HPC solutions for future demands in Oil & Gas, Wind Energy and Transportation Powertrain industries

WHITE PAPER

ENERXICO

INTRODUCTION

According to the Paris Agreement adopted in December 2015 (UNFCC, 2015) [1], the global energy demand is expected to experience a large growth over the 28 years from 2012 to 2040. However, energy scarcity or inefficient usage can lead to higher prices, which will have a critical impact on the economy, as emphasized by the Energy Challenge in the Horizon 2020 work program.

Specifically in Mexico, the energy independence index, which shows the ratio between national energy production and consumption, was equivalent to 0.72 in 2019 [2]. This result implies that more energy was consumed than was produced domestically. In this context, energy self-sufficiency is essential for the development of the industry in the country, which will have to face the new challenges of productivity and economic growth.

During the last years, High-Performance Computing (HPC) resources have undergone a dramatic transformation, with an explosion in the available parallelism and the use of special-purpose processors. The number of cores in an ExaFLOPS (10^{18} FLOPS) computer will be in the order of 100 million. This imposes severe pressure to increase the parallel efficiency of the applications. At the same time, this massive parallelism opens new opportunities to increase accuracy when simulating physical phenomena.

Due to their processing capabilities, supercomputers allow scientists to run more advanced simulations and data analyses applied to different energy sectors and help them bring these technologies quickly to the market. For instance, nowadays many energy companies are using HPC solutions to look for new sources of energy, such as oil.

In particular, during the COVID-19 pandemic, many industry sectors were negatively affected. Overall, the global energy demand fell by 5% in 2020 [3]. This was also the case of fossil fuels, which were less used and whose financing fell 9% the same year [4]. However, at the same time, the pandemic accelerated renewable energy usage. For instance, in the US the consumption of this energy increased by 40% in the first ten weeks of lockdown [5]. Therefore, even though it is clear that the production of fossil fuels requires a large infrastructure investment and it is difficult to have it in times of crisis, COVID-19 can probably reinforce the integration of renewable energy sources [6]. This situation strengthens the need to have stable sources and increase efficiency to face negative pandemic effects. In this context, HPC allows taking

advantage and adapting the technologies to the market to meet future demands more rapidly as well as not wasting energy sources.

In the framework of the EU and Mexico collaboration project ENERXICO, exascale HPC techniques have been applied to different energy industry simulations of critical interest in Mexico. The project has brought together the country's main stakeholders in the energy industry and European energy companies working in the Mexican market to develop beyond state-of-the-art high-performance simulation tools to give solutions for these very important industrial sectors: oil & gas, wind energy, and propulsion & power generation.

On the one hand, the research about oil & gas has given an extraordinary opportunity for the use of exascale technology in the industry. The scalability of the research's outcome allows the expansion of the developed techniques to areas of carbon sequestration and the development of low carbon energies. On the other hand, ENERXICO has also researched novel wind energy design tools to understand and predict atmospheric scales of motion for the operation and performance of wind turbines and farms in complex wind situations. Finally, the project has focused on the development and validation of predictive combustion simulation tools as well, with special attention to renewable fuels to optimise fuel design and fuel performance towards more sustainable and greener propulsion systems. The results help to provide a further understanding of properties that can assist in the decarbonisation in the transportation sector and achieve the best performance and minimal environmental impact.

In summary, ENERXICO has helped the different energy sectors by giving solutions for the modernisation of the whole Mexican energy industry and did its bit to improve the cooperation between industries from the EU and Mexico as well as leading research groups in both countries.

ENERXICO

Developing state-of-the-art HPC solutions for future demands in Oil & Gas, Wind Energy and Transportation Powertrain industries

COLLABORATION & COOPERATION BETWEEN EU AND MEXICO IN THREE AREAS:

Oil & Gas

- Using geophysical exploration for subsalt reservoirs to increase the oil & gas reserves.
- Carrying out massive data simulations for exascale computers for efficient exploration, production transport and refining of oil & gas.
- Exploiting in-silico models of new catalysts for crude oil refinement to reduce the dependence on traditional laboratory experiments.

Wind Energy

- Using exascale HPC platforms to improve the physical fidelity of the models used in wind farm design.
- Specifying important external parameters in complex flow situations to classify relevant wind conditions for wind turbine design.
- Assessing the developed meso and micro-scale coupling techniques in complex flow.

Transportation industries

- Developing predictive combustion simulation tools to optimise fuel design and performance to make more sustainable and greener propulsion systems.
- Using HPC to provide further understanding on the combustion process and emissions characteristics of new fuels.
- Establishing combustion and emission models for new industrial guidelines for engine operation and maintenance.

HPC IN OIL & GAS INDUSTRY

The business of subsurface exploration relies on determining, with a precision as high as possible, the location of new hydrocarbon reservoirs. Although several decades ago it was reasonable to just perforate in areas where oil was abundant and try your chances, the more these "easy oil" regions were exploited the less that strategy worked. Essentially in the '70s finding new oil became only possible employing geophysical exploration. The underlying principles were simple at first. By emitting seismic waves from the surface and listening to echoed waves, or reflected waves, several layers could be identified.

Such a data-driven strategy had its shortcomings. On the one hand, reflector locations at depth were difficult to identify because seismic velocities are not constant in the subsurface. Secondly, in areas where the geology was complex, the attained interpretation would fail. Both of the problems could be mitigated by employing different "offsets" for each emitter in the survey. By having a collection of receivers at different distances from the emitter, waves took different paths before and after being reflected on the surface. A careful, geometrical analysis of such differences in travel times reduced ambiguities. Essentially, by adding more data and some computing the subsurface could be better imaged. The trend towards more data augmented considerably to embrace 3D data acquisition in the 80's with several azimuths and more complex data processing workflows be-

coming standard in the early 2000's. Nevertheless, classical geophysical algorithms relied on strong assumptions, regarding the distribution of velocity in the subsurface or the flatness of reflectors. Too deep or too complex reservoirs were harder to characterize, and deep drilling became more and more expensive at the same time.

To circumvent most limitations in geophysics several algorithms were implemented efficiently so that they could be applied to geophysical surveys. This included Kirchoff prestack depth migration (PSDM) and reverse time migration (RTM) [7]. Both algorithms could map precisely reflectors in the subsurface but the computational cost associated was very high to be handled by regular workstations. In the first case, seismic waves were considered as bending rays that could interact with all subsurfa-

ce structures. As this had to be done for each source-receiver pair, the amount of ray-tracing was huge. PSDM already made a significant impact on the capacity of subsurface imaging. On the other hand, RTM started to make an impact in the last half of the 2000's. Instead of using an infinite-frequency approach (rays) for the seismic signals, it used the complete physical modelling of waves in 2D/3D media. By using such a rich physics-based kinematic modelling, geological models of almost arbitrary complexity could be imaged, as long as the underlying velocity models were sufficiently good. PSDM and RTM, if anything, marked the head start for the revolutionary uptake of HPC within the oil and gas industry. Since then several O&G companies have regularly appeared at the top500 list [8], which showcases the most powerful supercomputers in the world. Where many would have expected the aerospace, automotive, or big tech sectors to be the stars in such lists, it has been O&G systems, devoted mostly to seismic imaging, the lead industry contributions in a list that is mostly filled with state-owned supercomputers devoted to public R&D. Far from being an anomaly in time, still in 2020 we could count 9 O&G HPC systems among the world's 100 fastest computers (e.g. Eni, Total, Petrobras, PGS).

The reason for HPC's prevalence in O&G can be attributed to the continuous development of novel imaging techniques which have kept making good use of supercomputing. In particular, least-squares RTM and very especially full-waveform inversion (FWI) have had a transformative impact in the last decade [9,10].

Full Waveform Inversion enables the possibility of obtaining background velocity models by employing waveform modelling in inverse mode. This has a double advantage for the O&G industry. On the one hand, it tackles a necessity in imaging, for example for RTM, where good models result in better images.

On the other hand, it is algorithmically similar to RTM, therefore many developments related to the optimization and improvement of RTM, attained in years of development, can be readily applied to FWI. However, FWI is not devoid of limitations and issues. Firstly, it took a long time to adapt the algorithms to datasets different from first-arrivals, which excluded several data acquisition types from inversion and also limited the impact of the technology in its early days. Secondly, to converge, FWI needs very low frequencies which were not customary in O&G surveys before FWI. Thus, seismic surveys adapted to the algorithmic requirements by recording longer offsets and lower frequencies.

Finally, yet importantly, being an optimization problem, many strategies can be used to improve inversions, some of them using different cost functions, others using more complex physical models (e.g. anisotropic, elastic), and others trying to bridge the gaps between "background model" and "reflectivity model". The success of geophysicists in solving such problems will mark the continued relationship of HPC and O&G in the future, where novel technologies such as machine learning pose the possibility of either assisting, complementing, or perhaps even replacing several current geophysical technologies in the coming future. In any case, it is through R&D both in geophysics and HPC that the imaging problem can be further addressed and solved in years to come.

HPC IN WIND ENERGY INDUSTRY

The commercial development of wind energy in the last decades has been driven by continuous incremental improvements in all aspects of design, production, installation, and exploitation of wind turbines, leading to current-day turbines of up to 8 MW and 220 m height. The increase in size and height of modern wind turbines has challenged our current knowledge of the wind turbine operation conditions because the wind conditions change dramatically from 50-100 to 100-300 meters above the ground.

Compared to our current understanding of turbine technology itself, the understanding of the wind conditions to which large wind turbines are exposed is still lacking. This is due to the broad multiscale nature of wind variations, and their interaction with the environment. This knowledge gap has significant implications for the design and construction of wind turbines as well as the operation, maintenance, and power production during the lifetime of these machines. Addressing this critical knowledge gap is consequently a vital contribution to continue the reduction of wind energy costs that we have seen over the last decades. Hence, it will enable the EU to deliver on its renewable energy target of at least 27% of energy consumption and a commitment to continue reducing greenhou-

se gas emissions, setting a reduction target of 40% by 2030 relative to 1990 levels.

A possible enabler is the use of exascale HPC platforms where the massive increase of the computational power could be used to dramatically improve the physical fidelity of the models used in wind farm design. Hereafter, a possible exascale application for the improvement of wind farm designs is discussed along with some of the main algorithmic bottlenecks to solve before.

A possible application is to consider wind power forecasts at different time scales, from short-term (meso-micro dynamical or statistical downscaling), to seasonal or decadal using glo-

bal climate models. A feasible direction in this area is to use the disruptive coupling methodologies between meso-scale and microscale models developed in ENERXICO project using Alya (BSC CFD code) [11,18], where both HPC related challenges and numerical/modelling aspects have been taken into consideration. Here a coupling approach that uses mesoscale tendencies [12], where the mesoscale coupling data from ERA5 or WRF is obtained by extracting mesoscale budget terms, has been developed. The microscale model code reads the horizontal-averaged data, which is interpolated and added as source terms to the microscalar momentum equations [13]. Periodic domains are used to develop local turbulence driven by the mesoscale forcing terms. At the microscale level, LES equations are discretized using a finite-element non-stabilized formulation, with a non-incremental fractional step method to stabilize pressure, allowing equal order interpolation for pressure and velocity. This non-dissipative formulation, using an energy-conserving discretization, highly enhances the accuracy of the LES [14]. The turbulent viscosity is closed in terms of the Vreman subgrid-scale model [15]. The following strategy has been tested in the context of the ALEX17 experiment [16]. This experiment was conducted to characterize the wind conditions across the Elortz valley to the north of the Alaiz Test Site for the validation of flow models.

From the intensive operational period, a case study has been selected for a Wakebench benchmark consisting of a series of diurnal cycles with relatively persistent winds from the North. The validation is centered around a 118-m mast at the test site, with six 80-m met masts in the valley and a 10-km long Z-transect constructed from five long-range Wind Scanners, at a constant height of 125 m height above ground level. WRF (a pure mesoscale model) can reproduce correctly the big scale features of the experimental signal, it is expected that with the

microscale coupling the high frequencies coming from the terrain will be added into the numerical solution. Four days of the experimental campaign are simulated in the present paper. As an example, for a microscale domain of 16 km x 16 km with a vertical resolution of 10 m and tangential resolution of 35 m, the coupled WRF+LES simulation presents a mean normalised BIAS vs the observations of less than 4%, which is a big improvement from the pure WRF simulations with a typical BIAS of 30% (using 3km resolution meshes and domains of 264km x 264km) [16]. The improvement is very much related to the superior capability of the coupled meso-micro scale model to introduce the high frequencies coming from the terrain to the overall mesoscale solution. However, to achieve this simulation 2000 CPUs are used and 2h of CPU time is needed to obtain 1h of simulation (hence 8 days of CPU would be needed to obtain the 4 days of the experimental campaign). Given that meshes of 20 Million nodes are used, the limit of the strong speedup is already achieved in about 4000 CPUs, therefore there are not good ways to improve the time to solution. This means that to get a one-year simulation basically, the user needs to wait one year!

A possible solution is to apply a par-in-time strategy, where the meso-scale data is grouped in 2 day blocks (one to spin up the local turbulence and the other is the actual simulation day), this strategy is also convenient because it minimizes the possibility of model drifting between the mico-scale and meso-scale domains.

One could run the 365 days of one year in 365 blocks with a cost of 2 days with about 1.5M CPUs (this is 70 M CPU hours). This is an application that is more likely to be implemented in an exascale machine, where 70M CPU hours could be available to use in just a couple of days. In current Petascale machines, these types of jobs are usually implemented in HPC projects of one year of duration. Therefore, the proposed strategy could enable annual wind resource assessments using meso-micro scale coupling in upcoming exascale machines, and if the ENER-XICO experiment (using just 4 days) holds, bias with observations could be significantly reduced (from 30% to 4% in extreme conditions).

However, to be able to run such calculations effectively in an upcoming exascale platform, the wind energy meso and micro-scale codes will need to improve in several aspects. Hereafter we focus on the micro-scale part and BSC CFD code Alya. Similar actions will be needed in other CFD packages; therefore, the discussion is quite general. To take advantage of exascale HPC platforms, the Alya code is undergoing a huge transformation. To solve incompressible flow problems using an explicit treatment of the momentum equation, the two main kernels are the assembly of the right-hand side (RHS) term for the momentum equation and the solution of a pressure Poisson equation. Improvements in these two main kernels need to be also accompanied by improvements in fault-tolerance, input-output, in-situ visualization, and even artificial intelligence to make the best usage of exascale hardware.

Regarding the assembly of the RHS term, we have been working with a team of experts in node-level optimization from Friedrich-Alexander-Universität (FAU) to analyze and improve the code using their Likwid tool [17]. For the CPU implementation, they found that the RHS

assembly is not memory bound as one would initially expect for an unstructured grid code. This, and the fact that using AVX512, the CPUs run at about 25% of their peak floating-point throughput, indicates that Alya's implementation is already highly efficient. The FAU team has been able to find some slight code modifications to help the compiler that provides a 20% speed-up.

For the GPU implementation, the behaviour is not that satisfactory. The RHS assembly is strongly memory-bound. Moreover, the code exhibits low occupancy and low memory parallelism. This means that there is still work to be done. A mini-app that replicates the behavior of the RHS assembly has been created to facilitate the optimization process. After some optimization, the timing for the mini-app was reduced from 612ms to 175ms. Moreover, with some semi-automatic conversion to CUDA, the timing was further reduced to 53ms. This is a more than ten times reduction from the original timing. If these improvements can be translated to source code of Alya, the GPU version will be in a very competitive position.

Considering that most of Alya can run either in CPUs or GPUs, a possible direction is to develop a co-execution approach that makes better use of current pre-exascale supercomputers, which typically blend GPUs and CPUs. In this way, we make total usage of both GPUs and CPUs. CPUs are usually underused in such machines. A fast and scalable geometric mesh partitioning based on Space-Filling Curve (SFC) has been vital to enable the co-execution with a correct load balance between the GPUs and CPUs. At the beginning of the simulation, the SFC partitioning is called several times iteratively until an optimum partitioning of the mesh is obtained. In the first iteration, each MPI task (be it CPU or GPU) receives a specific portion of the mesh

according to some initial weights. With this partition, it calculates a couple of time steps. Based on the computational time taken by each MPI task, it adapts the weights and repartitions again. After a couple of iterations, each processor receives the correct amount of work so that they all take nearly the same time. GPUs receive a more significant chunk of the mesh than CPUs. In this way, the work done by the CPUs is spared in comparison to a pure GPU calculation. Tests for wind energy problems performed at the CTE-Power9 supercomputer at BSC, which has nodes with 2 POWER9 8335 CPUs and 4 NVIDIA Volta V100 GPUs, showed that the code scales correctly up to two million unknowns per node. Moreover, the GPU implementation runs 4.3 times faster than the CPU implementation, while the co-execution approach runs five times faster than the CPU implementation. Thus, the co-execution saves 18% of the time.

As already mentioned, the second key kernel is the solution of a linear system for pressure. Alya's most cited paper [18] recognizes that the lack of algorithmic scalability is one of its main limitations. This is due to the use of in-house coded Krylov-based solvers, whose number of iterations to converge to a certain tolerance is known to increase as the mesh is refined. Multigrid is the best solution for this problem. Since Alya deals with unstructured meshes, algebraic multigrid AMG is the obvious answer. Instead of trying to code our own AMG preconditioner, we have preferred to interface Alya to an external library. This work will present results obtained with PSBLAS and its multigrid package AMG4PSBLAS [19]. To perform the weak scalability study, we use an automatic mesh multiplication tool [20] available in Alya that allows partitioning elements into elements with half the size. We start with an unstructured mesh for the Bolund benchmark with 5.6 M nodes and apply the mesh multiplication twice to obtain a mesh with 358M nodes. We run the original mesh on one Marenostrum IV node with 48 cores. To maintain the average load per core when we go to the fine mesh, we use 64 nodes. The solver takes five iterations to converge on the coarse mesh, while it only takes four iterations on the fine mesh. This excellent algorithmic scalability is accompanied by CPU times increasing from 0.67s to only 0.72s when going from the coarse to the fine grid. Thus, we can say that this is the first time that Alya's solver attains correct weak scalability.

Although we only describe our experience with AMG4PSBLAS, we have also interfaced Alya with four other linear solvers AGMG [21], Maphys [22], Pastix [23], and Mumps [24]. The effort of interfacing with external linear algebra packages has been relatively small compared to what it would have cost to implement the same functionalities ourselves. This has been a paradigm change for Alya which relied mainly on in-house code in the past. The opening of Alya to the usage of external libraries has not concentrated only on linear algebra tools. Fault tolerance, input-output, and in-situ visualization libraries, such as FTI [25], PDI [26], Sensei [27],

Adios [28], and Catalyst [29], have also been interfaced with Alya lately. The strong interaction with experts from different fields is an efficient way of accelerating the path to exascale.

In summary, the use of exascale HPC platforms may improve dramatically the physical fidelity of the models used in wind farm design. A possible application is to consider wind power forecasts at different time scales, from short-term (meso-micro dynamical), to seasonal or decadal using global climate models. Small scale demonstrators (using WRF+LES dynamically coupled) indicate that errors in wind speed prediction could decrease from 30% to less than 4% by adopting these techniques in complex terrain and challenging weather conditions. However, the computational cost is still big, about 2 days of calculation using 1.5 Million CPUs together to produce an annual prediction of the energy of a wind farm. To be able to run such calculations effectively in an upcoming exascale platform, the wind energy meso and micro-scale codes will need to improve in several aspects. Improvements in both node-level and system-level performance need to be also accompanied by improvements in fault-tolerance, input-output, in-situ visualization, and even artificial intelligence to make the best usage of exascale hardware.

Finally, although Europe is currently a world leader in Wind energy with companies such as Vestas, Siemens-Gamesa and Iberdrola. However, its investments in HPC simulation for wind energy towards exascale is significantly smaller than what the US is investing. In the digital era it is well known that to out compute is to out compete. The US Wind vision recognizes that 'Without actions to improve wind's competitive position in the market, the nation risks losing its existing wind manufacturing infrastructure and a range of public benefits

as illustrated in the Wind Vision'. Advanced simulation has a key role in the US Wind Vision (Wind Vision: A New Era for Wind Power in the United States) instead in the EU Wind Vision (Wind Energy: A Vision for Europe in 2030) there is unfortunately no reference to simulation.

Herbert Owen, leader of the Wind Scientific challenge, within the EoCoE-II project has recently been in contact with Michael Sprague leader of the Exawind and Atmosphere to Electrons(a2e) US projects. While EoCoE and ENERXICO are the closest European counterpart to the a2e and Exawind US projects their Wind budget is one order of magnitude smaller. Moreover, the EU projects have a significantly shorter duration. With such huge differences in budget, it is inevitable to advance at a much lower pace in Europe than in the US. If EU really wants to be 'Fit for the 55' by 2030 and remain competitive with the US, this must change urgently. As President von der Leyen said at the European Parliament Plenary [30] 'At the heart of it is our mission to become the first climate-neutral continent by 2050. But we will not get there with the status quo – we need to go faster and do things better.' 'We are reaching the limits of the things we can do in an analogue way. Europe must now lead the way on digital.' 'If Europe is to move forward and move fast, we must let go of our hesitancies'. HPC and Wind Energy are two pillars of the transformation we want for Europe. HPC for Wind Energy needs to move fast in the correct direction since we are currently far behind the US.
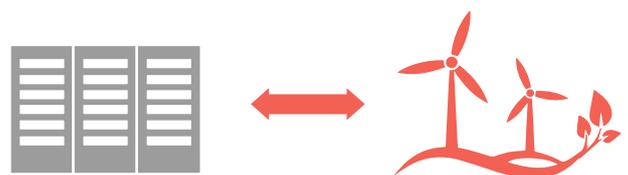
HPC IN THE TRANSPORTATION INDUSTRY

The transportation sector represents about 25% of the share in the energy sector and has a growth rate of 1.4%/year [31]. In the case of road transport, increased consumption of low-carbon fuels and deployment of hybrid or electric vehicles present viable alternatives, although the majority of vehicles operating today still use conventional fuels associated with internal combustion engines (ICEs). This scenario leads to issues with the energy supply for the transportation sector such as price rise, security and availability, in addition to the worldwide critical concern about the impact of future transportation related greenhouse gas emissions and pollutants on public health and climate change.

Current efforts in transportation aim at replacing internal combustion engines with electric motors in low power applications, such as motorcycles or passenger cars. For larger power applications aimed at the mobility of people, freight and air transport, doing without liquid fuelled powertrains is more difficult. One possible path is the replacement of current fossil fuels with renewable ones, which may help face the challenge of keeping a high energy density with the compromise of CO_2 neutrality. Decarbonization includes all options to reduce greenhouse gas (GHG) emissions and make road transport cleaner, including low-carbon energy carriers such as biofuels and power-to-liquid fuels, also known as electrofuels or e-fuels.

None of these solutions will be able to solve this challenge alone, and renewable transport fuels have an essential role in bridging the gap between GHG emission reduction targets and the prospected emission reductions [32].

According to the International Energy Agency (IEA) predictions, biofuel use for transportation should expand 24% over the 2019-2024 period [33]. Aside from biofuels, i.e. those derived from bioenergy, power-to-liquid fuels make up another interesting alternative. All these fuels are essentially starting from renewable hydrogen. Different options for their use in combustion engines are currently under intensive evaluation, including direct use of hydrogen, synthesis

processes combining hydrogen and CO2 and leading to a pledge of interesting oxygenated molecules (methanol, dimethyl-ether and oxy-methylene ethers), or even ammonia synthesis [34]. All these options should only be produced from renewable electricity, and their market penetration will depend largely on cost. Combustion characteristics with all such fuels need to be understood so that they can be applied in combustion systems with optimal efficiency and minimum pollutant emissions.

There are several constraints that limit the widespread use of renewable liquid fuels. Apart from the complexities in their production, meeting the stringent certification requirements and compatibility with the existing energy infrastructure are major challenges. Emissions are largely dependent on the fuel type and combustion system designs. Liquid fuel specifications usually are defined in terms of the required performance of the fuel, including characteristics such as the energy content, freeze point, and thermal stability. For automotive applications, it is often preferred to have liquid fuels to be "drop-in", which is defined as fuels that do not require any modification to the equipment or infrastructure. One of the major challenges in sustainable fuel supply is to balance the conventional fuel supply and economy for the next few decades, while satisfying ever tightening emission and safety requirements. In order to meet this challenge, new fuels need to be supplied and utilised efficiently with minimum environmental impacts. As these new fuel streams enter the market, a series of new engine technologies are under development, promising improved efficiency and cleaner combustion. To date, however, a fundamental understanding and strategic effort to match future liquid fuels with evolving injection systems and in-cylinder design is lacking. To provide the fundamental knowledge to enable the utilisation of these "e-fuels" in industry, the current knowledge on their combustion performance must be further explored. There has been a considerable number of activities in this area in the last few years leading to much faster progress than previously expected, with some fuel blends included in European/US specifications recently and several more expected to be included over the next few years. However, the existing specifications for the new generation of fuels are not fully established especially when emissions are concerned.

Combustion strategies involving these new fuel blends are likely to be characterised by higher in-cylinder pressure, higher level of dilution, minimum pollutant formation and CO2 emissions, and hence maximum efficiency. Combustion in unexplored thermodynamic environments where new physical and chemical fuel properties result in complex interactions with respect to fuel variability is not well understood. For instance, the fluid may well be in the supercritical regime for combustion at high pressures, where the fluid transport properties will be vastly different from those at ambient conditions. Both inter-molecular and intra-molecular energy transfer will impact the local reactive environment. Key aspects of the combustion kinetics of these fuel blends are also unknown.

Within this context, high-fidelity modelling capabilities are needed to guide the development of next-generation combustion technologies as well as the design and optimisation of advanced cleaner and highly-efficient blended fuel engine operation. This is especially important with future combustion engine applications being focused on heavy-duty road and marine transport, where large engine sizes make experimental work much more challenging and expensive. A practical framework connecting the sub-process models and providing a full description from the macroscopic scales through to the very detailed microscopic kinetic and turbulent mixing scales is essential. The access to Exas-

cale Computing provides a unique opportunity to accelerate this transition, as virtual testing platforms developed in the context of HPC can be used to carry out high-fidelity simulation of realistic engine conditions.

The use of advanced numerical simulations has enabled us to make important contributions for increasing cycle efficiency, reduction of pollutant emissions, and use of alternative fuels in practical applications [35]. Numerical tools are employed routinely every day to design and optimize combustion systems by aircraft, automotive or gas turbine manufacturers. Indeed, the increase in computing power over the last years has led to transition from Reynolds-averaged Navier-Stokes (RANS) to large-eddy simulations (LES) approaches in the design and development process, reducing the uncertainty of the models and increasing the reliability of the numerical predictions [36]. In this new era, when future exascale architectures become available, the use of LES or even DNS could transition to make a more important role on better understanding the performance of combustion systems in more complex conditions. The full characterization of practical combustion systems is a complex task and brings many challenges associated with different disciplines of engineering, chemistry, physics, computer science and mathematics, among others.

In particular, one of the most important parts is the modelling of the reacting spray under realistic conditions, as it includes the interaction of complex physical phenomena such as high-pressure liquid fuel injection, atomization, vaporization, fuel/air mixing and combustion. This interaction is not well understood and plays a major role in the overall performance of the system. With the use of future exascale architectures, the CFD codes will have to adapt to run efficiently on these machines with new algorithms and numerical methods, but they will

have the full potential of performing exaFLOPS operations, this means 1018, floating point operations per second.

One of the fundamental issues associated with the modelling and simulation of renewable

fuels under engine-like conditions is to account for the chemistry. In practice, traditional hydrocarbon fuels require the order of hundreds of species and more than thousands of reactions to characterize the combustion process at engine-like conditions [37]. In this type of problems, the cost of solving Navier-Stokes equations becomes negligible compared to the one needed for the chemistry calculations. These chemical kinetics schemes become even larger with realistic fuel descriptions and complex surrogates or blending of fuels, so this holds the investigation of alternative fuels and reduces the applicability of numerical simulations in this context [38, 39].

The correct prediction of the formation and destruction of species and radicals not only requires small time steps, but also high resolution in the computational domain, as some species can have short and large life times. In practice, the emissions are controlled by slow and fast reactions, i.e. NOx is governed by slow time scales, but CO coming from the partial

oxidation of CO2 is characterized by fast time scales, so both slow and fast processes are important and need to be taken into account. This means the numerical simulations will have to be run for long physical times using fine grids and small time steps. Such high requirements in terms of spatial and temporal resolution are even more demanding in the case of systems operated at moderate or high pressure. Under these conditions, the reacting layers are smaller and the computational cost of the simulations increases dramatically. There are strategies to simplify the chemistry problem, like chemical reduction, chemistry tabulation, transported PDF models or conditional closures, but those are only valid for well-known regimes and conditions, and their applicability to general problems can be questionable [40, 41, 42, 43, 44]. Another important limitation of the current modelling technologies is associated with the atomization process of the fuel [45, 52]. The formation of the spray is a complex physical phenomenon characterized by the break-up process of the liquid film into droplets and its subsequent atomization stages until the droplets are vaporized achieving the gas phase. The fuel injection process has a strong influence on stability limits, combustion efficiency and pollutant emissions. Nowadays, high-pressure injection system designers make great efforts to ensure the generation of small droplets to promote fast evaporation and more homogeneous mixtures, which in turn increase com-

bustion efficiency and reduce pollutant emissions. The atomization process can be divided into primary and secondary breakup with different physical dynamics and are controlled by droplet collisions, coalescence, heat transfer, phase change along with turbulent interactions between phases [46, 47, 48]. This is especially complex at the conditions of the operation of the engines, which usually have high Reynolds and Weber numbers conditions with liquid fuel injection at high speed on small nozzle diameters.

There are many approaches to solve these problems, going from Lagrangian droplet methods to fully Eulerian approaches [49, 50]. The simplest approach to model primary atomization is the use of phenomenological models, which requires low computational cost, but offers low accuracy. These models are able to provide distribution of droplets downstream the atomizing edge with a small number of parameters and without simulating the actual atomization process. These are the most appropriate for combustion analysis where the main purpose is not to understand the atomization process, but to accurately reproduce the main characteristics of the spray that have an impact on the flame and subsequent pollutant emissions. The low CPU cost and memory requirements of this approach make it suitable for its integration in CFD codes in order to perform large-eddy simulations of the complete domain (atomizer and combustion chamber). This approach has been

recently used to develop a phenomenological model for the atomization process downstream an atomizing edge in a LES frame [50].

A second approach, based on the surface/density ELSA model [48], allows the study of more realistic configurations, but at higher CPU cost. In this approach, the liquid phase is described by its mass fraction and the mean interface surface area per volume. This information is coupled to a Lagrangian solver in the dilute region to follow the spray evolution. This method has given interesting results for airblast atomizers and high-pressure liquid jets for ICEs [51], but it also needs a preliminary geometry- dependent calibration. Finally, among the approaches with the highest degree of detail and cost are interface capturing techniques such as Level Set [53, 54, 55, 56].

These methods require Direct Numerical Simulation (DNS) to accurately predict the turbulence level at the origin of the primary instabilities in the liquid. They have demonstrated their capability to accurately describe the atomization process, but the large range of involved length scales requires grids with an extreme number of cells, and are restricted to reduced size and simplified academic configurations. These are particularly demanding problems that are not feasible with today's modelling strategies and current computational resources at conditions of practical interest.

In summary, the exascale era will provide a significant platform for making important contributions in the power and transportation sectors towards more efficient, more fuel-flexible and low emissions systems with direct impact on public health and climate change. With the use of these new architectures and with the corresponding advances on the codes to fully exploit this new chip capabilities, the challenges on propulsion technologies and power generation systems will be conveniently addressed allowing a transition to a greener and more advanced combustion systems based on low carbon fuels combined with renewable energy technologies that can operate jointly.
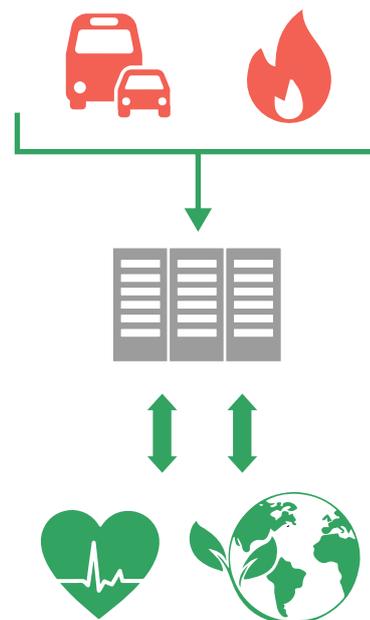
REFERENCES

[1] United Nations Framework Convention on Climate Change, 2015. Paris Agreement: UN Climate Conference.

[2] Balance Nacional de Energía 2019. Subsecretaría de Planeación y Transición Energética. Dirección General de Planeación e Información Energéticas. México, 2020.

[3] Global Energy Review 2020, ISBN 92-64-36216-9, IEA, 2020.

[4] https://www.spglobal.com/marketintelligence/en/news-insights/latest-news-headlines/fossil-fuel-financing-falls-in-2020-amid-covid-19-environmental-groups-find-63302209

[5] https://www.power-technology.com/news/covid-19-pandemic-accelerated-future-renewable-energy-usage-poll/

[6] Navon, A.; Machlev, R.; Carmon, D.; Onile, A.E.; Belikov, J.; Levron, Y. Effects of the COVID-19 Pandemic on Energy Systems and Electric Power Grids—A Review of the Challenges Ahead. Energies 2021, 14, 1056. https://doi.org/10.3390/en14041056

[7] Gray, S. H., Etgen, J., Dellinger, J., & Whitmore, D. (2001). Seismic migration problems and solutions. Geophysics, 66(5), 1622-1640.

[8] www.top500.org

[9] A. Long (2020). HPC and the Future of Seismic. Industry Insights (www.pgs.com)

[10] Virieux, J., & Operto, S. (2009). An overview of full-waveform inversion in exploration geophysics. Geophysics, 74(6), WCC1-WCC26.

[11] Avila M, Gargallo-Peiro A, and Folch A. A CFD framework for offshore and onshore wind farm simulation. Journal of Physics: Conference Series, 854(1):012002, 2017.

[12] Javier Sanz Rodrigo J, Churchfield M and Kosovic B. A methodology for the design and testing of atmospheric boundary layer models for wind energy applications. Wind Energ. Sci., 2, 1–20, 2017.

[13] Lehner, M, 2012, Observations and large-eddy simulations of the thermally driven cross-basin circulation in a small, closed basin, Ph.D. thesis. University of Utah.

[14] Lehmkuhl O, Houzeaux G, Owen H, et al. A low-dissipation finite element scheme for scale resolving simulations of turbulent flows. J. Comput. Phys., 390:51 − 65, 2019. ISSN 0021-9991.

[15] Vreman AW. An eddy-viscosity subgrid-scale model for turbulent shear flow: Algebraic theory and applications. Phys. Fluids, 16(10):3670−3681, 2004.

[16] J Sanz Rodrigo et al.The ALEX17 diurnal cycles in complex terrain benchmark. 2021 J. Phys.: Conf. Ser. 1934 012002.

[17] https://github.com/RRZE-HPC/likwid

[18] Mariano Vázquez, Guillaume Houzeaux, Seid Koric, Antoni Artigues, Jazmin Aguado-Sierra, Ruth Arís, Daniel Mira, Hadrien Calmet, Fernando Cucchietti, Herbert Owen, Ahmed Taha, Evan Dering Burness, José María Cela, Mateo Valero, Alya: Multiphysics engineering simulation toward exascale, Journal of Computational Science, Volume 14, 2016, Pages 15-27, https://doi.org/10.1016/j.jocs.2015.12.007.

[19] Pasqua D'Ambra, Fabio Durastante, and Salvatore Filippone. Amg preconditioners for linear solvers towards extreme scale, 2021.

[20] Guillaume Houzeaux, Raul de la Cruz, Herbert Owen, and Mariano Vzquez. Parallel uniform mesh multiplication applied to a Navier Stokes solver. Computers and Fluids, 80:142 − 151, 2013.

[21] Y. Notay. AGMG software and documentation. http://agmg.eu

[22] MaPHyS team. MaPHyS software and documentation. https://gitlab.inria.fr/ solverstack/maphys.

[23] PaStiX team. PaStiX software and documentation. https://gitlab.inria.fr/ solverstack/pastix.

[24] MUMPS team. MUMPS: Multifrontal massively parallel solver. http://mumps-solver. org/.

[25] https://github.com/leobago/fti

[26] https://pdi.julien-bigot.fr/master/

[27] https://sensei-insitu.org

[28] https://csmd.ornl.gov/adios

[29] https://www.paraview.org/Wiki/ParaView/Catalyst/Overview

[30] https://ec.europa.eu/commission/presscorner/detail/en/SPEECH_20_1655)-

[31] International Energy Outlook 2016 (IEO2016), With Projections to 2040, Energy Dept., Energy Information Administration, Office of Energy Analysis.

[32] D. Bacovsky, A. Sonnleitner, A. Brown, The Role of Renewable Transport Fuels in Decarbonizing Road Transport, AMF Annex 58 / IEA Bioenergy Task 41 Project 10, 2020.

[33] IEA (2019), "Renewables 2019", IEA, Paris
https://www.iea.org/reports/renewables-2019

[34] Masri, A. R. (2021). Challenges for turbulent combustion. Proceedings of the Combustion Institute, 38(1), 121–155.

[35] H.Pitsch, Large-eddy simulation of turbulent combustion, Annual Review of Fluid Mechanics 2006 38:1, 453-482.

[36] Kim WW, S Syed. 2004. "Large-Eddy Simulation Needs for Gas Turbine Combustor Design." In Proceedings of the 42nd AIAA Aerospace Sciences Meeting and Exhibit, pp. 3162-3175, Reston, Virginia.

[37] Dryer, F.L., Chemical kinetic and combustion characteristics of transportation fuels, Proc. Combust. Inst., 35 117-144, 2014.

[38] Jacobs, S., D¨ontgen, M., Alquaity, A. B., Kopp, W. A., Kr¨oger, L. C., Burke, U., Pitsch, H., Leonhard, K., Curran, H. J., and Heufer, K. A., 2019. "Detailed kinetic modelling of dimethoxymethane. Part II: Experimental and theoretical study of the kinetics and reaction mechanism". Combustion and Flame, 205, jul, pp. 522–533.

[39] Cai, L., Jacobs, S., Langer, R., vom Lehn, F., Heufer, K. A., and Pitsch, H., 2020. "Auto-ignition of oxymethylene ethers (OMEn, n= 2−4) as promising synthetic e-fuels from renewable electricity: shock tube experiments and automatic mechanism generation". Fuel, 264, p. 116711.

[40] Kundu, P., Ameen, M., Unnikrishnan, U., and Som, S., "Implementation of a Tabulated Flamelet Model for Compression Ignition Engine Applications," SAE Technical Paper 2017-01-0564, 2017

[41] Jochim, B., Korkmaz, M., and Pitsch, H. (2017). Scalar dissipation rate based multi-zone model for early-injected and conventional diesel engine combustion, Combust. Flame, 175, 138–154.

[42] D. Mira, O. Lehmkuhl, A. Both, P. Stathopoulos, T. Tanneberger, T. Reichel, C.O. Paschereit, M. Vázquez, G. Hozeaux, 'Numerical characterization of a premixed hydrogen flame under conditions close to flashback', Flow, Turbulence and Combustion, 104(2): 479-507, 2020

[43] Wehrfritz, A., Kaario, O., Vuorinen, V., and Somers, B. (2016). Large Eddy Simulation of n-dodecane spray flames using Flamelet Generated Manifolds, Combust. Flame, 167, 113–131.

[44] D. Mira, E. Pérez, A. Surapaneni, J.M. García-Oliver, J. Benajes, J.M. Pastor, D. de León, LES Study on Spray Combustion With Renewable Fuels Under ECN Spray-A Conditions, Internal Combustion Engine Division Fall Technical Conference ICEF2021, 67745.

[45] Faeth G.M., Hsiang L.-P., Wu P.-K., "Structure and breakup properties of sprays", Int. J. of Multiphase Flow 21 (1995) 99–127.

[46] Hiroyasu H., Spray breakup mechanism from the hole-type nozzle and its applications, Atomization and Sprays", vol. 10, pp. 511–527, 2000.

[47] Pitsch, H. and Desjardins, O., Detailed numerical investigation of turbulent atomization of liquid jets, Atomization Sprays, Vol. 20, pp. 311–336, 2010.

[48] Vallet A., Burluka A.A., Borghi R., Development of a Eulerian model for the atomization of a liquid jet, Atomization and Sprays 11, 619–642, 2001.

[49] F. Jaegle, J.-M. Senoner, M. Garcıa, F. Bismes, R. Lecourt, B. Cuenot, T. Poinsot, Eulerian and Lagrangian spray simulations of an aeronautical multipoint injector, Proc. Combust. Institute 33 (2011) 2099–2107

[50] Chaussonnet, G., Gepperth, S., Holz, S., Koch, R., Bauer, H.J., "Influence of the ambient pressure on the liquid accumulation and on the primary spray in prefilming airblast atomization", Int J Multiphase Flow 125, p. 103229, 2019.

[51] S. Jay, F. Lacas, S. Candel, Combined surface density concepts for dense spray combustion, Combustion and Flame, Volume 144, Issue 3, 2006, pp. 558-577.

[52] Menard, T., Tanguy, S., Berlemont, A., 2007. Coupling level set/VOF/ghost fluid methods: validation and application to 3D simulation of the primary break-up of a liquid jet. Int. J. Multiphase Flow 33, 510–524.

[53] O Desjardins, V Moureau, H Pitsch, An accurate conservative level set/ghost fluid method for simulating turbulent atomization, Journal of computational physics, 2008, Vol. 227, pp. 8395-8416.

[54] Fuster, D., Bagu , A., Boeck, T., Le Moyne, L., Leboissetier, A., Popinet, S., Ray, P., Scardovelli, R., Zaleski, S., 2009. Simulation of primary atomization with an octree adaptive mesh refinement and VOF method. Int. J. Multiphase Flow 35, pp. 550–565.

[55] Shinjo, J., Umemura, A., 2010. Simulation of liquid jet primary breakup: dynamics of ligament and droplet formation. Int. J. Multiphase Flow 36, 513–532.

[56] Herrmann, M., 2010. A parallel Eulerian interface tracking/Lagrangian point particle multi-scale coupling procedure. J. Comput. Phys. 229, 745–759.
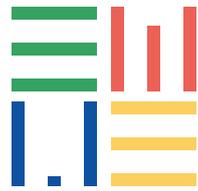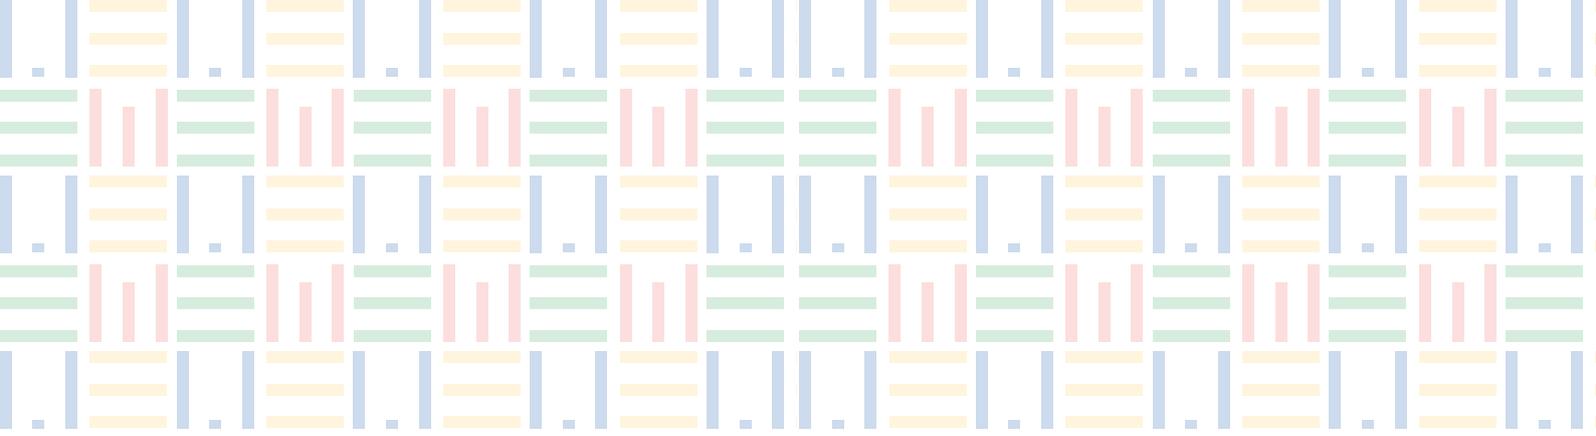
ENERXICO

enerxicoproject@bsc.es