

Finally, ExaFLOW contributed to standardisation processes where applicable. Through partnerships with other projects, ExaFLOW connected with relevant stakeholders such as technology platforms, standardisation bodies and governance stakeholders. These partnerships and strategic actions positioned ExaFLOW at an important juncture between research, infrastructure building and impact generation.

2 Dissemination (D4.8)

2.1 Goals and tools of dissemination

In the ExaFLOW project dissemination and communication have been managed concurrently. This is because the two are aligned around strategic timing and targets, leverage social media and are monitored for success with KPIs. Moreover, cross-fertilization of dissemination and communication is important and cost-effective (e.g. blogging or tweeting in parallel to publishing or presenting results) because although the means and audiences are distinct, the overall impact can be multiplied.

An overall dissemination strategy was drawn-up in the Initial Dissemination Plan in deliverable D4.6. We stated by identifying target groups and the most appropriate means to target them. We also defined the level of interaction that is necessary through time from awareness over understanding up to action/collaboration. Dissemination is of similar importance for the academic/scientific communities as well as for the industrial partners. For the latter group the goal is that dissemination activities transition into collaboration and finally exploitation activities towards the end of the project.

ExaFLOW complies with guidelines outlined in the "The EU guide to communication" published by the Research and Innovation Directorate-General of the European Commission. In the paragraphs below we report on Dissemination KPIs and show progress achieved using each dissemination tool separately.

2.2 Progress with dissemination KPIs

To monitor the dissemination progress of ExaFLOW, a number of Key Performance Indicators (KPIs) have been identified for the project (Table 1). The EB monitors the progress of the project against these KPIs on a monthly basis and reports on them in the project reports.

Table 1 - Dissemination KPIs for the project

KPI	Target	Progress M1 – M36
Journal papers (see section 2.8)	8	15
Conference proceedings (see section 2.8)	20	21
Conference presentations (see section 2.10)	25	30
Whitepapers (see section 2.8)	2	2
Press releases (see section 2.6)	2	2
Presence at events (see section 2.10)	2	19
Trainings (see section 2.9)	1	1

Website visits (see section 2.3)	3000 p.a., 40% spend > 2 min on site	average of 10.000 p.a.
Social media (see section 2.5)	Twitter account with bi-weekly updates	Twitter account with an average of 3 Tweets per week and special campaigns for the blog posts
Blog entries (see section 2.4)	15	32

2.3 Website

Within the first 18 months of the project runtime, the project website has been updated with short scientific papers describing the use cases, still making these pages the most popular pages of the website (see figure below).

Figure 1 - Most popular pages on the ExaFLOW website – Status 20.08.2018

Following the request of the invited reviewers in the PM 18 hands on meeting in Lausanne, the project website was modernized taking the measures described and depicted below:

1. Adapting the colour palette of the website to the colours of the ExaFLOW Logo
2. Including a more user-friendly menu navigation moving the menu bar from the right to the top and including a transparent drop-down menu structure where needed
3. Inserting a banner of photos leading to different areas of the website

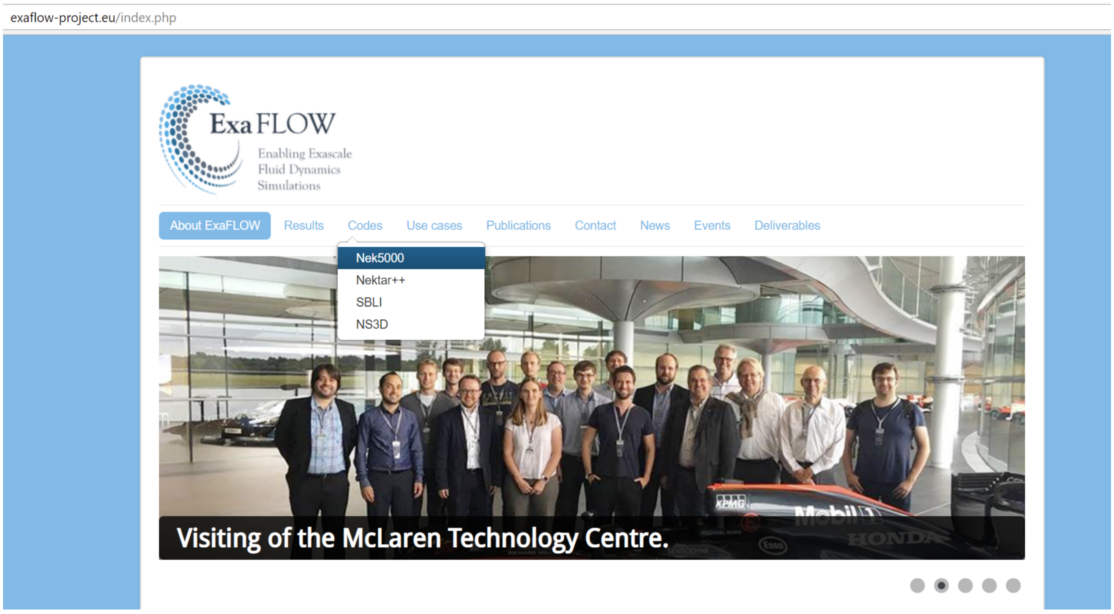

Figure 2 - Website modernization measures according to points 1, 2, and 3

4. Including a moving banner on the bottom of the homepage presenting all project partners

Figure 3 - Website modernization measures according to point 4

5. Adding an Event calendar, a list of publications and preliminary results as items to the menu, which has been updated on a regular basis

For extra traffic to the website, blog posts and events have been promoted via the @ExaFLOWproject Twitter channel.

2.4 Blogging

Within the first year of the project, irregular updates were posted online focusing on Events, Workshops, and Meetings held in connection with the ExaFLOW project.

After the all-hands meeting in London in September 2016, we decided to start a series of blog publications on topics relevant to ExaFLOW work, as well as on topics of major interest to specific partners. Around the time of the last all-hands meeting in Southampton in May 2018, the KPI on blogposts was met and has now been exceeded. The blogging strategy shifted again after the meeting: previously, all technical areas were covered by partners in turn, the focus was now on the flagship runs in particular, and how impact was created for both science and industry, leading to a lower frequency of published posts in the last few months of the project.

Table 2 - Blogging activities during the project runtime

Date	Title	Category
12.11.2015	ExaFLOW Kick-off Meeting	Project News
04.04.2016	Simulation Driven Design for Computational Fluid Dynamics Workshop	Event
04.07.2016	ExaFLOW at the ISC 2016 in Frankfurt	Event
28.09.2016	Third all-hands project meeting in London	Project News
17.10.2016	ExaFLOW Resilience @ HPDC'16 Symposium	Event Science Update
02.11.2016	Nektar++ Aorta Test Case on ARCHER: Improving I/O Efficiency for ExaFLOW Use Cases	Science Update
15.11.2016	OpenSBLI codegen framework for modelling with finite difference methods	Science Update
15.11.2016	Implementation of h-type refinement in Nek5000	Science Update
07.12.2016	ExaFLOW at the 8th annual general meeting of asc(s	Event
14.12.2016	Weak Dirichlet Boundary Conditions and Hybrid DG on Groups of Elements	Science Update
19.01.2017	Performance evaluation of explicit finite difference algorithms with varying amounts of computational and memory intensity	Science Update
02.02.2017	Data compression strategies for exascale CFD simulations	Science Update
16.02.2017	Preparation of industrial test cases to benchmark algorithmic improvements in ExaFLOW	Science Update
07.04.2017	How quickly time flies!	Project News
05.05.2017	Apply now and Combine Art with Technology in ExaFLOW Project	Project News
19.05.2017	Fault Tolerance at Exascale: A memory-conservative approach to resilience in CFD tools	Science Update
02.06.2017	Error indicators for finite difference methods using spectral techniques	Science Update
08.06.2017	ExaFLOW mini-symposium at ParCFD	Event
04.07.2017	Empirically determining energy- and runtime-efficient CPU clock frequencies	Science Update
04.07.2017	Comparison between high-fidelity CFD and PIV data for the isolated McLaren Front-wing	Science Update

24.07.2017	Update on h-type AMR implementation in Nek5000	Science Update
01.08.2017	Performance Analysis of an ExaFLOW code with the automotive use case	Science Update
15.08.2017	Multi-level Diskless Checksum Checkpointing for HPC Resilience at ExaScale	Science Update
26.10.2017	ExaFLOW Consortium Sets Agenda for Final Year	Project News
06.11.2017	Heterogeneous modelling: Combining LES and quasi-DNS to simulate turbulent channel flow	Science Update
27.11.2017	What's all this Use Case stuff anyhow?	Project News
08.12.2017	Adjoint error estimators for the spectral element method	Science Update
12.01.2018	Weak Dirichlet Boundary Conditions and Hybrid DG on Groups of Elements	Science Update
31.01.2018	Strong scalability test of an automotive use case performed with Nektar++ and Fluent (DES)	Science Update
22.02.2018	Multi-level Diskless Checksum Checkpointing for HPC Resilience at ExaScale, Part 2	Science Update
21.06.2018	ExaFLOW Flagship runs underway	Science Update
26.09.2018	Paving the Way to Industry-relevant Simulations	Science Update

2.5 Social media

Social media has been used as a tool for three main purposes:

1. To increase traffic to the website
2. To create a community interested in exascale HPC CFD applications
3. To inform the community about participation in events, new journal releases, and to encourage knowledge exchange with a view to strengthening the impact of research conducted by the partners within the scope of the ExaFLOW project

We used twitter campaigns for each new blog post on the website, tweeting about each post for several weeks in a row. This had very positive effects on the website allowing traffic, which we hold partly responsible for surpassing our target for the number of visitors. Moreover, partners helped to create conversations on twitter by posting about new software releases and publications.

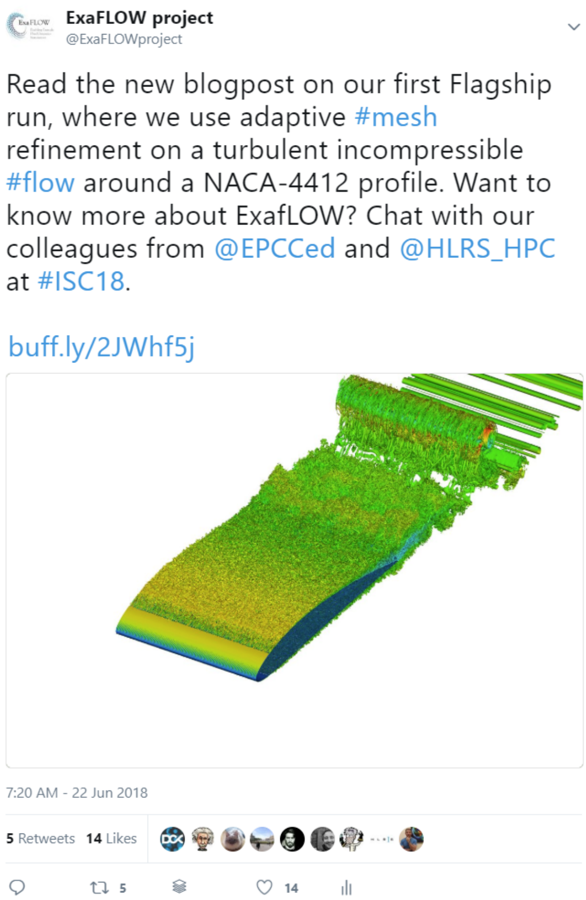

Figure 4 - Example for highly engaging tweet announcing new blog post

For a quick overview on the overall performance of the @ExaFLOWproject Twitter account, the figure below shows the development of daily impressions, followers, and engagement rate throughout the project runtime.

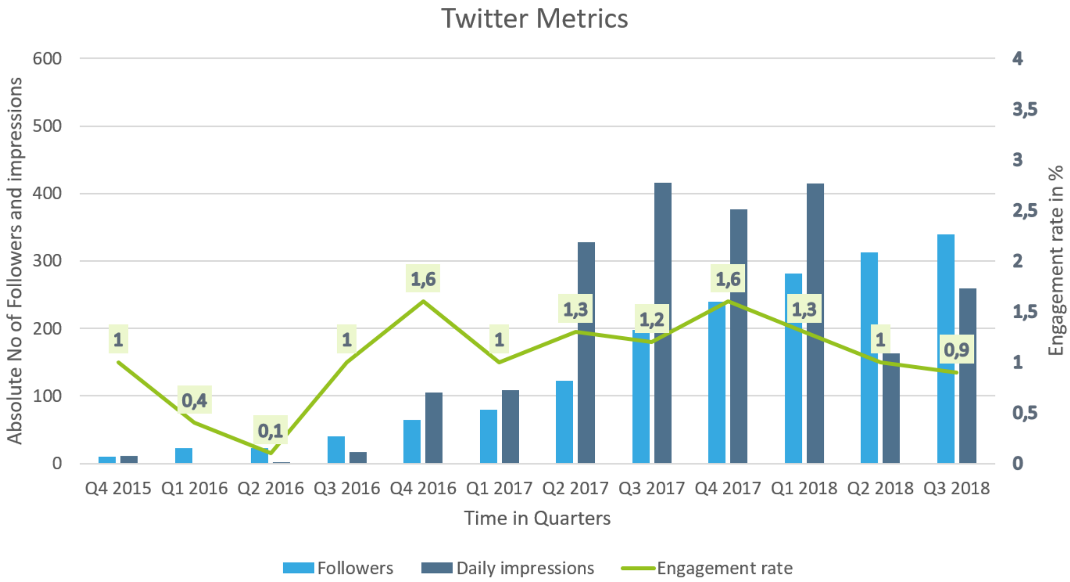

Figure 5 - @ExaFLOWproject performance overview

The left axis shows the number of followers at the end of the respective Quarter and the average number of impressions of one post achieved each day depicted as blue bars. The right axis shows the average interaction rate per Quarter in percent depicted by trend line. Generally, with the number of followers also the number of impressions rises, as more people get to see a post. In contrast, the higher the number of impressions, the harder it gets to maintain high engagement rates, as impressions don't rise proportional to the number of people interacting with a post by liking or sharing. The fact that we managed to keep the interaction rate (mostly) over 1% while impressions rose significantly suggests that the @ExaFLOWproject Twitter account must have provided its followers with content that was of value for them.

This hypothesis was supported by a study on online communication strategies of FET HPC projects[1]. Mazzollo found: "The most influential HPC projects on Twitter were found to be ExaFLOW, EXDCI, Mont-Blanc 3 and NEXTGenIO. In general, HPC influential initiatives are identified by a number of followers and a ratio followers/following larger than ~100 and ~2, respectively."

Noteworthy is that ExaFLOW follows the official Twitter account for the Horizon

[1] Mazzollo, G. (2017), European FET Research Projects on High Performance Computing and Quantum Technologies: Analysis of the Online Communication Strategies. Masterthesis.

2020 programme @EU_H2020 and participates in the community of projects on social media, using the newly created hashtag #ResearchImpactEU when announcing “breaking news”.

2.6 Media Relations

In November 2016 ExaFLOW published their first official press release, timing it strategically just before the SC16 Conference. The main goal of this press release was to announce the presence of the project at the conference and point interested audience to the booths of two ExaFLOW partners: UEDIN and USTUTT. The second official press release was published end of September 2018 with the goal of informing the community about the main achievements of the ExaFLOW project. Potential

Apart from the mentioned press releases, other activities such as reports from events and blog posts have been regularly picked up by the press. All appearances of the project in the press so far are listed below:

- Scientific Computing World, 28 October 2015, Europe spends to strengthen its HPC: https://www.scientific-computing.com/news/analysis-opinion/europe-spends-strengthen-its-hpc
- Inside HPC, 28 October 2015, European Commission Steps Up Funding of HPC: http://insidehpc.com/2015/10/european-commission-steps-up-funding-of-hpc/
- HPCwire, 12 July 2016, ISC Workshop Tackles the Co-development Challenge: https://www.hpcwire.com/2016/07/12/isc-workshop-tackles-co-developments-thorny-challenges/
- Scientific Computing World, 3 August 2016, Exploring energy efficiency: https://www.scientific-computing.com/news/analysis-opinion/exploring-energy-efficiency
- Inside HPC, 5 August 2016, Adept Project Explores HPC Energy Efficiency: http://insidehpc.com/2016/08/exploring-energy-efficiency/
- Inside HPC, 8 November 2016, ExaFLOW Funds CFD Research for Exascale: http://insidehpc.com/2016/11/bringing-europes-cfd-community-one-step-closer-to-exascale/
- HPCwire, 9 November 2016, ExaFLOW Project Secures Funding: https://www.hpcwire.com/off-the-wire/exaflow-project-secures-funding/
- Scientific Computing World, 10 November 2016, Bringing Europe’s CFD community one step closer to exascale: https://www.scientific-computing.com/news/bringing-europe%E2%80%99s-cfd-community-one-step-closer-exascale
- CFD online, 8 November 2016, Bringing Europe’s CFD community one step closer to exascale: https://www.cfd-online.com/Forum/news.cgi/read/101657

- Primeur Magazine, 8 February 2017, ExaFlow on data compression strategies for exascale CFD simulations: http://primeurmagazine.com/flash/AE-PF-02-17-3.html
- Primeur Magazine, 6 July 2017, ExaFLOW blog: Empirically determining energy- and runtime-efficient CPU clock frequencies: http://primeurmagazine.com/flash/AE-PF-07-17-6.html
- Rich Report, 12 July 2017, Interview: Mirren White from EPCC on the importance of communications for HPC: https://www.youtube.com/watch?v=flDp6cQJBQo&t=1s
- Inside HPC, 17 July 2017, Exaflow: Collaborating on the Computational Challenges of CFD: https://insidehpc.com/2017/07/exaflow-collaborating-computational-challenges-cfd/
- Primeur Magazine, 9 August 2017, Performance Analysis of an ExaFLOW code with the automotive use case: http://primeurmagazine.com/flash/AE-PF-08-17-7.html
- Primeur Magazine, 26 October 2017, ExaFLOW exascale project consortium sets agenda for final project year: http://primeurmagazine.com/flash/AE-PF-10-17-22.html

2.7 Printed materials

Posters were used during the presentations of the project at different events and over a hundred flyers were distributed at two major HPC events during 2016: ISC16 and SC16. An updated version of the flyer was released at the beginning of year 3 in October 2017, to showcase first results of the project, and has been distributed again at events such as SC17, ISC18, PASC18, and at visits to other Universities.

2.8 Publications and whitepapers

Below are listed the publications that members of the ExaFLOW project have authored during the project runtime:

Table 3 - Publications to appear (presumably in 2019)

Publication Type	Citation
Journal Paper	N. Offermans, A. Peplinski, O. Marin, P. Schlatter. Adjoint error estimators and adaptive mesh refinement applied to low Reynolds number flows. Computers & Fluids, 2018. To be submitted
Conference Proceedings	A. Peplinski, N. Offermans, P. Fischer & P. Schlatter. Nonconforming elements in Nek5000: Stability and Implementation. In proceedings of Spectral and High Order Methods ICOSAHOM 2018, London. To appear.
	N. Offermans, A. Peplinski, O. Marin, P. Schlatter. Efficient preconditioning for the pressure equation in Nek5000 using the

	Hypre library. In proceedings of Spectral and High Order Methods for Partial Differential Equations ICOSAHOM 2018, London. To appear.
	O. Marin, R. Vinuesa, E. Merzari, P. Schlatter. Data compression effects on coherent turbulence structures. In proceedings of Spectral and High Order Methods ICOSAHOM 2018, London. To appear.
	M. Vymazal, D. Moxey, C. Cantwell, R. M. Kirby & S. Sherwin. On a mixed CG-HDG formulation for high-order simulations. In proceedings of Spectral and High Order Methods ICOSAHOM 2018, London. To appear.
	C. Cantwell & A. Nielsen. A minimally intrusive low-memory approach to resilience for existing transient solvers. In proceedings of Spectral and High Order Methods ICOSAHOM 2018, London. To appear.

Table 4: Publications in 2018

Publication Type	Citation
Journal Papers	N. Jansson, N. Johnson, M. Bareford. Efficient Gather-Scatter Operations in Nek5000 Using PGAS. Journal of Parallel and Distributed Computing, Special Issue on Exascale Applications and Software, 2018. Submitted
	Otero E., Gong J., Min M., Fischer P., Schlatter P., and Laure E., OpenACC acceleratior for Pn-Pn-2 algorithm in Nek5000. Journal of Parallel and Distributed Computing. Submitted
	Eichstädt, J., Green, M., Turner, M., Peiró, J., & Moxey, D. (2018). Accelerating high-order mesh optimisation with an architecture-independent programming model. Computer Physics Communications, 229, 36-53. Doi: 10.1016/j.cpc.2018.03.025
	Nielsen, A. S., Brunner, G., & Hesthaven, J. S. (2018). Communication-aware adaptive parareal with application to a nonlinear hyperbolic system of partial differential equations. Journal of Computational Physics. DOI: https://doi.org/10.1016/j.jcp.2018.04.056
	Moxey, David, Shankar P. Sastry, and Robert M. Kirby. "Interpolation Error Bounds for Curvilinear Finite Elements and Their Implications on Adaptive Mesh Refinement." Journal of Scientific Computing (2018): 1-18. Doi: https://doi.org/10.1007/s10915-018-0795-6
	Jacobs, Christian T., et al. "An error indicator for finite difference methods using spectral techniques with application to aerofoil simulation." Computers & Fluids 168 (2018): 67-72. Doi: https://doi.org/10.1016/j.compfluid.2018.03.065

	Lusher, David J., Satya P. Jammy, and Neil D. Sandham. "Shock-wave/boundary-layer interactions in the automatic source-code generation framework OpenSBLI." Computers & Fluids (2018). Doi: https://doi.org/10.1016/j.compfluid.2018.03.081
	Cantwell, C. D., & Nielsen, A. S. (2018). A Minimally Intrusive Low-Memory Approach to Resilience for Existing Transient Solvers. Journal of Scientific Computing, 1-17. Doi: https://doi.org/10.1007/s10915-018-0778-7
	Otero, E., Vinuesa, R., Marin, O., Laure, E., & Schlatter, P. (2018). Lossy data compression effects on wall-bounded turbulence: bounds on data reduction. Flow, Turbulence and Combustion, 1-23. Doi: https://doi.org/10.1007/s10494-018-9923-5
Conference Proceedings	M. Zauner, C. T. Jacobs, N. D. Sandham (2018). Grid refinement using spectral error indicators with application to airfoil DNS. In Proceedings of the 6th European Conference on Computational Mechanics (ECCM 6) and 7th European Conference on Computational Fluid Dynamics (ECFD 7), 11-15 June 2018, Glasgow, UK. Link: https://bit.ly/2MZnaTy
	S. P. Jammy, C. T. Jacobs, D. J. Lusher, N. D. Sandham (2018). Energy consumption of algorithms for solving the compressible Navier-Stokes equations on CPU's, GPU's and KNL's. In Proceedings of the 6th European Conference on Computational Mechanics (ECCM 6) and 7th European Conference on Computational Fluid Dynamics (ECFD 7), 11-15 June 2018, Glasgow, UK. Link: https://bit.ly/2Oc8GVn
	N. Jansson, N. Johnson and M. Bareford. Efficient Gather-Scatter Operations in Nek5000 Using PGAS. In Proceedings of the 5th International Conference on Exascale Applications and Software, Edinburgh, 2018. Link: http://www.easc2018.ed.ac.uk/conference-programme/
	Otero E., Gong J., Min M., Fischer P., Schlatter P., and Laure E., OpenACC acceleratior for Pn-Pn-2 algorithm in Nek5000, EASC 2018, 17-19 April, Edinburgh, UK. Link: http://www.easc2018.ed.ac.uk/conference-programme/
Whitepaper	Government Office for Science (2018), Modelling in Business and Manufacturing , In: Computational Modelling: Technological Futures, p. 73 – 80, https://bit.ly/2xsvfv9
	J. Peiró, D. Moxey, M. Turner, G. Mengaldo, R. C. Moura, A. Jassim, M. Taylor, S. J. Sherwin. spectral/hp element methods for under-resolved dns: paving the way to industry-relevant simulations. In ERCOFTAC Bulletin 89. To appear

Table 5: Publications in 2017

Publication Type	Citation
Journal Papers	Jacobs, C. T., Jammy, S. P., & Sandham, N. D. (2017). OpenSBLI: A framework for the automated derivation and parallel execution of finite difference solvers on a range of computer architectures. Journal of Computational Science, 18:12-23, doi: http://doi.org/10.1016/j.jocs.2016.11.001
	Sandham, N. D., Johnstone, R., & Jacobs, C. T. (2017). Surface-sampled simulations of turbulent flow at high Reynolds number. International Journal for Numerical Methods in Fluids, 85(9), 525-537. DOI: http://dx.doi.org/10.1002/fld.4395
	Turner, M., Peiró, J., & Moxey, D. (2017). Curvilinear mesh generation using a variational framework. Computer-Aided Design. Doi: 10.1016/j.cad.2017.10.004
	S. P. Jammy, C. T. Jacobs, D. J. Lusher, N. D. Sandham (Submitted). Energy efficiency of finite difference algorithms on multicore CPUs, GPUs, and Intel Xeon Phi processors. Pre-print: https://arxiv.org/abs/1709.09713
Conference Proceedings	Jacobs, Christian, Sandham, Neil and De Tullio, Nicola (2017) An error indicator for finite difference methods using spectral techniques with application to aerofoil simulation At 29th International Conference on Parallel Computational Fluid Dynamics, Glasgow, United Kingdom. 15 - 17 May 2017. 2 pp. https://eprints.soton.ac.uk/id/eprint/407655
	M. Vymazal, D. Moxey, C. Cantwell, S. Sherwin, and M. Kirby (2017). Towards combined cg-dg for elliptic problems. In SIAM 2017. Abstract: https://bit.ly/2PSRCRq
	N. Offermans, A. Peplinski, O. Marin, P. Fischer & P. Schlatter. Towards adaptive mesh refinement for the spectral element solver Nek5000. In Proc. Direct and Large-Eddy Simulation 11 (DLES-11), Pisa, Italy (2017). To appear.
	Otero, E., Marin, O., Vinuesa, R., Schlatter, P., Siegel, A., Laure, E.: The effect of lossy data compression in computational fluid dynamics applications: resilience and data postprocessing, In Proc. Direct and Large-Eddy Simulation 11 (DLES-11), Pisa, Italy (2017). To appear.
	O. Marin, E. Merzari, P. Schlatter, & A. Siegel, (2017). Proper orthogonal decomposition on compressed data. In TSFP 2017. Link: https://bit.ly/2DsuZSi
	M. Turner, D. Moxey, J. Peiró, M. Gammon, C. R. Pollard and H. Bucklow. A framework for the generation of high-order curvilienar hybrid meshes for CFD simulations, in Procedia Engineering, 2017, 203, pp. 206–218. Doi: 10.1016/j.proeng.2017.09.808

Other	Jacobs, C. T., Jammy, S. P., Lusher, D. J., Sandham, N. D. (2017). KNL Performance Comparison: OpenSBLI. ARCHER technical report.
	Offermans, N. (2017). Towards adaptive mesh refinement in Nek5000 (Doctoral dissertation, KTH Royal Institute of Technology).

Table 6: Publications in 2016

Publication Type	Citation
Journal Paper	S. P. Jammy, C. T. Jacobs, and N. D. Sandham, Performance evaluation of explicit finite difference algorithms with varying amounts of computational and memory intensity, J. Comput. Sci. (2016), http://dx.doi.org/10.1016/j.jocs.2016.10.015
Conference Proceedings	Michael Bareford, Nick Johnson, and Michèle Weiland. 2016. On the trade-offs between energy to solution and runtime for real-world CFD test-cases. In Proceedings of the Exascale Applications and Software Conference 2016 (EASC '16). ACM, New York, NY, USA, Article 6, 8 pages. doi: http://dx.doi.org/10.1145/2938615.2938619
	D. Moxey, C. D. Cantwell, G. Mengaldo, D. Serson, D. Ekelschot, J. Peiró, S. J. Sherwin and R. M. Kirby, Towards p-adaptive spectral/hp element methods for modelling industrial flows, International Conference on Spectral and High-Order Methods, 2016. Abstract: https://bit.ly/2puCkH3
	A. S. Nielsen and J. S. Hesthaven, 2016. Fault Tolerance in the Parareal Method. Proceedings of the Fault Tolerance for HPC at eXtreme Scale Workshop at the 25th ACM Symposium on High-Performance Parallel and Distributed Computing. doi: 10.1145/2909428.2909431
	Offermans, N., Marin, O., Schanen, M., Gong, J., Fischer, P., Schlatter, P. On the Strong Scaling of the Spectral Element Solver Nek5000 on Petascale Systems. In Proceedings of the Exascale Applications and Software Conference 2016. doi: https://doi.org/10.1145/2938615.2938617
	S. Sherwin, J.-E. Lombard, D. Moxey, J. P. R. Moura, and G. Mengaldo. Implicit les spectral/hp element modelling of flow past complex geometries related to formula 1, International Conference on Spectral and High-Order Methods, 2016. Abstract: https://bit.ly/2puCkH3
	Peplinski, A., Fischer, P. F., & Schlatter, P. (2016, April). Parallel performance of h-type Adaptive Mesh Refinement for Nek5000. In Proceedings of the Exascale Applications and Software Conference 2016 (p. 4). doi: 10.1145/2938615.2938620

2.9 Training

One training day has been delivered on the use of Nektar++ in the form of a workshop in June 2016 by IMPERIAL. Similarly, a Summer School on numerical methods in CFD (see http://www.flow.kth.se/?q=node/255) was organized by KTH in April 2016, where one day was devoted to Nek5000 and its implementation.

2.10 Dissemination and collaboration events

In the table of collaboration actions partners have participated since the beginning of the project. It is compiled based on partner inputs. Events organized by ExaFLOW, such as Minisymposia at established conferences, are highlighted in green.

Table 7 - Dissemination and collaboration events overview (M1-M36)

Date	Location	Event description
23 September 2015	Rome, Italy	European eXtreme Data and Computing Initiative workshop in Rome for FET projects and Centres of Excellence (CoE) related to HPC
23 February 2016	Leinfelden-Echterdingen, Germany	Key note presentation and poster presentation at the ASCS workshop "Simulation Driven Design for Computational Fluid Dynamics"
8 March 2016	Reading, UK	Presentation at the European Centre for Medium Range Weather Forecasts
16-17 March 2016	Sendai, Japan	Uwe Küster presented ExaFLOW as part of his presentation on "Spectral structures for nonlinear operators on arbitrary compact spaces" at the 23th Workshop on Sustained Simulation Performance (WSSP)
26-29 April 2016	Stockholm, Sweden	Exascale Application & Software Conference EASC 2016: Three talks
9-12 May 2016	Prague, Czech Republic	European HPC Summit Week organised by European Extreme Data & Computing Initiative
12 May 2016	Leinfelden-Echterdingen, Germany	Presentation of the ExaFLOW project at the ASCS general assembly 2016

31 May - 4 June 2016	Kyoto, Japan	Talk on Fault Tolerance in the Parareal Method by EPFL at the 25th international symposium on high performance parallel and distributed computing (HPDC'16)
7-8 June 2016	London, UK	Nektar++ workshop organised by ICL
8-10 June 2016	Lausanne, Switzerland	Platform for Advanced Scientific Computing Conference (PASC'16): EPFL talk titled "Space-Time Parallelism for Hyperbolic PDEs"
19-23 June 2016	Frankfurt, Germany	ISC High Performance 2016
24-28 June 2016	Rio, Brazil	ICOSAHOM: International Conference on Spectral and High Order Methods Methods
21-26 August 2016	Montreal, Canada	ICTAM: The wing test case was presented in the context of the ExaFLOW project by Philipp Schlatter
7 – 9 September 2016	London, UK	Presentation by SOTON about OpenSBLI at the UK Fluids Conference at Imperial College London.
13-18 November 2016	Salt Lake City, USA	SC 2016, International Conference for High Performance Computing, Networking, Storage and Analysis
23 November 2016	Oxford, UK	Seminar presentation by SOTON about OpenSBLI at the University of Oxford.
21-23 November 2016	Toulouse, France	1st TILDA Symposium and Workshop on Industrial LES & DNS., CERFACS. Presentation by SOTON on "High-order simulations of shock-wave/boundary-layer interactions: Current state of the art and software future-proofing"
5-6 December 2016	Stuttgart, Germany	Two contributions "Data compression strategies for exascale CFD simulations" and "Spectral decomposition of nonlinear Trajectories" at 24th Workshop on Sustained Simulation Performance (WSSP)
24 February 2017	Kobe, Japan	Presentation of ExaFLOW at the 7th AICS International Symposium
5-6 April 2017	Hanau, Germany	Poster presentation at the Automotive CAE Grand Challenge 2017

15-19 May 2017	Barcelona, Spain	European HPC Summit Week 2017
15-17 May 2017	Glasgow, UK	ParCFD 2017: Full workshop "Towards Exascale in High-Order Computational Fluid Dynamics". Presentation project work on error indicators and OpenSBLI developments by SOTON
29-31 May 2017	Pisa, Italy	Two talks at ERCOFTAC Workshop Direct and Large Eddy Simulation DLES11 on adaptive mesh refinement and lossy data compression
01 June 2017	Leinfelden-Echterdingen, Germany	Presentation of the ExaFLOW project at the ASCS general assembly 2017
18-22 June 2017	Frankfurt, Germany	ISC High Performance 2017: Full workshop "Interdisciplinary Challenges Towards Exascale Fluid Dynamics"
26-28 June 2017	Lugano, Switzerland	Minisymposium: "Enabling Exascale Fluid Dynamic Simulations" at PASC 17 Conference
10-14 July 2017	Pittsburgh, USA	Participation in a minisymposium on exascale at the SIAM Annual Meeting
30 July – 3 August 2017	Waikoloa, Hawaii	ASME FEDSM: Panel and Keynote on Exascale Methods in industrial CFD
4 – 5 September 2017	London, UK	Presentation at UK Turbulence Consortium Annual Review 2017
18 – 21 September 2017	Barcelona, Spain	Talk on High-Order Curvilinear Hybrid Meshes at 26th International Meshing Roundtable and User Forum
10 – 11 October 2017	Stuttgart, Germany	Presentation on OpenSBLI developments by SOTON at HLRS Workshop on Sustained Simulation Performance
12 - 17 November	Denver, USA	SC 2017, Supercomputing Conference
6 December 2017	Tsukuba, Japan	EPPC visit of Center for Computational Sciences, University of Tsukuba
6 March 2018	Kobe, Japan	Niclas Jansson at 133rd AICS Cafe on Efficient Gather-Scatter Operations in NEK5000 using PGAS / Philipp Schlatter at 133rd AICS Cafe on Enabling Exascale Fluid Dynamics: Adaptive Mesh Refinement

9 March 2018	Tokyo, Japan	Minisymposium: "Approaches Towards Exascale Computational Fluid Dynamics" at SIAM Conference on Parallel Processing for Scientific Computing
17-19 April 2018	Edinburgh, Scotland	Three talks on PGAS, Nek5000, and data compression algorithms at the Exascale Applications and Software Conference (EASC)
11-15 June 2018	Glasgow, UK	Two contributions to Minisymposia on compressible Navier-Stokes equations and spectral error indicators
25 - 28 June 2018	Frankfurt, Germany	HLRS and EPCC with booths at ISC 2018
2 July 2018	Basel, Switzerland	Minisymposium: "On the Road to Exascale Computing" at PASC 2018 Confernce
11 July 2018	London, UK	Minisymposium: "High-Order Spectral Elements in Computational Fluid Dynamics: The ExaFLOW" Project at ICOSAHOM 2018

3 Collaboration activities (D4.11)

3.1 Collaboration on European and national levels

3.1.1 Collaboration with EU projects

In this section we provide a list of EC funded research projects (FP7 and H2020 program), whose outcomes have been identified as potentially valuable for ExaFLOW in terms of knowledge transfer, as well as the exploitation of specific components or approaches.

As the FP7 projects CRESTA, EPiGRAM, and IDIHOM were running when ExaFLOW started, development in these projects could be more directly exploited. The other four H2020 projects started running at the same time as ExaFLOW, and as such the collaboration was more of a dialogic nature, exchanging experience on overlapping issues such as fault tolerance and I/O.
The consortium intends to contact the groups/alliances/communities formed after the finalisation of each project (if any) in order to contribute to the reuse and sustainability of results.

Table 8 - Results of collaboration with EU-projects

Project	Research area	Collaboration results/Form of collaboration	Partner
CRESTA	Co-design	Initial work on AMR and exploitation of GPUs in Nek5000.	KTH, USTUTT, UEDIN

EPiGRAM	Programming Models	Efficient usage of MPI, hybrid models, experiences with PGAS approaches served as basis for ExaGS. Work on communication kernel optimization to be included in Nek5000.	KTH, UEDIN
IDIHOM	Industrial application of high order methods	Parallelisation for explicit DG methods and mesh generation	IMPERIAL
Intertwine	Programming Models	Knowledge transfer regarding efficient use of hybrid programming models went into ExaGS	KTH, UEDIN
NextGenIO	Memory & Storage	Knowledge transfer on I/O and fault tolerance work	UEDIN
SAGE	Memory & Storage	Knowledge transfer on I/O and fault tolerance work	KTH
ESCAPE	Accelerators	Application of Accelerator technology to explicit DG methods	IMPERIAL

3.1.2 Collaboration with EU technology platforms

Several members of the consortium are already involved in the European Technology Platform for High Performance Computing (ETP4HPC) technology platform, helping shape its Research Agenda. These partners constitute the channel to advocate for ExaFLOW's interests. The objective of the strategic research agenda (SRA) of ETP4HPC is to outline a roadmap for the implementation of a research programme aiming at the development of European HPC technologies. The approach followed to structure the SRA has been to combine ETP4HPC members technical and market knowledge (internal) with that coming from external experts and sources alike. In its activities ETP4HPC gets inputs from other important HPC networks such as PRACE. At the same time input is obtained from HPC end-users and ISVs.

During 2016 ExaFLOW has contributed to ETP4HPC and their SRA, through completing surveys and providing material for their Bof at SC'16.
In 2017, ExaFLOW contributed to the ETP4HPC European HPC Handbook, a publication detailing the HPC Technology, Co-design and Applications Projects within the European HPC ecosystem. With ExaFLOW covered as a HPC Technology Project (other categories were Co-design Projects and Centres of Excellence in Computing Applications) the Handbook was distributed at the SC17 Conference in Denver.

3.1.3 Collaboration with national level projects

As far for projects or initiatives allowing for collaboration on a national level the IMPERIAL and SOTON are part of the UK Turbulence Consortium (UKTC). During the first year of the project Imperial College performed a number of simulations under the auspices of UKTC using ARCHER resources. Neil Sandham of SOTON was the Principal Investigator of UKTC until 2018 and will pursue this opportunity for collaboration in the months to come.
KTH is the lead partner of the Swedish e-Science Research Center, SeRC and the Linné FLOW center. Philipp Schlatter is director of the Linné FLOW Centre, and Dan Henningson the Director of the Swedish e-Science Research centre. Through SeRC, collaborations on efficient implementations and exascale technologies will be pursued, while ExaFLOW is connected to the Linné FLOW Centre both through the research areas of "e-Science" and "Turbulence"; the latter is mainly relevant for the physical interpretation of the ExaFLOW test cases (wings, jet in crossflow).
ASCS presented ExaFLOW within the AINET project, a networking project with the aim to support and promote young scientists and young engineers in the field of simulation and HPC.

3.2 Collaboration with the wider world

This section refers to collaboration of ExaFLOW with projects geographically outside EU. ExaFLOW will collaborate with non non-EC projects where there is mutual benefit from so-doing. This engagement is likely to be a lightweight 'monitoring' and exchange of public domain information but may extend further.

The above set of guidelines enables partners to make appropriate judgements as and when appropriate opportunities for collaboration arise. At the current time the following organisations have been identified outside the EEA with which ExaFLOW can collaborate.

- NCSA, UUIC: collaboration on Nek5000
- University of Utah: Mike Kirby Collaborators on Nektar++
- Honda, Cray and NEC: collaboration as members from the ASCS network
- RIKEN AICS & CCS: collaboration on extreme scale simulations

4 Exploitation activities (D4.3)

4.1 Market outlook

HPC technologies and computer-aided engineering (CAE) software are helping manufacturers drive faster time-to-market by quickly resolving design challenges, forecasting real-world performance, and testing fewer prototypes. HPC innovations are revolutionizing the manufacturing industry with scalable, high-performance software and higher-quality products. Companies are striving to accommodate these transformative technologies; however, the performance, memory, and storage shortcomings of legacy systems are no longer sufficient to power the complex simulations required by today's manufacturers. As the demand for compute capacity steadily rises, companies are looking to optimized platforms to accelerate IT transformation and maximize productivity.

4.1.1 CFD market trends

Computer simulations, in a context of scientific modelling, simulation of technology for performance optimization, or simulation of testing, nowadays represents an everyday activity in most of the research and design organisations around the world. Such simulations can be used to appraise the performance of systems too complex or too expensive to test in real life, explore new technologies, or to simply accelerate the design phase of a new product. An example is a computational fluid dynamics (CFD) simulation, where computers are utilized to simulate the interaction of liquids and gases with bodies of interest.

CFD tools are able to predict the flow of the fluid, calculate the resulting forces and simulate the chemical reactions, fluid-structure interactions, ice-accretion, estimate mass, heat transfer and much more. They are widely used in automotive, aerospace, electronics, defence, material and energy industries, helping engineers and researchers to test and improve mistaken designs even before they get manufactured.

The effectiveness and impact of CFD on the design and analysis of engineering products and systems is largely driven by the power and availability of modern HPC systems. During the last decades, CFD codes have been formulated using distributed memory message passing (e.g., MPI) and shared memory threading (e.g., OpenMP) software models for expressing parallelism to run as efficiently as possible on current generation systems. However, with the emergence of truly hierarchical memory architectures having numerous graphical processing units (GPUs) and coprocessors, new CFD algorithms and new implementations of existing algorithms may need to be developed to realize the full performance offered by such systems.

Some of the main growth drivers in the CFD market are:

1) The requirement to develop **products of higher quality** is therefore one of the major growth drivers in the CFD market[2]. CFD vendors try to avoid inaccurate prediction of fluid flow behaviour for highly complex geometries.
2) The need for the **reduction in time for CFD analysis**. In other words, the trend is to decrease the time an engineer spends preparing and running the simulation, letting him focus on the post-processing phase and analysing the key results of the simulation.
3) **The availability of more powerful high performance computing** enables the end user to obtain results of hundreds of simulations at the

[2] Computational Fluid Dynamics (CFD) Market - Global Industry Analysis, Size, Share, Trends, Growth and Forecast 2015 - 2023, Transparency Market Research

same time or to run computationally demanding simulations of higher fidelity. [3]

4) The **shift from manual to automatic optimization** workflows is accelerated by integrating geometry parameterization and mesh morphing tools into the optimization cycle[4,5], which allow for an automatic optimization loop.
5) **Having CFD tools embedded with CAD software to** enable a fast transition from design to manufacturing. By eliminating the data transfer gap between CAD and CFD applications the product development cycle time is reduced and simultaneously its efficiency increases.

4.1.2 HPC market

The market for high performance computing is huge and growing. IDC estimates growth in the HPC sector will continue to outpace that of the overall IT market for many years to come. Intersect360 Research projects in its latest five-year forecast that spending from 2016 through 2020 for aggregate HPC market will grow at a compounded annual growth rate (CAGR) of 5.2 percent. The current projected growth rate will put the total market value at $36.9 billion by the end of 2020. The forecast uses 2015 as a starting point, which saw $28.6 billion spent on HPC products and services, according to the Intersect360 data, a 2.7 percent increase from 2014. Servers continue to be the largest segment, claiming $10.6 billion of that total.

IDC claims that the growth in HPC system revenues and a big increase in spending by hyperscalers and cloud builders is what propelled the overall server market up 8 percent in 2015 to $55.1 billion; worldwide server shipments across all sizes, architectures, and segments were up 4.9 percent to 9.7 million machines.

HPC spending by industry and commercial businesses continue to drive the market with a five-year growth rate of 6.8 percent. Growth in both government and academia are projected to be less than half of that, with CAGRs of 3.0 percent and 3.2 percent, respectively. The slower expansion of public spending in HPC is attributed to continued government austerity, which the report says looks like it could become a long-term trend.

4.2 Progress on delivering the exploitable results

The exploitable results of the project have been developed in such a way that it is possible to exploit them as standalone solutions. This approach of clearly driven developments was intentional so that the project outcomes would reach a maximum marketability and penetrate niche markets effectively. Here we list the exploitable results and the progress achieved during the project. Next to the exploitable results we list the partners who plan to exploit them during the project

---

[3] https://ntrs.nasa.gov/archive/nasa/casi.ntrs.nasa.gov/20140003093.pdf
[4] https://www.caeses.com/
[5] http://www.esteco.com/modefrontier

as part of their individual exploitation plans, depending on the current needs of each organisation.

Table 9 - ExaFLOW project exploitable results

Exploitable results	Progress during M1-M36
ER1: Novel formulations for error estimators and resilient algorithms in realistic turbulent situations suitable for Exascale Priority for: **KTH, IMPERIAL, SOTON, EPFL** Secondary for: UEDIN, USTUTT, McLaren, ASCS	To develop a fully functional AMR solver based on SEM we considered all its key components starting from efficient nonconforming pressure preconditioners based on adaptive Schwartz and hybrid Schwartz-multigrid methods. These preconditioners relay on coarse grid solvers to solve global Poisson problem on reduced number of grid points. Two approaches were considered: XXT and AMG, and both were adapted for nonconforming meshes. In the case of error estimators we investigated a spectral error indicator for time dependent flows and an adjoint error estimator for steady flows. The second approach was adapted for SEM. A number of tools for grid modifications in Nek5000 were developed, e.g. a two-level partitioner using ParMETIS for optimal grid decomposition was introduced. All these components together allow Nek5000 to perform fully adaptive simulations. Error indicators are implemented in OpenSBLI, which can run on various architectures. These combine with mesh generation which refines or coarsens the original mesh based on the flow physics. Using such an iterative process reduces superfluous grid points and are very helpful for mesh generation in flagship calculations. A C++/MPI Library for multi-level checkpointing called Llama was developed. The library uses lightweight in-memory checkpoints to protect data. Lightweight checkpoints are enabled through the use of erasure codes, encoding and storing checksum parity code locally in groups. This approach serves to mitigate the impact of failures by decreasing the overhead of checkpointing and by minimizing the compute work lost upon failure. Automatic rollback is supported through ULFM MPI. Lightweight checkpoints, that are stored and encoded locally, has

	the advantage of being highly scalable as the resources for encoding and storage scale linearly with the size of the machine. This opposed to traditional checkpoints that are limited by the capacity of the parallel-file-system. Erasure codes have the unfortunate property that if more data blocks are lost than parity codes created, the data is effectively considered unrecoverable. In the context of light-weight checkpointing, this is unfortunate as it forces the application to sometimes rely on the parallel-file-system for recovery. In an attempt to reduce the frequency of such events, a new method for partial information recovery for incomplete checksums was developed. Under the assumption that some meta knowledge exists on the structure of the data encoded, we demonstrated that the data lost may be recovered, at least partially. This result is of interest not only in HPC but also in data-centers where erasure codes are widely used to protect data efficiently. IMPERIAL & EPFL have developed, implemented and tested a minimally intrusive approach to adding fault tolerant implementation of the incompressible solver in Nektar++. A paper has been accepted for publication in the Journal of Scientific Computing and a development branch is available in the Nektar++ framework.
ER2: Novel mixed CG-HDG algorithms <u>Priority for:</u> **IMPERIAL, EPFL** Secondary for: KTH, SOTON, UEDIN, USTUTT, McLaren, ASCS	IC have developed and tested a CG-HDG technique for imposing weak Dirichlet boundary conditions in Nektar++ in a relatively parameter-free manner as compared to existing approaches. This implementation was a first step into providing a CG-HDG mixed solver; however, a theoretical analysis of the full solver highlighted how memory bandwidth/movement would be a limiting factor in such an approach. A paper on the formulation and testing of the weak boundary conditions has been submitted to the Journal of Computational Physics, and a technical report has been made available on the theoretical analysis under arvix. The implementation of the weak boundary conditions are openly available within the Nektar++ open source framework and are due to be merged into our master branch shortly.

ER3: Novel I/O strategies based on feature extraction <u>Priority for:</u> **UEDIN, USTUTT** Secondary for: KTH, IC, SOTON, McLaren, ASCS	The DMD-like mechanism for analysis of time dependent data was further developed and refined with respect to the theory and to a program for filtering of data. The program is enabled to read some formats, especially the VTK-format. It writes data in VTK-format for the visualization of animated data by Paraview. The program is not yet parallelized and not yet optimized. In collaboration with McLaren, IC has overcome a parallel scaling bottleneck with the XML default input format which previously made partitioning/initialisation on above 10K cores in Nektar++ very challenging. A new parallel input format based on HDF5 has been developed and tested and is showing significant improvement and is currently available in a branch of Nektar++awaiting review to go into Master
ER4: Efficient Open Source Pilot Implementations <u>Priority for:</u> **KTH, IMPERIAL, SOTON, UEDIN, USTUTT, McLaren, ASCS**	HDF5 input and fault tolerance work are available in the open source Nektar++ framework under gitlab.nektar.info. Based on performance analysis undertaken at UEDIN, IMPERIAL also identified performance bottleneck arising from adopting a “universal” degree of freedom format for the solution which necessitated too much communication on the elliptic solver setup. IMPERIAL have therefore been redesigning this aspect of the solver, keeping the solution degrees of freedom in a local elemental format until much deeper into the solver algorithm. Future developments may look to keep this format throughout the solver to help maintain data locality and related performance. The I/O routines in Nektar++ have been analysed and alternative output options have been provided to reduce the time taken in load and checkpoint phases. The physical accuracy of Nektar++ compared to the commercial simulation software FLUENT has been investigated. In Nek5000, work has been undertaken to replace original AMG implementation with HYPRE library removing dependence on MATLAB and the preprocessing step. Nek5000 has been extended to nonconforming meshes, that together with number of tools for grid management (based on p4est library) and dynamic domain decomposition (based

	on graph partitioner ParMETIS) allows to perform fully adaptive CFD simulation. This code will become a part of the offical Nek5000 release at the end of 2018. Finally, a replacement implementation of the gather-scatter communication library (GS) used by both Nek5000 and Nektar++ has been developed which reduces latency and improves performance at scale with up to 10%.
ER5: Energy-efficiency driven algorithm designs Priority for: **UEDIN, USTUTT** Secondary for: KTH, IMPERIAL, SOTON, McLaren, ASCS	USTUTT (with some input from UEDIN) has produced a comprehensive set of results showing the implications of energy saving measures on runtimes for CFD use cases. UEDIN foresees the potential of an advisory/consultancy service based on this capability.

4.3 IPR and Licensing

The ExaFLOW consortium defined all matters related to confidentiality and IPR handling in the Consortium Agreement. This agreement formalises project management procedures, IPR issues, and exploitation of results. This agreement ensures that the IPR of third parties will be respected by registering and tracking all use of third party components to ensure that license conflicts are not generated. The IPR of participants will also be protected by ensuring that the impact of the licenses of any third party software is analysed prior to its use (particularly concerning GPL and other viral open source licenses). All background owned by members is clearly stated in the project's Consortium Agreement and the IPR policies clarified. IP created during the project will be the property of the partner who creates it, however, all IP created during the project will be available to other partners for use on the project without payment. Use of that property following the conclusion of the project will be subject to the normal considerations. All deliverables produced by the project that do not include financial information or security-related issues will be made public, and the project follows an open source policy for the technologies and interface developed in the project to facilitate technology transfer to the business sector and encourages the uptake of the results in third-party products.

The initial background included the existing CFD methods as implemented in the pilot codes as well as use case geometries. This particularly includes:

- Nek5000 code, available under BSD license at https://nek5000.mcs.anl.gov/
- Nektar++ code, available under the MIT license at www.nektar.info
- McLaren front wing geometry, protected by McLaren, will be made available under NDA to the academic consortium members.

- NACA 4412 airfoil geometry, standardised geometry published by the National Advisory Committee for Aeronautics (NACA) Report 460, 1935.
- Jet in crossflow geometry, published by Peplinski, Schlatter and Henningson, Eur. J. Mech. B/Fluids (2014).
- Opel Astra rear part geometry, protected by ASCS and its member Adam OPEL AG, could be made available under NDA to the consortium members.

Most recent releases/updates are:
- SBLI code, currently available for collaborators from the University of Southampton. An open-source release which features automated code generation capabilities, called OpenSBLI, has been released under the GNU General Public License on GitHub. See https://opensbli.github.io
- ExaGS, a new low latency communication kernel for Nek5000, released under BSD license on GitHub
- Nek 5000 has moved to GitHub (https://github.com/Nek5000)
- Nek5000 v17 released in December 2017
- Nektar++ v4.4.1 released in October 2017

The licensing of the project results has been considered by the consortium from the very beginning of the project in a joint effort made by technical and exploitation teams. The consortium has decided that the project should be "as open as it can be". However, due to the manifold underlying licenses as well as the commercial partners ASCS and McLaren, there are obstacles which lead to a more complex licensing and IPR management. As a result, in the table below we list the licence types the consortium considers applying to each of its exploitable results.

Table 10 - ExaFLOW project exploitable assets' licences

Exploitable results	**Licence type**
ER1: Novel formulations for error estimators and resilient algorithms in realistic turbulent situations suitable for exascale	MIT Licence
ER2: Novel mixed CG-HDG algorithms	Open Literature, MIT licence
ER3: Novel I/O strategies based on feature extraction	MIT License
ER4: Efficient Open Source Pilot Implementations	Open source
ER5: Energy-efficiency driven algorithm designs	Copyright of the consortium

4.4 Industrial stakeholders outreach

We plan to identify and contact stakeholders demonstrating to them the value of the results through example tools, documents and code as well as giving an introduction into "How to Adopt" the ExaFLOW project exploitable results. We will seek the possibility to jointly initiate a sustainability path, thus establishing and nurturing relations and collaborations with third parties.

Exploitation is of clear importance to the impact generation, in particular with regards to uptake of open source project results. The results of ExaFLOW will be communicated to industrial beneficiaries, both through the industrial links of the partner organisations and the ETP4HPC. Table 8 below lists the main industrial contacts of the project partners.

Table 11 - Main Industrial Contacts

Partner	**Main Industrial Contacts**
KTH	SAAB, Scania, Vattenfall, Huawei
IMPERIAL	McLaren, Airbus, BP, British Gas
SOTON	Airbus, Vestas
UEDIN	Rolls Royce, ICON CFD, Prospect FS, BAE Systems
USTUTT	Porsche, Daimler, RECOM
EPFL	HyperComp Inc (US)
ASCS	Porsche, Daimler, Opel, HONDA, TECOSIM, Altair, CD-adapco, MentorGraphics, ESI

Moreover, ASCS has established a series of workshops for keynote speakers on the subject of automotive simulation using the HPC technologies. One of the main aims of the workshops was the promotion of the results of Exaflow, with an accent on the last developments of innovative CAE and HPC methods and the creation of synergies between the companies active in this field.

4.5 Individual exploitation reports

At the beginning of the project, all partners described in their individual exploitation plans, what they wanted to achieve with the ExaFLOW project, and how their involvement in ExaFLOW could benefit their institutions. These individual exploitation plan can be found in the Grant Agreement document and in the WP 4 deliverable in PM 18.
Now, the paragraphs below show for each partner how their individual exploitation has in fact turned out during the project runtime and plans on how to exploit and further develop their achievements in the future.

KTH

KTH has integrate the results of the project into the Nek5000 code, which is used in a number of academic and industrial research projects and one of the most used codes on the KTH HPC resources. Thanks to the open source nature of Nek5000 the outcome of the project is readily available to the wider research community and KTH will continue developing the code after the end of the project. KTH is particularly committed to further exploit ER1 and ER4 and ER2, ER3, and ER5 all which will likely have an impact on KTH's future development of Nek5000. KTH has also exploited the results via its HPC Center, PDC, which made the modified Nek5000 publicly available on its resources and used the results, specifically ER1 and ER2 in technology transfer activities with its many academic and commercial CFD users. A first tangible exploitation result is the European Commission funded centre of excellence Excellerat, the European Centre of Excellence for Engineering Applications. The centre, which comprises the groups from KTH, USTUTT, UEDIN participating in ExaFLOW together with ten European universities and research institutes aims to directly apply the ExaFLOW results to industrially relevant CFD problems. The outcome will also be exploited and further developed by SeRC, the Swedish e-Science Research Centre, the Swedish Foundation for Strategic Research (SSF) Infrastructure Fellow program and through the Wallenberg Academy Fellow programme. During the project KTH achieved the following:

- Efficient nonconforming pressure preconditioners based on adaptive Schwartz and hybrid Schwartz-multigrid methods were developed and implemented in Nek5000. An important finding was proper use of the transpose and inverse interpolation operators in the direct stiffness summation of the Schwartz operator.
- Interface to p4est and ParMETIS libraries was re-implemented to improve code efficiency. Special attention was paid to the optimal domain decomposition that is crucial for efficient coarse grid solver setup. To achieve it two-level graph partitioning was introduced allowing to place entire graph branch on a single computational node and to reduce internode communication in the coarse grid solver.
- To improve coarse grid solver an AMG library HYPRE was integrated with Nek5000 tools and tested.
- The under-resolved mesh regions are identified using a spectral error indicator for time dependent flows and an adjoint error estimator for steady flows. The second approach was adapted for SEM.
- In cooperation with UEDIN developed a new low latency gather-scatter communication library for Nek5000 and Nektar++

IMPERIAL

IMPERIAL have integrated the results of the project into the open source Nektar++ framework (www.nektar.info), which is used in a number of academic and industrial research projects. Components of this code are also linked to the NEK5000 code. Thanks to the open source nature of Nektar++ and NEK5000 the results will also be readily available to the wider research community and IC will continue developing the code after the end of the project. During the project IMPERIAL achieved the following which is now available in the Nektar++ framework:

- IMPERIAL have developed and tested a CG-HDG technique for imposing weak Dirichlet boundary conditions in Nektar++ in a relatively parameter-free manner as compared to existing approaches. This implementation was a first step into providing a CG-HDG mixed solver; however, a theoretical analysis of the full solver highlighted how memory bandwidth/movement would be a limiting factor in such an approach. A paper on the formulation and testing of the weak boundary conditions has been submitted to the Journal of Computational Physics, and a technical report has been made available on the theoretical analysis under arvix. The implementation of the weak boundary conditions are openly available within the Nektar++ open source framework and are due to be merged into our master branch shortly.

- In collaboration with McLaren we have overcome a parallel scaling bottleneck with the XML default input format which previously made partitioning/initialisation on above 10K cores in Nektar++ very challenging. A new parallel input format based on HDF5 has been developed and tested and is showing significant improvement.
- In collaboration with EPFL we have developed, implemented and tested a minimally intrusive approach to adding fault tolerant implementation of the incompressible solver in Nektar++. A paper has been accepted for publication in the Journal of Scientific Computing and a development branch is available in the Nektar++ framework.
- IMPERIAL has also considered memory layouts and parallel frameworks to help enable best use of many-core architectures in Nektar++. Our initial efforts along these lines have been published in a journal article in Computer Physics Communications, which examine the use of the Kokkos framework (developed by Sandia National Labs) for the parallelisation of a high-order mesh optimisation technique using both CPUs and GPUs. Following on from this work, we have performed a further study comparing OpenMP 4.5, OpenACC and Kokkos, in combination with an adjusted memory layout, for an elliptic solve that forms a key part of the incompressible Navier-Stokes equations. This work is presently being submitted for publication to Computer Physics Communications. Finally, together with Prof. Mike Kirby (University of Utah), we are examining in detail a matrix-free implementation of the Helmholtz operators with this new data layout, which we believe will significantly improve computational performance on modern architectures.
- Based on performance analysis undertaken at EPCC we also identified performance bottleneck arising from adopting a “universal” degree of freedom format for the solution which necessitated too much communication on the elliptic solver setup. IC have therefore been redesigning this aspect of the solver, keeping the solution degrees of freedom in a local elemental format until much deeper into the solver algorithm. Future developments may look to keep this format throughout the solver to help maintain data locality and related performance.

SOTON

SOTON has integrated the project in the automatic source code generation software OpenSBLI. This is a high-level implementation of the concepts of the compressible flow SBLI code, with ease of implementation on various architectures. The core features are made available in the Year 1, and a latest beta version of the developments is released[6] as a branch in the repository, so that they are available to the wider community . The code is being used for a number of mainly academic projects and is one of the main codes of UK Turbulence Consortium. These various components are used in conjunction with SBLI code, for the flagship calculations.

UEDIN

The University of Edinburgh, through its supercomputing centre EPCC, has set exascale computing research as its key computational science research priority over the next decade. Power efficiency of HPC and software scalability on a massive scale are two crucial parts of the research that needs to be undertaken on the path to exascale.

UEDIN firmly believes that a long-term solution can only be found in close collaboration with experts from across all computing and science segments; thus ER3, ER4 and ER5 are of particular interest. ExaFLOW and the work done as prt of the project is an important component in our strategy to achieve our goal of moving HPC into the exascale era. For UEDIN, multiple exploitation channels exist for the output of ExaFLOW. These include publications, the production use of the software developed in the project, continued research and development of the software and ideas developed during the project, and the transfer of the results into the scientific and industrial domain.

In more detail we have:

- Publications: published our research results in academic journals, and in presentations and papers at conferences, for example at CUG2018 and EASC2018;
- Production use: the ExaFLOW algorithms and prototype implementations have been tested on UEDIN HPC services and are available (given their open source nature) to all users. This exploits the expertise gained in understanding power usage of HPC algorithms and performance of IO methods in our day-to-day work as a supercomputing centre.
- Continuing R&D: we will continue to take forward the results gained from ExaFLOW in subsequent and on-going projects and other activities such as novel hardware design projects with vendors, based on knowledge gained from efficient algorithm design and implementation. For example, the work on establishing roofline analysis method directs us to the best use of time in further optimisation rounds.
- Technology-transfer: we will also use the results in our work with our many academic and commercial customers who use CFD applications in their daily business.

USTUTT

---

[6] OpenSBLI main website: https://opensbli.github.io ; Project source code on GitHub: https://github.com/opensbli ; Documentation: https://opensbli.readthedocs.org

ExaFLOW is helping USTUTT-IAG and -HLRS in first instance by allowing for further scaling of CFD codes. HLRS as provider of facilities for large scale capability jobs has a vital interest of the further development of scalability of computational fluid dynamic codes because these take the major part of the computing cycles of all machines of the centre and thus are a major consumer of resources and energy. USTUTT-IAG will be able to increase the overall performance of their analysis by the proposed approach (by providing and using techniques of doing faster IO and by showing ways of extracting the physically and technically relevant flow features of an unsteady flow). This approach which will at the end lead to a reduction of storage needs is also a vast interest for USTUTT-HLRS as this will avoid unnecessary investments in excessive increases of storage capabilities. Finally, the energy efficiency aspects of ExaFLOW will provide another aspect for the Green IT and energy reduction activities of HLRS to reduce costs and minimize the environmental impact. In addition, USTUTT-HLRS will use its synergies with the Institute of High Performance Computing (IHR) at the university and embed the research and technical objectives as subject in its lectures and of doctorate thesis work. Furthermore, the results and experiences of this project will enrich USTUTT-HLRS´ training activities which will be enhanced and increased with the construction of a dedicated training centre and the release of new, improved training plans.

EPFL

EPFL will focus on the development and analysis fault tolerant and resilient algorithms at the exascale, including the development of suitable approaches for in-situ model development and fault detection strategies. EPFL will also be involved in the development of new scalable solvers and their implementation. The developments will be transitioned to the Nektar++ code and, through shared code base, to NEK5000.

A library for fault tolerance trough checkpointing, called Llama, was developed. It is the first library to support both automatic rollback without restart, and the use of an arbitrary number of checkpoint levels with topology aware checksum checkpoints of arbitrary group size and number of parity code blocks. We've demonstrated that the library can achieve in excess of 300GB/s data encoding when running on 256 nodes.

The library is still under development and there are several points that needs to be addressed in future versions. At the moment, only the use of spare nodes to replace failed workers has been implemented, i.e., there is no support for spawning new ranks as needed. Another issue that remains is that upon injecting a large number of errors, the failure-detection and communicator repair may fail. Identifying the underlying cause of these issues is ongoing work. Once the issues have been identified and removed, the library will be tested on applications written using Nektar++ whilst injecting failures at rates predicted to occur at Exascale.

In addition to the work on developing the library, we've investigated the idea of partial information recovery in incomplete checksums. We proposed a new method, and tested its application on a weighted checksum scheme for floating

point numbers. Our preliminary finding is that it is indeed possible to partially recover the data otherwise considered lost when having some knowledge of the underlying structure of the data encoded.
The method developed has certain limitations. The data recovered is not of machine accuracy with respect to the original data as is otherwise normally the case, so for the approach to be of practical use in many applications, one would need some way of quantifying the accuracy expected of the data recovered. In addition, the method was tested using Gaussian matrices to encode checksums on floating point numbers. For the method to have practical relevance outside of HPC, it must be extended to the case where the checksum matrix elements are from a Galois field, i.e. Reed-Solomon codes. There are no obvious reasons to believe that the fundamental idea of alternating between enforcing the constraint of the checkpoint equations, and applying some filter, should not work, but a method for enforcing constraints for incomplete Reed-Solomon has yet to be developed.
If successfully extended, this would mean that failures that would otherwise result in complete data loss would instead only result in loss of data fidelity. This could in turn potentially lead to relaxed requirement on data protection for a wide number of applications in HPC fault tolerance and data-centers.

McLaren
The application of time averaged RANS CFD simulation was the lead technology behind McLaren Racing's aerodynamics design at the time where the ExaFlow project was granted.
Given the inherently transient nature of the flow field around open wheel race cars, a natural evolution of this important design tool is the introduction and wider application of transient flow modelling to achieve even greater aerodynamic performance.
A first step towards this goal is the use of Unsteady RANS and Detached Eddy Simulations, now common tools in the Automotive Industry, and the next is implicit Large Eddy Simulations.
For this type of modelling to be applied to a full car simulation undertaking complex manoeuvring conditions within a suitable timescale for design will require Exascale levels of compute based on the tools provided by ExaFlow.
Successful application of this technology was therefore seen as not only a key enabler in McLaren Racing's race winning objectives but will also feed into our other applications of luxury automotive car design and technology transfer through McLaren Advanced Technology division.
Over the course of the ExaFlow project and thanks to the work undergone by partners of the consortium, we have been able to assess the feasibility of High Order Methods, greatly improve their use and automation within our process, and prove that they could be used on geometries as complex as ours.
We have also undergone Experimental measurements on a realistic Formula One front wing and simulated the same geometry to assess the benefit of such techniques on the accuracy of the solution, both on the time-averaged and transient sense.

Starting from weeks of user intensive preparation for a relatively simple geometry (by our Industry standards) to reach our clusters, and quite often with limited success, and thanks again to the work undergone by the consortium, we have reached the point where a detailed full car could be prepared within a couple of days with limited user interaction. This was a fundamental step to bridge the gap between the theoretical interest of the technique and their potential use in our field given our tight development cycles.
If their cost is still too prohibitive to be used in production as it stands given the current Formula One regulations which cap CFD simulations (to roughly 3000 cores of nominal usage for all Aerodynamics design simulations), the situation is bound to evolve to follow the trends in other Industrial sectors in the coming years.
As the project comes to an end, we have been able to train R&D engineers to use those techniques and made the simulation of simple test cases an almost routine task when increased accuracy is required. We expect to do more work internally in the coming years to improve our methodology and keep collaborating with other partners of the consortium.
The overall cost of those simulations has also been greatly reduced, and we could see how other Industrial sectors not subject to simulation caps could within the next few years use them more systematically.

ASCS
As an applied research institute, the ASCS is interested in the results of ExaFLOW in order to maintain the position of a transfer platform in the field of future-oriented virtual vehicle development. The current very strong trend of increasing computing time in CFD simulations using scale resolving modelling methods benefits the penetration of new exascale technologies in the automotive sector. The ASCS wants to increase the awareness and knowledge of exascale solutions for this application. The ASCS looks forward to find new approaches as specific HPC-simulation codes, to push the usage of HPC clusters.
The gained expertise will be the basis for further conception and implementation of research projects in the field of CFD simulations for the future-oriented vehicle development process. Furthermore, the conflation of forces engaged in research with industrial practice for the purpose of reciprocal exchange on current issues, the dissemination of scientific results relating to modelling and simulation in exascale environments, to be used in practical applications including the method-oriented support of users is focused with this project.
The ASCS use the project expertise to support companies in improving their development process by accelerating the pre- and postprocessing as well es generating new calculation methods and consistent workflows to reduce their development costs.

5 Innovation management (D4.5)

5.1 The innovation management process

This section presents the ExaFLOW project innovation management process, which is also depicted in Figure 6. It consists of an innovation cycle with four steps: Idea generation, Architecture development, Technical Implementation and Increase in TRL. During this process the Executive Board (EB) of the project makes sure that feedback from the market, the end users and the scientific and open source communities makes its way to the development teams. The WP4 leader who is part of the EB is informed regularly on the progress made in regards the Exploitation Results of the project and updates the IPR status of each ER and/or other components/outputs of the project.

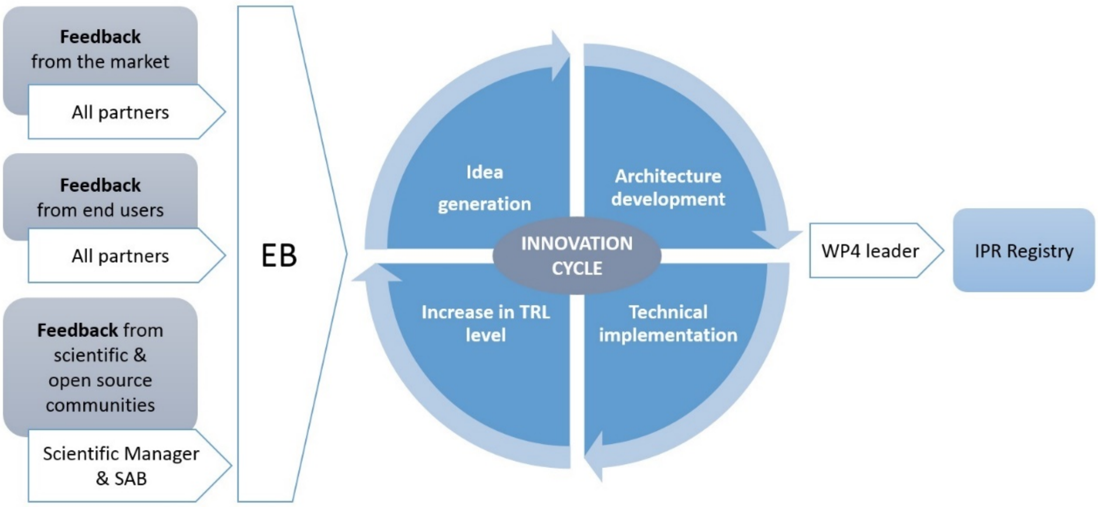

Figure 6 - ExaFLOW Innovation Management Process

5.2 Summary on innovation metrics

The goal of ExaFLOW is to address key algorithmic challenges in CFD to enable simulation at exascale, guided by a number of use cases of industrial relevance, and to provide open source pilot implementations.

ExaFLOW produces a number of clearly defined innovations in the area of exascale computing and CFD. In the table below we describe the progress made in year 1 (Y1) of the project in regards each innovation.

Table 12 - Innovations of the ExaFLOW project and their metrics

Innovations	Success Metrics

<table>
<tr><td>Innovation 1: Mesh Adaptivity, Heterogeneous Modelling, and Resilience</td><td>Reduction of simulation costs of up to 50%.</td></tr>
<tr><td colspan="2">Progress:
KTH: Within ExaFLOW fully functional SEM solver capable of dynamical modification of computational gird according to estimated computational error was developed. Special attention was paid to efficient nonconforming pressure preconditioners, so two different approaches: additive Schwarz and hybrid Schwartz-multigird were investigated. We considered as well two different coarse grid solvers for nonconforming pressure preconditioner: XXT and AMG. In the case of AMG solvers HYPRE library was found to be an optimal solver and it was implemented in Nek5000. Other aspect was implementation of efficient tools for mesh management using p4est and ParMETIS libraries. To improve grid partitioning we have introduced two-level partitioning scheme, where distribution of cores between different computational nodes using MPI3 features is taken into account. The computational error is estimated using spectral error indicator (for time dependent problems) and adjoint error estimator (for steady flows). All mentioned modifications were implemented in Nek5000 turning it into AMR solver. We have performed number of 2D and 3D test cases (e.g. tilted lid driven cavity and flow past circular cylinder) debugging the code and verifying results. The final test case is our flagship test case which is the turbulent flow over a NACA-4412 wing profile. In this case we fulfilled our success metric by reducing the effective computation time by more than 50%. It was achieved by reducing the number of elements from 335664 (conforming reference case) to 224328 (nonconforming mesh), which is more than 50% of conforming case, but. the nonconforming mesh was three times wider than the conforming one. The wider domain gives number advantages allowing to study more complex flow cases and e.g. shortening simulation time necessary to collect statistics.
SOTON: A heterogeneous modelling approach has been implemented in an incompressible flow solver, and evaluated using a channel flow simulation. The approach uses a combination of LES and quasi-DNS. For the lower Reynolds number channel flow simulations only 25% of the grid points are used compared to the reference data to get a reasonable agreement with the published data. This method has shown some limitations for higher Reynolds number flows and requires further investigations.</td></tr>
<tr><td>Innovation 2: Strong scaling at Exascale through a mixed CG-HDG</td><td>Improved absolute performance will be greater than, or comparable to, existing individual schemes for current levels of parallelism and improved for larger levels of parallelism available on exascale systems.</td></tr>
<tr><td colspan="2">Progress
IMPERIAL:</td></tr>
</table>

<table>
<tr><td colspan="2">The algorithm was explored and work divided into two parts:
1) assembly of global CG-HDG system that couples together multiple mesh partitions
2) implementation of the ‘interior solve’ from HDG, but applied to a single mesh partition here. This amounts to the imposition of weak Dirichlet boundary conditions for the continuous Galerkin method.

Task 2 was fully implemented and tested and is publicly available as a part of Nektar++.
After additional performance analysis, it was found that the first part of the algorithm can’t be faster than traditional CG or HDG method, because it requires repeated matrix-vector multiplications with dense matrices whose rank is too large for current hardware architectures and the algorithm would therefore suffer from memory bandwidth limitations. In addition to that, the time complexity to solve the coupled system scales with a power of basis polynomial degree p that grows faster than for CG or HDG.
The first part of the algorithm was therefore not implemented in Nektar++.

The weak boundary condition (part 2) is useful in itself and a paper with detailed description has been submitted for publication. A technical report with performance analysis will be eventually uploaded to arXiv or similar repository.</td></tr>
<tr><td>Innovation 3: I/O data reduction via filtering</td><td>Reduction of operation time for the complete workflow, i.e. simulation (data generation), data processing and data I/O</td></tr>
<tr><td colspan="2">Progress:
Some data compression and decomposition methods have been developed at USTUTT. The data generated by project partner SOTON have been tested with the developed methods before I/O. The alleviation of the I/O bottleneck by considerable data-reduction before I/O has been successful for smaller data sets using Singular Value Decomposition (SVD) as a data compression method as well as the wavelet based compression algorithm, which could perform feature extraction in raw data.</td></tr>
<tr><td>Innovation 4: Energy Efficient Algorithms</td><td>Quantifiable reduction of the total energy consumption and instantaneous power draw for the co-design applications. We expect 20% reduction (depending on the application).</td></tr>
<tr><td colspan="2">Progress:
UEDIN + USTUTT: runs of industrial and baseline test cases showing energy saving versus runtime trade-offs for different clock frequencies across different phases of codes. 20% reduction achievable in compute phases at 5% runtime increase.
SOTON: energy/power analysis of 5 different finite difference algorithms of varying computational and memory intensity on CPU's, GPU's and Intel KNL's. Applied to the Taylor-green vortex test case an energy reduction of 40-50 % was achieved by reducing the memory access and increasing the computational</td></tr>
</table>

intensity compared to the traditional memory intensive algorithms in which the finite-difference derivatives are stored to global work arrays.

All ExaFLOW innovations are clearly targeted to enhance the efficiency and exploitability of an important class of applications on large-scale (Exascale) systems.

5.3 Progress of TRL

Within ExaFLOW we developed fully functional CFD solver based on spectral element method (SEM) and capable of dynamical modification of computational mesh according to estimated computational error. In WP1 we considered all the key components of such AMR solver paying special attention to the parallel efficiency of the developed algorithms. Extending previous work of EU project CRESTA we improved mesh regeneration and partitioning phases using standard libraries p4est and ParMetis for mesh management and graph bisection respectively. An important enhancement was here introduction of a two-level partitioning scheme, that significantly reduced setup and solve times of the coarse grid solver, and improved code scalability. This was critical for the solution of the pressure equation, which is a main source of stiffness for the incompressible Navier-Stokes equations. Development of the efficient nonconforming preconditioner to the pressure problem was one of the main tasks of WP1. Two different approaches: adaptive Schwarz and hybrid Schwarz-multigrid methods were considered and adapted for nonconforming solvers. Their important part is coarse grid solver, which provides global solution to the Poisson problem on reduced number of the degrees of freedom. In this case we investigated two different solution methods: XXT and AMG.

The AMG solver is preferred for large cases (more than 100,000 spectral elements) on large systems (more than 10,000 cores). The native version of the AMG solver has been replaced by the Hypre library for linear algebra. This new version allows to perform the setup phase at run time, at a fraction of the setup time of the original version. Furthermore, preliminary tests have shown that similar solver times are observed for both versions.

Another important aspect is error indicator, that marks the domain regions with insufficient resolution. Within ExaFLOW we considered two different methods: spectral and based on the dual problem, and the adjoint error estimator for steady flows for SEM was developed in WP1.

The adjoint error estimators aim at minimizing the error in the computation of a functional of interest to the problem in the most efficient way. These estimators are goal-oriented and consider both the local properties of the solutions and some sensitivity map given by the resolution of an adjoint problem. The adjoint error estimators have been applied to steady cases so far and have been shown to reduce the total error in the functional of interest by up to one order of magnitude compared to the spectral error indicators for a similar degree of refinement.

All the mentioned algorithms were implemented in Nek5000 and tested within WP2 and finally examined in WP3 by running big flagship test case. Regarding AMR the flagship test case was the turbulent flow over a NACA-4412 wing profile for two Reynolds numbers 200,000 and 1,000,000. These runs were performed on two supercomputers Beskow (PDC) and Hazel Hen (HLRS) and are described in WP3 deliverable. They show the project to be mature and to reach technology readiness level TRL 6.

6 Standardisation

The objective of the standardisation activities was to assess, track and contribute to relevant standardisation possibilities and acts as a collaboration gateway to relevant standardisation bodies.

The project partners embraced and enhanced open standards where possible in order to allow for the best possible exploitation of the project results. While for the developments made in this project there are no standardization bodies as such, the community is impacted through large community events like ECCOMAS, ERCOFTAC, EUROMECH, ICOSAHOM, and ICCFD, which have particularly been targeted by the dissemination activities described in chapter 2. This was complemented through national efforts like the UK Turbulence Consortium (UKTC). In addition, the developments of ExaFLOW used and followed the developments of standards and community best practices for exascale computing such as MPI, OpenMP, OpenACC/CUDA, etc.

7 Conclusions

This deliverable is a merge of several deliverables of WP4 to be published in PM36. It contains the final reports on Exploitation D4.3, Dissemination and Communication D4.8 and Collaboration D4.11 as well as Innovation management D4.5.

The overall goal of activities in WP4 was to enable academic as well as commercial pre-competitive sharing of knowledge and its preservation in the form of standards where possible. This has been achieved through extensive participation in community events that focus both on academia (such as ICOSAHOM) and on industry (such as the ASCS general assembly), as well as close collaboration with mostly EU funded projects which were especially valuable for the exploitation of developments in related research areas. Most dissemination KPIs could be exceeded, but especially conference presentations and presence at events. This highlights the overall importance of research areas related to large-scale CFD. The high engagement rates on Twitter support the notion that the general interest on the research topics of the ExaFLOW project was high.

In this document we also outlined our innovation management process and connected it to the different management bodies of the project. Two flagship runs performed on large scale suggest that Technology-Readiness Level (TRL) 6 could be reached. However, as exascale machines are not yet in operation, some concepts could only be approximated.