THEME PAPER TEMPLATE FOR ARCHER2

TITLE: Combating the challenges of energy efficiency and environmental friendliness in new generation combustors by utilising advancements in high-performance computing

VISION: The efficient and safe use of energy and the understanding of its impact on health, transportation and climate change pose major societal and economic challenges. The improved fundamental understanding and modelling of turbulent reacting flows play pivotal roles in the effective utilisation of energy resources and manipulating the combustion process to ensure environmental friendliness. The UK Consortium on Turbulent Reacting Flows (UKCTRF) intends to maintain its excellent track record in world-leading research by (i) exploiting HPC resources for world-leading turbulent reacting flow research involving Reynolds Averaged Navier-Stokes (RANS), Large Eddy Simulation (LES) and Direct Numerical Simulation (DNS) to ensure that the UK remains at the forefront of turbulent reacting flow research; (ii) utilising fundamental physical insights from simulations to develop high-fidelity modelling methodologies for gas turbine, automotive and fire-safety engineering; and (iii) ensuring a forward-looking software strategy to develop computationally efficient algorithms, and effectively exploit current and future HPC hardware.

CURRENT AND FUTURE DEMAND:

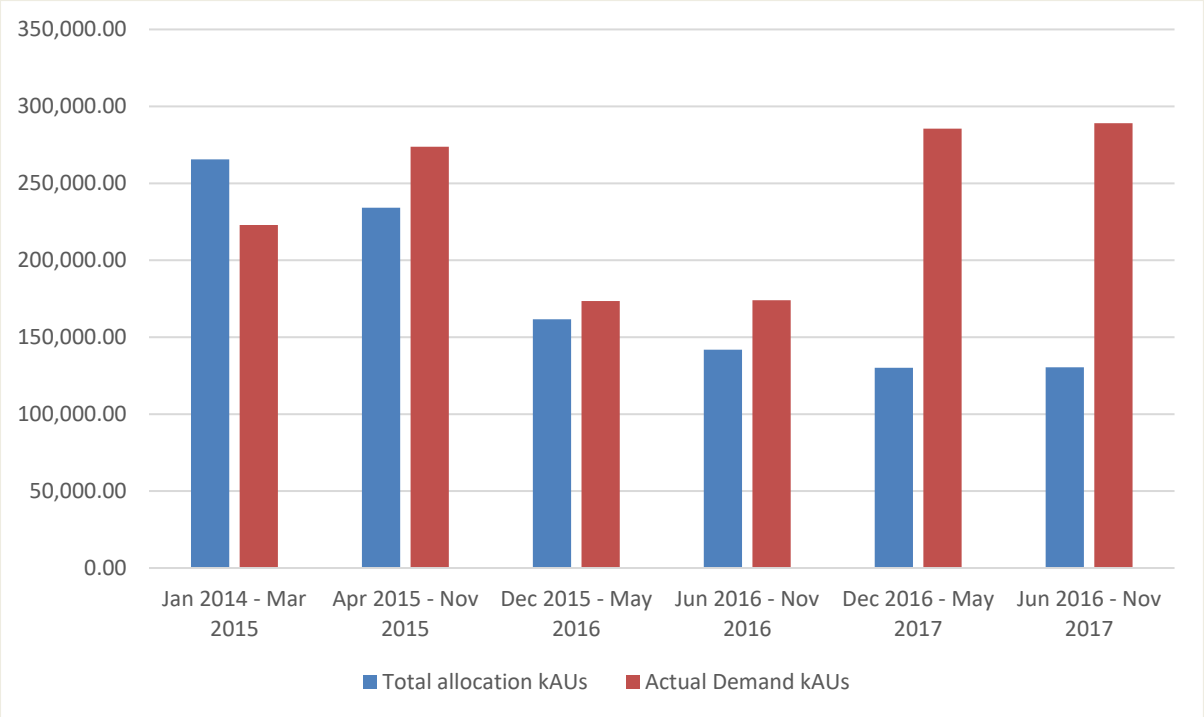

Figure 1: Evolution of computational time allocations to the UKCTRF and demand of its members during the lifetime of the consortium (2014 Jan- To date).

The evolution of computational time allocations to the UKCTRF and demand of its members during the lifetime of the consortium (2014 Jan- To date) are shown in Fig. 1. It can be seen that the demand was smaller than the available allocation only for the first allocation period (December 2013-March 2015) but for every subsequent allocation period the demand of the consortium has been greater than the available allocation. It is important to note that UKCTRF was launched much later (about 6 months) than other consortia, which were funded through

the previous High-End Computing (HEC), call in 2012. The UKCTRF research community did not have access to high-end computational resources for a long time so the members of the consortium were reliant on regional HPC facilities (e.g. N8, Cambridge HPCS facility etc.) and on computational time allocations from their existing research grants. Thus, some research groups within the consortium were using the computational time allocation in their own research grants before applying for computational time from the consortium. This contributed to the under-utilisation in the first allocation period. As a result of continuous encouragement by the management team of the UKCTRF the level of utilisation improved over time and since March 2015 the demand surpassed the available allocation. It is worth noting the available allocation amount is communicated to the members at the beginning of the allocation period so that they can tailor their applications accordingly. In spite of that the demand in last two allocation periods was 3 times of the available allocation. It is worth noting that the demands shown in Fig. 1 are the figures based on the KAUs requested in the application forms but by that stage the consortium members have already moderated their demands based on the knowledge of the available allocation. Thus, the actual demand is much greater than the demand shown in Fig. 1. Under this circumstance the management team needed to take actions so that the available computational time is fairly distributed among the applications. In the last allocation period, once the scientific excellence of the applications is established based on the internal peer review, the priority was given to the applications from the named academic investigators who did not receive allocations in the past and those who requested for a small amount (<5000 KAUs) and the remaining computational time was distributed equally among all other investigators (excluding those who applied for the first time or asked for a small amount or in cases where reviewers raised serious concerns). This allowed us to allocate 13250.5 kAUs to most of the named academic investigators (and their groups) who made an application for computational time during the last allocation period (Jun-Nov 2017). Moreover, the management team of the UKCTRF has decided to limit the number of applications to 2 per academic investigator's group to ensure fairness of the allocation process under the situation of limited availability of the computational resource. The management team actively monitors the utilisation of the allocated computational time and reallocates the time allocation of any user who fails to utilise their computational time significantly after two warnings. Furthermore, the management team takes a flexible approach in altering the computational allocation, wherever possible, to allow PhD students to complete their analyses for their theses. However, this moderation of computational time has a detrimental effect on the ambitious research activities of the academic investigators within the consortium who are either altering the number and size of the computations to work within the allocated computational time or waiting a long time to get the computation finished as only a small part of computation is done within a particular allocation period or they are utilising their international contacts to run large simulations on supercomputers based in other countries (e.g. PRACE, K-Computer, Gauss).

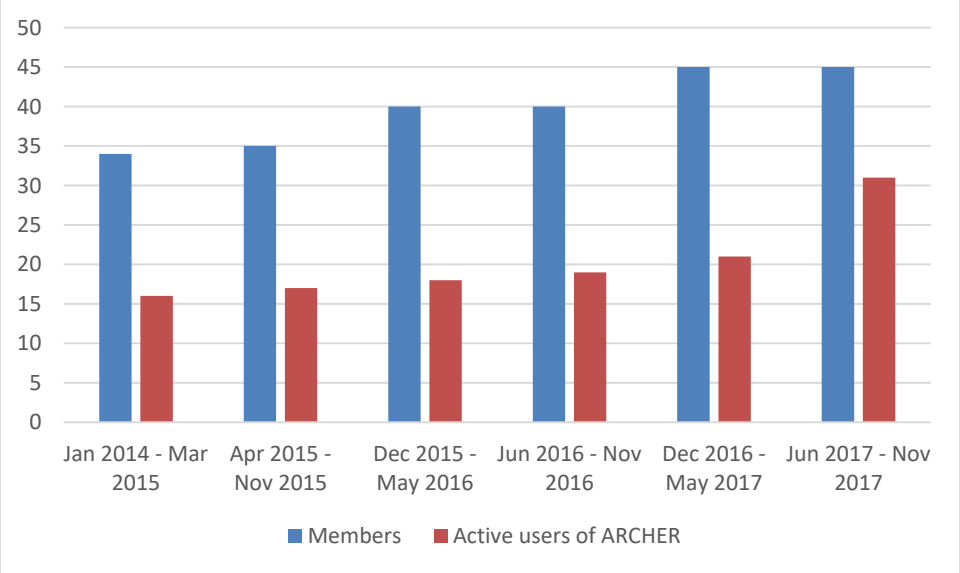

Figure 2: Evolution of the number of academic investigators and active ARCHER users during the lifetime of the consortium (2014 Jan- To date).

The evolution of the number of academic investigators and active ARCHER users during the lifetime of UKCTRF are shown in Fig. 2, which shows that the number of investigators and active ARCHER users are increasing with time. It is worth noting that the majority of active users (99%) are early career researchers and thus it is expected that some of them will obtain academic positions and join the consortium in the near future. This suggests that the demand for computational time within the UKCTRF is expected to rise significantly in the near future. At least 2 times of the current available allocation is needed for the consortium to be internationally competitive in research and innovation but realistically no major step change in the state of computations can be expected with 5 times of current allocation because the increased allocation will be distributed larger number of HPC users in the future.

SCIENCE PROJECTIONS: Over the next 3 years (2017-2020), the goals of the UK Consortium on Turbulent Reacting Flows (UKCTRF) are to: (i) exploit HPC resources to conduct world-leading turbulent reacting flow research involving Reynolds Averaged Navier-Stokes (RANS), Large Eddy Simulation (LES) and Direct Numerical Simulation (DNS), ensuring the UK's world-leading status in turbulent reacting flow research ; (ii) extract fundamental physical insights from simulations to develop high-fidelity modelling methodologies to analyse of turbulent reacting flows with relevance to gas turbine, automotive and fire safety engineering; and (iii) ensure a forward-looking software development strategy to develop computationally efficient algorithms, and effectively exploit current and future developments of HPC hardware. The long-term research (2019-2024) of the consortium will build on the strengths of UKCTRF and will address universal challenges of energy efficiency, environmental friendliness and high-fidelity fire safety. The progress in HPC will enable the consortium to reinforce existing strengths, but also address the following multi-physics, multi-scale challenges, which are yet to be analysed in detail due to the limitations of computational power (in terms of both storage capacity and computational time):

- Simulation and modelling of multi-phase (e.g. droplet and pulverised coal/biomass combustion) reacting flows;
- Combustion analysis of biogas and gases derived from coal gasification;
- Flame-wall interaction
- Combustion at elevated pressures

The case studies based on the above topics are given below:

Simulation and modelling of multi-phase (e.g. droplet and pulverised coal/biomass) reacting flows

Model situation with 5x allocation: This scale of resource increase will allow the consortium to maintain its push towards a deeper understanding of turbulent droplet-laden and particle-laden reacting flows where the inherent numerical methodology will not change but the simulation domain for DNS will become more complex allowing for more detailed chemical mechanisms to be used so that the effects of flow shear, extinction and re-ignition can be simulated. These DNS data can be processed to develop high-fidelity flamelets, Probability Density Function (PDF) and Conditional Moment Closure (CMC) models for industrial engineering calculations based on RANS and LES.

Model situation with 10x allocation: At this scale of allocation, the consortium will be able to retain and expand its international leadership by tackling outstanding challenges in e.g. including atomisation processes in spray combustion. To date, no DNS and LES have been carried out for atomisation of sprays in a reacting environment. With this capability, DNS of droplet combustion with full details of spray atomisation could be addressed in detail but new software is required for this challenge. An adaptive mesh refinement (AMR) based code, HAMISH, is currently being developed under the EPSRC Flagship Software Grant (EP/P022286/1) that will provide the necessary capability to undertake this ground-breaking work.

Model situation with 20x allocation: This level of capability would enable a revolutionary change in the status of the simulation of multi-phase turbulent reacting flows. This will enable DNS and LES of industrial combustors and furnaces where spray, coal and biomass combustion will be analysed without any limiting assumptions which are currently used in state-of-art simulations [20,22]. Furthermore, it will be possible utilise Volume of Fluids (VOF) [17,18] and Discrete Elements Method (DEM) [23] in conjunction with chemical processes associated with spray and coal/biomass combustion and simulate multi-phase reacting flows without any limiting assumptions regarding atomisation, particle shape and size distributions. This level of capability will avoid the necessity for 'trial and error' based design methodologies and reduce the empiricism in the analysis method and the need for expensive experimentation. The expected level of capability will make Computational Fluid Dynamics (CFD) simulations not

only an indispensable but also the most reliable tool for the development of new generation energy-efficient, environmentally friendly IC engines, gas turbines and industrial furnace to meet increasingly demanding environmental legislations.

Combustion analysis of biogas and gases derived from coal gasification

Model situation with 5x allocation: Coal-derived gases in the form of syngas, Synthetic Natural Gases (SNGs), Compressed Natural Gases (CNGs), and biogas from sustainable sources, are increasingly going to play an important role in the evolving fuel landscape. Compressed syngas, SNGs and biogases find their applications in light automotive sector as CNGs, which are expected to replace diesel oil in many engineering applications. Thus, modern combustors in the UK (where coal for power generation and petrol and diesel for automotive applications will be phased out by 2025 and 2040 respectively) need to be designed in such a manner so that they can operate on a range of different fuels to take advantage of this evolving fuel landscape. Although fuel-flexibility of combustors opens an opportunity of producing energy in a cleaner form than conventional means, the chemical compositions of syngas and biogas offer different challenges, as the presence of H_2 in syngas resists blowout and promotes flashback, whereas CO_2 in biogas makes it more susceptible to blowout. Thus, a flexible fuel combustor operating on homogeneous fuel-air mixture (which is necessary to control burned gas temperature to limit NOx emission) needs to be optimised with respect to combustor blowout, flashback, auto-ignition and dynamic stability of the combustion processes. With 5 times more allocation than the current one, it will be possible to simulate atmospheric combustion of biogas, syngas, SNGs and CNGs using DNS in canonical configurations with reduced mechanisms using DNS and LES.

Model situation with 10x allocation: The gases derived from coal gasification (e.g. syngas, SNGs and CNGs) are going play increasingly important role in reducing greenhouse gas emission, but the analysis of their combustion needs appropriate modelling of differential diffusion. Most existing models do not account for differential diffusion and thus either the existing models need to be modified or new models need to be developed in order to account for combustion of syngas, SNGs and CNGs. In order to achieve this objective, detailed parametric analyses need to be carried out based on detailed chemistry DNS for a range of different values of turbulence intensity, length scale ratio and H_2/CO_2 blending. This data will subsequently explicitly Reynolds averaged/LES filtered for the purpose of a-priori DNS analyses in order to develop high-fidelity RANS and LES models. These new models will need subsequent a-posteriori validation based on actual RANS and LES studies. In order to achieve this one will need to carry out routine simulations involving 1000 cores.

Model situation with 20x allocation: In most industrial applications combustion takes place under elevated pressure. The flame becomes thin, and turbulent Reynolds number increases with increasing pressure and thus the resolution requirements become increasingly demanding for high pressure values. Thus, three-dimensional DNS and LES of biogases, syngas, SNGs and CNGs at elevated pressures will need of the order of 10,000 cores. The knowledge gained from these simulations will help in the development of next-generation fuel-flexible combustors in process industries, and environment-friendly IC engines.

Flame-wall interaction

Model situation with 5x allocation: The increasing industrial demand for light-weight combustors and micro-combustors makes flame-wall interaction (FWI) an inevitable factor in these applications, which will have implications on energy-efficiency, pollutant emission in IC engines and flashback in the mixing chamber in gas turbines. However, FWI is yet to be analysed in detail both experimentally and numerically because of small length and time scales associated with near-wall flame- and fluid-dynamics. However, current advances in HPC can be utilised to carry out DNS and LES of FWI, and bypassing all the above limitations. Even with 5 times of current computational time allocation, simple or reduced chemistry FWI-DNS for gaseous phase premixed and non-premixed combustion will be limited to simple configurations such as head-on quenching of statistically planar flames and oblique quenching of V-flames for moderate values of turbulent Reynolds number involving typically 1000 processors.

Model situation with 10x allocation: With 10 times of the current computational allocation, it will be possible to carry out DNS and LES of FWI in a range of different configurations using detailed chemistry. Currently, detailed chemistry DNS of FWI of head-on quenching is about 70-80 times more expensive than single-step head-on quenching DNS, and one such case was not possible to run within a reasonable physical time-scale, just relying upon the allocated computational time and the author of this report needed to utilise HPC support from Gauss

Centre of Supercomputing and Leibnitz supercomputing centre with the help of his collaborator Prof. M. Klein based at the University of Bundeswehr München. An AMR based CFD solver called HAMISH, which is currently under development as a part of the EPSRC Flagship Software Grant (EP/P022286/1) [19], will be utilised for production runs of FWI.

Model situation with 20x allocation: With this level of computational power it will not only be possible to carry out high-fidelity simulations of FWI for gaseous phase flames for a range of different fuels (including un-conventional fuels) but also to simulate droplet and coal/biomass particle interactions with combustor walls without any major limiting approximations. This step-change in the computational analysis of turbulent reacting flows will allow for integrating advanced methodologies such as VOF [17,18] (for spray break-up), DEM [23] (for solid particle combustion) and Discrete Ordinates Method (DOM) [24-28] (for radiative heat transfer) with AMR based high-fidelity CFD solver. These computations and associated methodologies will play a pivotal role in the design-cycle of (i) light-weight compact combustors in automotive engines and gas turbines, (ii) micro-combustors for Drones and domestic applications where the high-energy density cannot be supplied by batteries, and also to achieve a reduction in pollutant (such as unburned hydrocarbon) emissions. Thus, the knowledge gained from these simulations will help us to meet ever-growing challenges of maintaining high-energy density without compromising environmental friendliness.

High-pressure combustion

Model situation with 5x allocation: In most industrial applications, such as in IC engines and gas turbines, the combustion processes take place under elevated pressure. However, the flame becomes increasingly thin with increasing pressure. Moreover, the turbulent Reynolds number increases with increasing pressure for a given set of values of root-mean-square velocity fluctuation and integral length scale. Thus, the resolution requirements become increasingly demanding and small grid spacing will be necessary for both DNS/LES of combustion at elevated pressures. Thus, DNS and LES of high-pressure combustion are computationally expensive, so with 5 times of the current computational power will only be sufficient to carry out DNS of high-pressure combustion in conventional canonical configuration only in the presence of simple/reduced chemistry, and LES in simple burner configurations using flamelet methods.

Model situation with 10x allocation: An access to 10 times of the current computational time allocation will enable the UKCTRF to carry out DNS of conventional hydrocarbon fuels using detailed chemical mechanisms and LES of elevated pressure combustion using computationally demanding CMC and PDF methodologies for swirl flows, bluff body stabilised flames and dumped combustors. Most conventional models are proposed for atmospheric conditions so explicit filtering/averaging of DNS data will provide insights into the validity of the existing models under elevated pressure conditions in the context of LES and RANS respectively, and based on this exercise appropriate modifications to the models will be suggested for high-pressure applications. The aforementioned research activities are expected to require of the order of 10000 cores on a routine basis. The author of this report is using his collaborative network with Japanese colleagues (Prof. R. Kurose from Kyoto University and Prof. H. Watanabe from Kyushu University) to simulate high-pressure turbulent spherical expanding H_2-air flames using DNS on K-computer in Japan (http://www.aics.riken.jp/en).

Model situation with 20x allocation: It will be possible to carry out three-dimensional detailed chemistry DNS of conventional premixed and non-premixed combustion at elevated pressure values in more complex configurations (e.g. V-flame, Bunsen burner and bluff body stabilised flame) than the conventional canonical flame in a box configuration if 20 times of the current computational time allocation is available to the consortium users. Moreover, with this computational power, it will be possible to simulate combustion in gas turbine combustors and IC engine cylinders under elevated pressure conditions realised in these applications using LES without any major limiting conditions but these activities will be computationally demanding. This level of capability would enable a step-change in simulating practical engineering combustion problems under engine/gas turbine operating conditions without any major limiting approximations so that design of new combustors are capable of meeting global energy-related challenges.

Impact of the aforementioned research activities:

The vision of UKCTRF is closely aligned with the 'Energy' research theme, and falls under the 'Combustion Engineering' research area of EPSRC's delivery plan. Major beneficiaries are UK based industries (e.g. Rolls Royce, Shell and Siemens etc.) engaged in developing new technologies for low-pollution and high efficiency IC engines (urgently needing the knowledge of new fuel combustion to meet the challenges posed by government's decision to ban Petrol and Diesel cars by 2040) and gas turbines (which will need to be increasingly environment-friendly). The impact of the aforementioned research activities, in terms of new product and wealth creation in the UK, will be felt in a time-scale of 10-20 years because of the long-term nature of the design-cycle of IC engines and gas turbines, and the time required to build confidence in the community. The technological advancements will also help in designing energy-efficient and environment-friendly combustors especially for UK industries, which will bring a long-term benefit (in a time scale of 10-20 years) for society (consistent with Productive and Resilient Nation delivery plans of EPSRC). The research outcomes will also be particularly beneficial for the CFD software community, who use state-of-the-art combustion RANS/LES models in their codes to yield accurate predictions, and this benefit will be realised in 5 -10 years. This research activity will also help in developing a highly skilled UK-based workforce in the form of RAs and PhD students, who will eventually carry the expertise gained in the course of the project in their future roles.

SCIENTIFIC CODES

Mostly used current codes

The codes, which have been used on the national facilities such as HPCx, HECToR and ARCHER, are listed below in Table 1:

Code	Machines
SENGA+ (in-house code for DNS)	HECToR, HPCx, JAGUAR, ARCHER
HAMISH (under development for DNS and LES but will be made open source)	ARCHER
OpenFOAM (open source code for RANS)/LES	HECToR, ARCHER

Table 1: List of major codes used by the UKCTRF members

SENGA+ and OpenFOAM receive EPCC support as a part of on-going UKCTRF activities. Moreover, the EPSRC Flagship Software Grant (EP/P022286/1) supports the development of HAMISH. SENGA+ and OpenFOAM can be run in production mode, and the number of cores used on ARCHER for these codes range from 24 to 20,000 for up to an order of 1,000 hours of wall-clock time. Note, SENGA+ can already scale beyond 50,000 cores as the science requires.

Please note that SENGA+ was a procurement benchmark code for both HECToR and ARCHER. OpenFOAM (and variants such as FireFOAM) is another CFD code commonly used on ARCHER.

SENGA+ is an explicit high-order CFD code with excellent scaling performance, used in many earlier class 1a, 1b and RAP projects. The results shown below were obtained by the Cray Centre of Excellence for two sets of test cases of turbulent combustion using HECToR and Jaguar. A sample speed up table for SENGA+ is listed in Table 2, and the speed up is graphically shown in Fig. 3.

	Runtime (s)				
Cores	Run 1	Run 2	Run 3	Mean	Speedup
48	2746.7	2766.4	2765.7	2759.6	1.00
240	518.9	521.0	523.1	521.0	5.30
480	290.2	271.6	267.4	276.4	9.98

1200	103.0	103.4	101.1	102.5	26.92
3600	35.4	35.0	38.2	36.2	76.23
7200	18.8	19.8	18.4	19.0	145.24
10800	14.2	13.1	12.6	13.3	207.49

Table 2: An example of speedup of SENGA+.

Figure 3 demonstrates nearly perfect scalability for SENGA+ up to 65,536 cores on ARCHER (the simulations were carried out by Jason Beech Brandt from Cray) using a 4-step methane-air reaction mechanism [29].

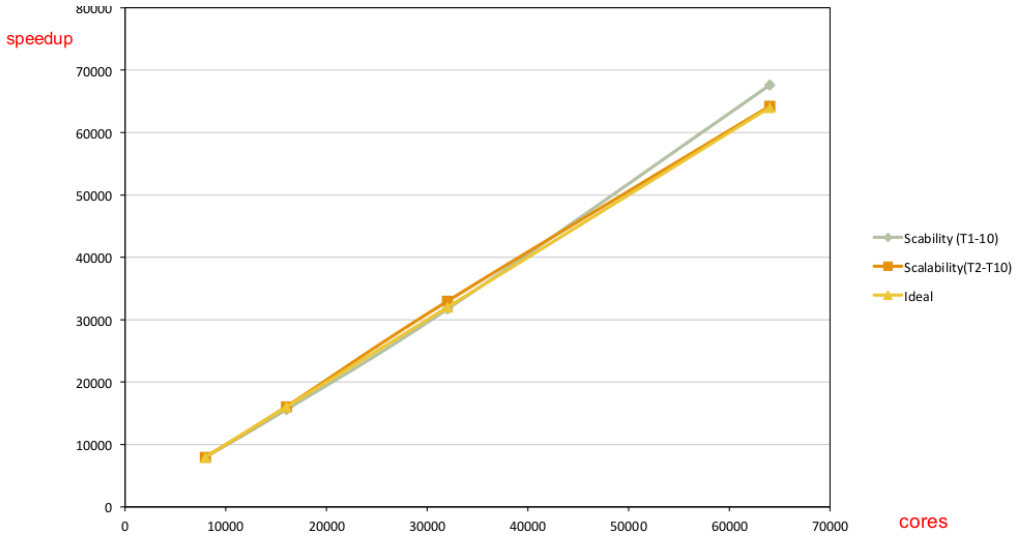

Figure 3: An example of speed up performance of SENGA+. The legends (T1-T10) and (T2-T10) refer to Cray's internal terminology based on their architecture and these legends are not important for the purpose of this document.

OpenFOAM is a state-of-the-art C++, open-source toolbox for CFD, which is also available on ARCHER. This code has been widely used on HPC facilities around the world. The scalability of this code has been tested on ARCHER and the results are shown Fig. 4 for three sets of test runs while investigating turbulent lifted jet flame and a model gas turbine combustor (source: Prof. N. Swaminathan, University of Cambridge).

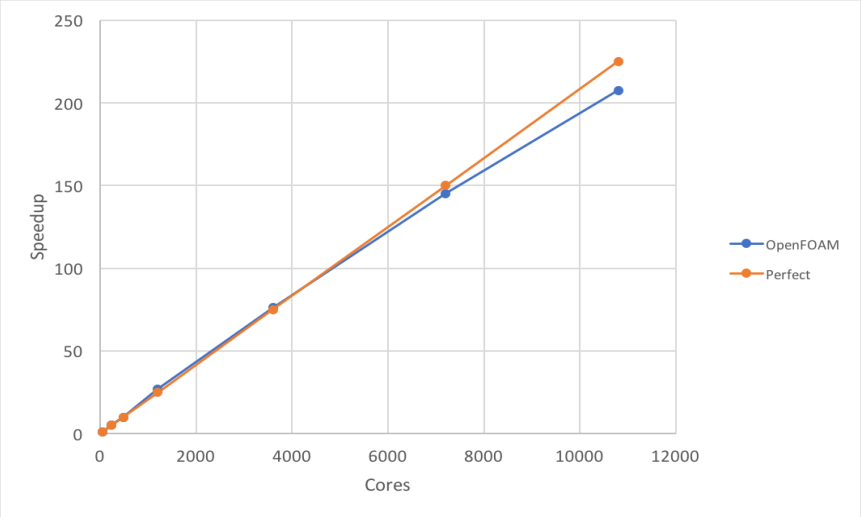

Figure 4: An example of speed up performance of OpenFOAM. The Y-axis has been scaled by 48 for both the simulation and the ideal performance

The HAMISH code has been tested in its MPI version, using a fixed mesh (no adaptive mesh refinement on) on the iDataPlex Cluster (X86_64 IvyBridge E5-2697) of the Hartree Centre (http://community.hartree.stfc.ac.uk/wiki/site/admin/resources.html#idataplex) for the Taylor-Green Vortex case using 1283 cells. This is a relatively small configuration for this type of code. However, linear speed-up can be observed up to 392 cores. Considering each core is busy with only about 5,000 cells, the result is quite good. Note that some speedup is observed up to 1,024 cores, with only 2,000 cells per core. The architecture of the iDataplex Cluster (Blue Wonder Phase2) is very similar to that of ARCHER. Therefore, one can expect similar performance on ARCHER for the same case. As the code is explicit in time (for fixed meshes), one can expect the performance to increase linearly as a function of the number of cells. A sample speed up table for HAMISH is exemplarily provided in Table 3, and the data in Table 3 is graphically shown in Fig. 5.

Cores	Runtime / s	Speedup	Perfect Speedup
8	1632.8	1.00	1.00
16	910.6	1.79	2.00
24	670.5	2.44	3.00
32	529.1	3.09	4.00
48	345.1	4.73	6.00
64	263.2	6.20	8.00
96	184.0	8.87	12.00
128	138.0	11.83	16.00
192	102.9	15.87	24.00
256	84.82	19.26	32.00
392	62.3	26.20	49.00
512	46.3	35.26	64.00
768	42.2	38.71	96.00
1024	41.5	39.38	128.00
1536	46.3	35.23	192.00

Table 3: An example of speedup of HAMISH.

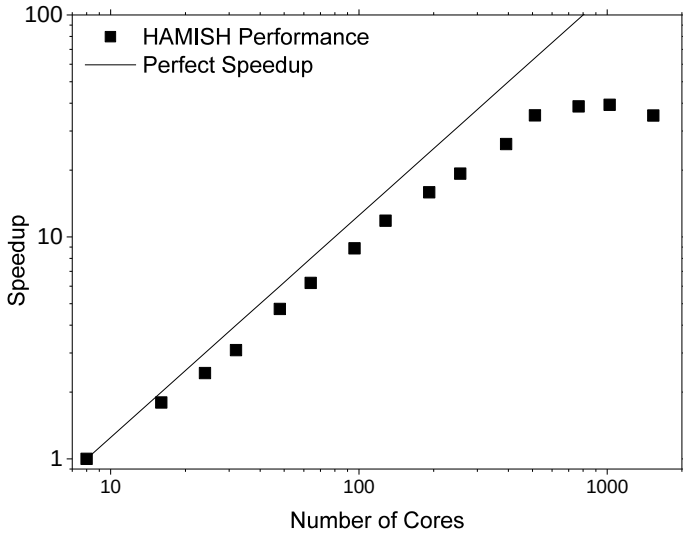

Figure 5: An example of speed up performance of HAMISH.

It is worth noting the runs representative of the kinds, which the consortium is expected to undertake, are considered in above examples to demonstrate the actual scaling in practice.

Emerging code

A next-generation fully parallelised computational fluid dynamics (CFD) code called HAMISH based on adaptive mesh refinement (AMR), is currently under development under the EPSRC Flagship Software Grant (EP/P022286/1). This grant focuses on the development, validation and documentation of HAMISH, which will enable high-fidelity Direct Numerical Simulations (DNS) of advanced turbulent reacting flows such as flame-wall interaction, localised ignition, and droplet combustion including atomisation processes. Such simulations cannot be achieved at present without limiting simplifications due to their prohibitive computational cost. AMR for large-scale highly-parallel simulations of compressible turbulent reacting flows is a significant new functionality which will offer major benefits in terms of computational economy for problems involving thin fluid-mechanical structures, e.g. resolution of both the flame and the boundary layer in flame-wall interaction, droplet surfaces in atomisation in spray combustion, shock waves in localised forced ignition, etc. Such flow features have either been ignored or simplified severely in previous work due to the prohibitive computational cost of fixed global meshes, thus limiting the usefulness of the simulations.
Hence AMR will offer a step-change in capability for the computational analysis of turbulent reacting flows, and will provide data with the degree of detailed physical information which is not currently available from simulations using existing CFD codes.

This new-generation software will be validated with respect to the results obtained from the well-proven uniform-mesh DNS code SENGA2, which has already been ported to ARCHER and is currently widely used by members of the UKCTRF. The newly developed code, HAMISH, will not only be ported to ARCHER, but also be prepared for architectures supporting accelerators thanks to OpenMP 4.5, which will support OpenACC, targeting a POWER8 cluster. As a part of this project, a detailed user guide will be produced at each new release of the code. This user guide will be made available on a website for public download along with the open-source version of the code and the associated documentation on code validation and user tutorials.

The work package 3 of the UKCTRF will focus on the development of the next-generation software that will utilise HPC resources of the future. In recent years, hierarchical forms of hardware (e.g. multicore) have been developed to deliver performance, and extracting this new level of concurrency from these hardware platforms has emerged as a major technical challenge. Moreover, future platforms will involve multicore CPUs with floating point accelerators provided by Graphical Processing Units (GPUs). Thus, the HPC users of this consortium need to

respond to this challenge by designing and re-tooling algorithms so that the complex memory hierarchy of these new computing platforms can be addressed. This WP enables consortium codes to successfully adapt to the evolving hardware and take advantage of future high-end computing. This activity is relevant for ARCHER and its successor, IBM's Blue Joule, located at the Hartree Centre, and HPC platforms available through PRACE.

REFERENCES

[1] Miller, R.S., Bellan, J. (1999) J. Fluid Mech., 384, 293-338.
[2] Reveillon, J., Vervisch, L. (2000) Combust. Flame, 121, 75-90.
[3] Wang, Y., Rutland, C. J. (2005) Proc. Combust. Institit.,30,,893-900.
[4] Wang, Y., Rutland, C. J. (2005) Combust. Flame, 149, 353-365.
[5] Reveillon, J., Demoulin, F.X. (2007) Proc. Combust. Institit., 31,, 2319-2326.
[6] Reveillon, J., Demoulin. F.X. (2007) J. Fluid Mech., 583, 273-302.
[7] Schroll, P, Wandel, A.P., Cant, R.S. and Mastorakos, E. (2009) Proc. Combust. Inst., 32, 2275-2282.
[8] Wandel, A., Chakraborty, N., Mastorakos, E. (2009), Proc. Combust. Institit., 32, 2283-2290.
[9] Xia, J., Luo, K. H. (2010) Flow Turb. Combust., 84, 397-422.
[10] Neophytou, A., Mastorakos, E., Cant, R.S. (2010) Combust. Flame, 157, 1071-1087.
[11] Neophytou, A., Mastorakos, E., Cant, R.S. (2011) Proc. Combust. Inst., 33, 2135-2142.
[12] Luo, K., Pitsch, H., Pai, M.G., Desjardins, O. (2011) Proc. Combust. Inst., 33, 2143-2152.
[13] Neophytou, A., Mastorakos, E., Cant, R.S. (2012) Combust. Flame, 159, 641-664.
[14] Wacks, D., Chakraborty, N., Mastorakos, E. (2016) Flow, Turb. Combust, 96, 573-607.
[15] Wacks, D., Chakraborty, N., Flow Turb. Combust., 96, 1053-1081.
[16] Wacks, D., Chakraborty, N. (2016) D.H. Wacks, N. Chakraborty, Combust. Sci. Technol., 188, 2149-2177.
[17] Fulgosi, M., Lakehal, D., Banerjee, S., De Angelis, V. (2004) J. Fluid Mech., 482, 319–345.
[18] Klein, M. (2005) Int. J. Heat and Fluid flow, 26, 722-731.
[19] http://www.ukctrf.com/flagship-software-grant/
[20] Brosh T, Chakraborty N. (2014) Energy & Fuels, 28(9), 6077-6088.
[21] Brosh T, Patel D, Wacks D, Chakraborty N. (2015) Fuel, 145, 50-62.
[22] Luo K, Wang H, Fan J, Yi F. (2012) Energy & Fuels, 26,6128-6136.
[23] Kawaguchi, T., Tanaka, T., Tsuji, Y. (2015) Powder Tech. 96, 129-138.
[24] Modest, M. F. Radiative heat transfer. Academic Press, 2013.
[25] Franchetti BM, Cavallo MF, Navarro-Martinez S, Kempf AM. (2013) Proc Combust. Inst 34,2419-2427.
[26] Rabacal, M., Franchetti, B.M., Cavallo Marincola, F., Proch F., Costa, M., Hasse, C., Kempf, A.M. (2015) Proc. Combust. Inst. 35. 3609-3617.
[27] Kurose, R., Makino, H., Suzuki, A., (2004) Fuel , 83, 693–703.
[28] Chui, E.H., Hughes, P., Raithby,G. (1993) Combust. Sci. Technol., 92, 225–242.
[29] Peters, N., Williams, F.A. (1987) Combust. Flame, 68, 185-207.