Response of the UK Consortium on Turbulent Reacting Flows (UKCTRF) to the questions for "The Future of Compute Review – request for evidence on exascale computing" initiated by the Department of Digital, Culture Media and Sport (DCMS)

1. ***Key drivers (sectors & applications)***

a. Who would use the system in your research discipline, and what would they use it for? Please provide details on specific use cases

b. Who would the early adopters be?

c. What use cases can we expect to emerge over the coming decade- ie. what area of science isn't possible now but will be? What will this contribute to policy priorities?

d. Must the UK have its own system to do this work? With reference to the Own, Access, Collaborate framework, would access to international systems meet this demand?

The world energy demand is predicted to increase by 28% by 2035 [R1]. Importantly, all energy outlooks acknowledge that combustion (including hydrogen, ammonia, and biofuels) will play an essential role in the energy sector, especially [R1-R4]. Thus, the demand for high-fidelity computational methodologies for turbulent reacting flows will increase significantly to offer the improved physical understanding needed, which will be essential for the design of new generation combustors that can take advantage of the evolving fuel landscape that will address the challenges of energy efficiency and pollution control faced in the UK and worldwide. The UK is home to major manufacturers of energy conversion devices, employing over 1M UK citizens and contributing to ~20% of UK exports. The number of vehicles/aircrafts is anticipated to nearly double in 20 years with approximately 75% of them relying on some form of combustion [R1]. The technological advances in computational simulation and modelling will help in designing energy-efficient and environmentally friendly combustors, bringing a long-term benefit to society, which is consistent with the UK Government's "Road to Zero" strategy [R5]. However, this goal cannot be achieved without the advancements in the computational analysis of turbulent reacting flows, which is indispensable for addressing the challenges posed by new generation carbon-neutral fuels such as hydrogen and ammonia for utilisation in power generation and propulsion without compromising energy efficiency and pollutant (e.g. NOx) emission.

In turbulent reacting flows, complex interactions occur between chemical reactions, turbulence, heat transfer, radiation, fuel atomisation, and vaporisation, which are further influenced by the separation of length and time scales. These problems exhibit a variety of multi-physics phenomena that are at the forefront and often beyond experimental capability. The UK Consortium on Turbulent Reacting Flows (UKCTRF) community can perform Direct Numerical Simulations (DNS) that can resolve all length and time scales, without physical approximations, to produce the equivalent of a time-resolved experiment. Moreover, the fundamental insights obtained from DNS data are already used by UKCTRF to develop high-fidelity models for Reynolds Averaged Navier-Stokes (RANS) and Large Eddy Simulations (LES), which in turn have industrial applications (e.g. Rolls Royce, Siemens, etc.) in developing future generation energy efficient and environmentally friendly combustors. Furthermore, the high-fidelity simulation data will be utilised to train Machine Learning algorithms to inform future model developments, identification of early signs of explosions and combustion instability, and inform the design of experiments and associated diagnostics. However, to realise the aforementioned capabilities and maintain UK's world-leading status in turbulent reacting flow research requires access to substantial amounts of computational time, far beyond what is provided through the existing grant-based allocations. This is summarised in Table 1, shown below.

Prepared by Prof. Nilanjan Chakraborty (Nilanjan.chakraborty@ncl.ac.uk), Newcastle University & Prof. David Emerson (david.emerson@stfc.ac.uk), STFC, Daresbury Laboratory

	Timescale (~ 2022)	Timescale (~ 2024)	Timescale (~ 2025/26)
Sample problems to tackle	Simulation & modelling of single-phase and multiphase turbulent reacting flows involving conventional fuels and high-Hydrogen content fuels in canonical configurations under atmospheric or slightly elevated pressures for the simulations with the highest fidelity (i.e. Direct Numerical Simulations (DNS)). Most industrial simulations are currently based on Reynolds Averaged Navier-Stokes (RANS) methodologies, which are heavily dependent on the accuracy of strong assumptions (i.e. models) but they provide useful (mostly qualitative) design inputs. More accurate methodologies for large-scale unsteadiness, such as Large Eddy Simulations (LES) are becoming increasingly popular, but they are yet to be routine calculations for engineering design purposes.	The research themes in the left-hand column but in laboratory-scale geometries for DNS involving especially hydrogen-based fuels (e.g. H2-air, Ammonia-air), e-fuels, biofuels, and low calorific fuels derived from coal gasification at high-pressures involving multi-scale multi-physics phenomena (e.g. flame-wall interaction, phase change, plasma ignition, etc.). LES is expected to be more routine than the current situation and will be equipped with high-fidelity models developed based on improved physical insights extracted from DNS data. These simulations will be needed to design new generation combustors for power-generation and propulsion sectors, which are crucially needed to switch from the conventional energy generation mode to new carbon-free/neutral fuels to meet net-zero targets.	Systems-level high-fidelity simulations (e.g. DNS and LES) for high-pressure combustion of alternative carbon-free/neutral fuels, e-fuels involving multi-scale multi-physics phenomena (e.g. flame-wall interaction, phase change, plasma ignition, etc.). At this stage, the accuracy simulation data will be able to replace the need for expensive experiments at least for some configurations. The simulation data will be accurate enough to train Artificial Intelligence and Machine Learning techniques to construct digital twins of actual industrial combustors in both power and propulsion sectors so that the decisions regarding operating conditions can be optimised based on the requirements of energy efficiency, environment regulations, net-zero targets without expensive experimentation. This will inform future government energy policies based on scientific evidence.
Grid points needed	10^{10}	$10^{12\sim13}$	$10^{14\sim16}$
Supercomputing requirements	100 Pflop/s	300 Pflop/s	1.0 Eflop/s

The flagship codes of the UKCTRF are currently going through transformations to make them ready for future exascale infrastructure and several management team members of the UKCTRF are involved in the ongoing EXCALIBUR programme in different capacities. Therefore, a large section of the

UKCTRF community, especially those engaged in DNS and large-scale LES research activities, will actively engage with an exascale system and early career researchers (e.g., postgraduate research students and research associates) under the supervision of the UKCTRF members will be early users of the UK's exascale computing facility. This will eventually give rise to the development of highly skilled UK-based personnel who will contribute positively to this country's industry and academia, which will be beneficial for the UK's economy and innovation, along with the socio-economic benefits that the computational research on turbulent reacting flows will offer in terms of carbon neutrality and environmental friendliness.

The international reviewers of the exascale Science case commissioned by UKRI in 2019 [R6] commented that the UKCTRF community is internationally leading, and this position has been achieved by innovations in terms of software development for simulations and postprocessing of simulation results to extract important physical information. Although some of the members in the UKCTRF community can access exascale systems (e.g. Fugaku, Frontier) through their ongoing collaboration with Japan and USA-based colleagues, the UK-based software development will not be able to sustain its world-leading position in the absence of a UK-based exascale infrastructure. Moreover, a UK-based exascale system will offer the UKCTRF community the opportunity to engage in the design of pioneering flagship simulations, which are yet to be attempted. These are relatively more difficult to achieve on an overseas platform accessed through collaboration with international colleagues. Finally, an absence of a UK-based exascale system will have a severely detrimental effect on the development of the future generation of Research Software Engineers (RSEs) and the highly skilled workforce that was mentioned earlier.

2. Infrastructure

What kind of exascale system is needed in your research area? What are the requirements for:

- ***hardware***
- ***software***
- ***architectures***

UK's exascale system will likely have a hybrid architecture, at least from the point of view of power requirements. Hybrid infrastructure is also desirable for the community because most flagship codes are designed for CPU-based systems and the usage of Graphical Processing Units (GPUs) for large-scale fluid dynamics problems is yet to be optimised and there are uncertainties regarding future support of GPU programming languages. However, all flagship codes are currently going through a transformation process to take advantage of hybrid platforms and the advantages offered by floating-point accelerators provided by GPUs. Non-intrusive language directives (OpenMP 4.5+ and OpenACC) could play a vital role in easing the transition to exploiting hybrid architectures. In addition, vectorisation to exploit multicore, manycore, and GPU-based accelerators for discretisation techniques and improvement of Input/Output (I/O) performance for large-scale execution for the improvement of overall performance on large core counts, are critically required. The current ExCALIBUR activities, where the UKCTRF members are involved, include reengineering of flagships using Domain Specific Languages, and generic activities related to the implementation of parallel I/O, code coupling, and uncertainly quantification for future Exascale systems. These ongoing activities are expected to yield results and new capabilities will be ready to demonstrate scientific use cases within 3 years.

3 Preparedness

a. How ready is the academic/research community for exascale? (ie. what work is able to use GPU-enabled codes? How long before these workflows can be scaled/are ready?)

b. What are the advantages/disadvantages of a phased approach to exascale?

c. What can government do to help enhance user preparedness for exascale?

The readers are referred to the response to the above question regarding the readiness of the UKCTRF research community. It takes time to (i) reengineer the codes to make them exascale-ready, (ii) induce a culture change in the community, industry and academia, and (iii) produce enough skillsets in the form of RSEs to take full advantage of an exascale system. The current activities under the umbrella of ExCALIBUR are addressing the aforementioned challenges in anticipation of a UK-based exascale system around 2025 as a successor of ARCHER2. In this respect, a phased approach has already been taken and the current activities in anticipation of an exascale facility will give rise to a wide range of innovations in terms of software developments, which will eventually translate to the subject domain-specific advancements. It will be beneficial to have provisions for follow-up funding once the current ExCALIBUR programme is over to not only sustain the current investment and development activities, so that they can be utilised to demonstrate their capabilities in high-priority use cases on exascale platforms, but also look to accelerate the move towards exascale computing through additional investment in researchers, RSEs, and upskilling industry through training and knowledge transfer. This requires government funding for a future exascale machine as a successor of ARCHER2 and commitment of funds which will enable UKRI, Met Office and UAEA to sustain the innovations which are currently in the nascent stage and bring them to fruition. The government funding is also necessary to have a sufficient number of Tier-2 and Tier-1 machines in the UK so the future exascale machine is used only for large-scale parallel computations which merit exascale computing.

4. Investment

a. What are the economic benefits of exascale?

b. What are the key interactions government would need to take an interest in/invest in?

The economic benefits of exascale, which are relevant to the UKCTRF community are already discussed in detail in the first paragraph of the response to the first question and interested readers are referred to that response for part (a). The key interactions of government should include increased fund allocation to UKRI's e-infrastructure programme to invest in the development of highly skilled personnel (RSEs, PhD students, Research Associates from diverse backgrounds) and funding for an exascale machine. More UKRI funding opportunities for the science use cases enabled by exascale computing will accelerate innovation and offer socio-economic benefits to UK society. A follow-up of the SPF funding of the kind of ExCALIBUR programme will also be immensely beneficial for computational research in the UK.

5. Environmental sustainability

Please comment on how important environmental sustainability of such a facility would be for your research community, and what considerations you would like to see incorporated into design and operations of such a facility.

The users of the community need to be educated about the carbon footprint of their simulations. They also need to understand that some choices in chemical and turbulence models can make the computation more expensive (and therefore more carbon footprint) without offering additional benefits in terms of

Prepared by Prof. Nilanjan Chakraborty (Nilanjan.chakraborty@ncl.ac.uk), Newcastle University & Prof. David Emerson (david.emerson@stfc.ac.uk), STFC, Daresbury Laboratory

physical information. This education and culture change will need sustained effort on the part of the consortia chairs, CCP chairs and knowledge exchange coordinators. EPCC has already done some work on CO_2 emission for some flagship codes of the UKCTRF and there is scope for improvement in this regard by profiling operations on the codes. It will also be useful if notional CO_2 emission is mentioned in the technical assessment form (even for ARCHER2) to raise the awareness level. The alteration of job scheduling and incentives depending on carbon footprint and performance of certain codes (e.g. running it at a time when the cooling requirements are not high, e.g. night time or winter) by the exascale system will help reduce the carbon footprint of such a facility in the future.

6. *International*

a. Is there an international/competitive advantage in having such a system? Please specify the advantage to be gained by owning an exascale system.

b. Does the UK's lack of an exascale system affect our international standing? If so, how?

Exascale scientific computing will be critical for the propulsion and power generation sectors of the UK to meet the net-zero and other pollution control targets and also to retain and reinforce UK's leading position in combustion research. The lack of sufficient supercomputing provision in the UK is placing the UKCTRF community at a competitive disadvantage relative to international partners. Several countries (e.g, China, Japan, USA, Germany France, etc.), including those which did not traditionally have a track record of computational research (e.g. Saudi Arabia), have recently invested heavily in supercomputing infrastructure. The supercomputers based at King Abdullah University of Science and Technology, Saudi Arabia (Saudi Aramco, which is 18^{th} in June 2022 Top500 list), and Tokyo University (20^{th} in the June 2022 Top500 list, but not one of the top 10 fastest systems in Japan) were faster ARCHER2 and these are only two examples out of several possible ones. Although ARCHER2 is currently the 25^{th} fastest supercomputer in the world, this resource is shared by several communities and the share of the UKCTRF is rather modest in this respect. Moreover, the number of supercomputing facilities that the UKCTRF community can access in the UK is much smaller than in China, Japan and USA, where several university-owned facilities are comparable to ARCHER2. This has implications for the quantity and quality of research outputs and the level of innovation. For example, DNS of turbulent reacting flows involving 35 million CPU hours on 10,000 cores were hero simulations in 2009 [R7] but can now be routinely done by some US and Japanese groups (e.g. the same group in [R7] reported simulations involving 50 million CPU hours on 72,000 cores in a recent publication [R8]). These simulations can play significant roles in offering fundamental physical information which can be utilised for high-fidelity model developments that can play key roles in the development of new generation propulsion and power generation devices for the net-zero economy. However, these simulations are yet to be routine research in the UK and remain hero simulations for UK researchers. Many of our international competitors will deploy pre-exascale and exascale systems over the next two years, and without such an ambition the UK risks becoming unattractive as a destination for researchers in computational turbulent reacting flows and losing leadership in fields in which the UK has traditionally been a key global leader. The access to PRACE facilities is going to be associated with EuroHPC membership, and thus the access to European computing facilities for UK researchers is expected to end following the departure of the UK from the EU. Therefore, UK-led computational turbulent reacting flow projects will not be able to make use of the pre-exascale or exascale systems which will be deployed across Europe under the auspices of EuroHPC. Investments in UK-based exascale computing provisions along with the investments in code refactoring for exascale platforms and training of personnel to utilise exascale computing will enable the UKCTRF community to remain internationally competitive and address increasingly challenging multiphysics, multiscale turbulent reacting flow problems. This will accelerate innovations and technological step changes, with positive impacts on energy efficiency, environmental friendliness, wealth creation, industrial competitiveness and the development of highly skilled personnel in the UK.

References

R1. International Energy Agency (2015), "World Energy Outlook 2015".
R2. Committee on Climate Change (2018), "Reducing UK Emissions: 2018 Progress Report to Parliament".
R3. EUCAR (2013) "Research & Innovation Roadmaps for EUCAR Strategic Pillar: Sustainable Propulsion".
R4. Department of Energy, US (2015), "Annual Energy Outlook 2015 with Projections to 2035", DOE/EIA-0383.
R5. 'Road to Zero' strategy of the UK government
R6. UKRI: Science case for UK supercomputing
R7. J. H. Chen, et al., 2009, Comput. Sci. Discov., 2, 015001.
R8. A. Gruber, et al., 2018, Phys. Rev. F, 3, 110507.

Prepared by Prof. Nilanjan Chakraborty (Nilanjan.chakraborty@ncl.ac.uk), Newcastle University & Prof. David Emerson (david.emerson@stfc.ac.uk), STFC, Daresbury Laboratory