

Parallel computation of aeroacoustics of industrially relevant complex-geometry aeroengine jets

Zhong-Nan Wang^{a,*}, James Tyacke^a, Paul Tucker^a, Peer Boehning^b

^a*Department of Engineering, University of Cambridge, Cambridge CB2 1PZ, UK*

^b*Aeroacoustics, Rolls-Royce Deutschland, Blankenfelde-Mahlow 15827, Germany*

Abstract

Jet noise is still a distinct noise component when a commercial aircraft is taking off. A parallel high-fidelity simulation framework for industrial jet noise prediction is presented in this paper. This framework includes complex geometry meshing and Ffowcs Williams-Hawkings (FW-H) surface placement during preprocessing, a parallel hybrid RANS-LES flow solver coupled with an FW-H acoustic solver in the simulation and mean and unsteady data processing after the simulation. The use of this framework is demonstrated through two jet noise prediction cases: in-flight heated jets and installed ultra-high bypass-ratio (UHBPR) engines. These simulations can provide more insight than experimental tests into jet flow physics for engineering model improvement. Additional advantages are also shown in the cost and turn-around time. Thus there is great potential for high-fidelity jet noise simulations to partly replace rig tests for industrial use in the future.

Keywords: parallel computation, hybrid LES-RANS, jet noise prediction

1. Introduction

The world aircraft fleet is expected to double by 2036 relative to 2016 with increased demand for air transport [1]; reducing the environmental impact of aircraft is thus becoming more urgent. Aviation noise, as one such impact, affects the airport neighboring residents and is an important aircraft certification index. Jet noise is one of the most significant noise components when a civil aircraft is taking off. The noise certification process requires propulsive jets to be evaluated with forward flight streams at different speeds. Experimentally testing this is extremely difficult, because acoustic measurements are easily contaminated by the background noise of the flight stream wind tunnels [2]. In addition, some operating conditions are not always achievable even at a large expense. High-fidelity numerical approaches are therefore pursued as an alternative to generate reliable data for product design validation [3].

*Corresponding author

Email address: znw22@cam.ac.uk (Zhong-Nan Wang)

In recent years, successful attempts have been made to numerically predict jet noise from first principles using Large-eddy Simulation (LES). LES has been demonstrated to capture the large coherent structures responsible for major components of sound generation and its use in jet noise prediction is well reviewed [4, 5, 6]. Compared with experiments, this type of simulation can provide more insight into sound source physics and help with physics-based design improvements and low-order acoustic model development [7]. Chapman, the former director of aeronautics in NASA, envisioned this application of LES in the engineering design environment in 1979 [8]. His vision was based on wall-resolved LES grid requirements and the assumption of exponential growth in computational power and significant improvement in numerical algorithms. Although his vision is optimistic overall, the LES of jet noise of a real engine exhaust is close to reality because it largely comprises of free shear flow and is not strictly bounded by wall grid resolution except on nozzle surfaces. However, for industrial use, the simulation needs to be performed on real-world complex geometries and completed in a relatively short turn-around time. Therefore, the capability of handling complex geometries and efficient parallel computation is clearly necessary to make large-scale high-fidelity simulations meet the industrial design cycle.

Computational power has been increasing fast following Moore’s law since 1970s. The recent advances in high performance computing (HPC) technology are about to bring us into a new era of computer performance - exaFLOP/s (10^{18} floating-point operations per second) by around 2020, see Fig. 1. As CPU frequency stalled near 2005, multi-core architectures, such as GPUs and coprocessors, have emerged to continue the trend via massive parallelism and also reduce the energy consumption per FLOP. This indicates a decrease in computational cost. Although Moore’s “law” is challenged to hold further ahead, new computing concepts, such as quantum computing or biological computing, still offer a promising future for a step-change computing power increase. These all enable computational fluid dynamics (CFD) to tackle large-scale and real-world problems in the foreseeable future.

Figure 1 shows our experience of jet LES with computing power growth along a timeline. As computing power has increased exponentially over years, the capability of LES prediction of jet noise has also developed rapidly. This can be represented by the increase of Reynolds number, mesh size and geometrical complexity of simulated jets. It starts with simple pipe nozzle geometry [9, 10] and then moves towards explicit real geometries (single round/serrated nozzles [11, 12, 13, 14, 15], dual stream round and serrated nozzles [16, 15, 17], nozzle with pylon structures [18] and jet-wing configurations [19, 20, 21]). The simulation technology itself is also becoming mature through this process. The computation techniques develop from jet noise alone simulation to component coupling simulation, i.e. installation noise from jet and wing/flap interactions even including upstream fan wakes [19]. The technology is now passing the validation phase and entering the prediction phase [6]. In this stage, new product validation and novel noise reduction concepts can be pursued.

In this paper, the main process chain of this jet noise prediction is shown, in-

cluding pre-processing, flow and acoustics simulation and data post-processing. Parallel computing techniques with modern architectures are discussed to speed up the process to satisfy the short industrial design schedule. After that, examples of this type of simulation are demonstrated to show how all these techniques work together to solve real-world problems. Finally, the current technology level is summarized and future challenges with possible technology solution is envisioned.




Figure 1: Computing power growth and jet LES development

2. Simulation process for jet noise

The NASA CFD vision 2030 study [3] lists 4 grand challenge (GC) problems. Our research on the high-fidelity simulation of jet noise lies within the first two GC problems: simulation of powered aircraft over the full envelope and the turbofan transient state at off-design points. In order to solve them, there are clearly a number of technical challenges, such as complex geometry handling, solver parallel scaling and multidisciplinary problem solving. This section introduces the process chain of jet noise simulation used in our research and demonstrates how the processes address the technical challenges. Future technology development is also envisioned.

Figure 2 shows the outline of this process. It demonstrates the key elements in three stages (pre-processing, simulation and post-processing). The arrows represent the data/information flow. Iteration happens among the stages to refine the simulation. For example, the preliminary simulation is first run on a

coarse mesh to obtain an initial estimate and then a final simulation is launched on the refined mesh to achieve a better result. When the elements in this process are built, they can be automated, and then used for the optimization in the future.

85 To meet the tight industrial design cycles, the simulation needs to be completed in a few weeks if used for design validation. Parallelization must be used to speed up the whole process. Parallelization is possible in two manners: First, each element in the process runs efficiently in parallel; Second, the data flow in the process chain is parallelized. In other words, each element does not only
 90 run in parallel, but potentially can also be connected in parallel. In this section, only element-level parallelization is discussed.




Figure 2: Simulation process for jet noise

2.1. Pre-processing

The objective of pre-processing is to provide what is needed to launch a simulation. Meshing is obviously the most important part of pre-processing.
 95 In addition, data collection on the near-field surface is also needed for far-field sound prediction in the Ffowcs William-Hawkings (FW-H) acoustic solver.

2.1.1. Mesh Generation

The jet nozzle geometry in real aeroengines is complex. The capability to mesh complex-geometry jet nozzles is necessary. When the engine is installed
 100 much closer to the wing for ultra-high bypass-ratio (UHBPR) engines, the coupling between the propulsive jet and the aircraft wing becomes more obvious causing additional noise. Therefore jet noise simulations should also be performed with wing configurations. This poses extra challenges for meshing. An example of this type of mesh is shown in Fig. 3.

105 Hybrid structured-unstructured mesh offers a good option for complex geometry meshing. It offers flexibility of meshing in different flow zones. Generally, hexahedral mesh is preferable for LES zones because it gives low numerical dissipation and smaller total cell count compared to tetrahedral mesh. Tetrahedral mesh fits better for acoustic wave propagation in the near-field acoustic zone.
 110 For jet shear layers, structured-unstructured mesh modules can be built to align grid lines with shear layers and also refine azimuthal resolution near the nozzle if needed, e.g. for serrated nozzle simulations, see Fig. 3. The modular concept




Figure 3: Hybrid structured-unstructured mesh for installed chevron jets

is to use hexahedral mesh for the majority of the jet plume and hexa-dominant
 115 prisms for shear layers in the near nozzle region. The hexa-dominant prism
 layers are generated on the wing surfaces. Tetrahedral cells are then used to
 connect these regions. This modular structured-unstructured mesh can allow
 the usage of the right type of mesh in the most suitable region and therefore
 the best mesh quality can be achieved. The mesh quality is carefully controlled
 in the LES region. A geometric expansion ratio is kept approximately 1% in
 120 the axial and radial directions with uniform radial grid distribution near the
 potential core. The mesh resolution follows the summarized simulation practice
 [6].

This modular structured-unstructured mesh can also be recorded as a tem-
 125 plate for a given topology. When the nozzle shape changes during the design
 phase, the mesh can be generated without human intervention. This supports
 the automation of simulation process.

2.1.2. FW-H surfaces placement

Flow data is recorded on near-field surfaces to calculate far-field acoustic
 pressure fluctuations using the FW-H method. The method will be described in
 130 detail in the section 2.2.2. A key aspect of using this method is the selection of
 FW-H surface locations. The surfaces have to enclose all sound sources and are
 outside the rotational vortical regions to avoid hydraulic signal contamination.
 Normally a series of surfaces at different radial positions (shown in Fig. 4) are
 used to find the convergent prediction of far-field sound. It is also challenging
 135 to close the FW-H downstream end surface in some cases, because vortical
 structures can pass the end surface and cause spurious noise. In order to remove
 this, a set of downstream end surfaces can be placed with incremental axial
 displacement and the filtering techniques have been developed to potentially
 remove hydraulic signals [22]. As presented here, the surface can also be left
 140 open leading to only a small amount of sound missing at low polar angles, but
 this can potentially be compensated by making the surface long downstream.

For isolated jet cases, it is easy to use axisymmetric surfaces for FW-H predictions, see Figure 4(a). For industrial complex-geometry jets, such as installed jets, it is more difficult to place FW-H surfaces. Level-set methods are used for setting the FW-H surfaces based on the walls (nozzle and wing), vorticity and turbulence intensity isosurfaces. Generally, the turbulence intensity is less than 0.25% and the normalized vorticity $|\Omega|D/U_c$ is less than 0.05 on the FW-H surface. An example of FW-H surface placement for this installed jets is shown in Fig. 4(b). Unstructured elements can be generated on these complex FW-H surfaces and flow variables are recorded on the element centroids to facilitate the integration of far-field sound prediction.




Figure 4: FW-H surface placement for jet noise prediction

2.2. Solvers

The solvers form the core part of the simulation process, so they must be parallelized most efficiently. As a result of the multidisciplinary nature of aeroacoustics, both a flow solver and an acoustic solver are employed with a one-way coupling.

2.2.1. Flow Solver

The flow solver is used to calculate the near-field jet flow, which constitutes the sound source. LES is used to model the jet plume turbulence with RANS for attached boundary layers adjacent to solid surfaces.

Numerics. The CFD solver used here is an unstructured, edge-based finite volume code for compressible flows [23]. For LES, low numerical dissipation is required to preserve turbulent eddies. A kinetic energy preserving scheme is used for convective fluxes

$$F^{conv} = F^{KEP} - \frac{1}{2}\varepsilon|A|[L(U_R) - L(U_L)] \quad (1)$$

where, F^{conv} represents convective flux, U is the conservative variable vector, $|A| = (\frac{\partial F}{\partial U})$ is the absolute Jacobian matrix, L is the pseudo Laplacian operator. The parameter ε controls the artificial dissipation level.

This scheme preserves kinetic energy evolution avoiding artificially dissipating it. It is helpful in maintaining the turbulence cascade process. The numerical stability can also be enhanced by preserving this physical quantity. The smoothing factor ε is kept very small or zero in the LES region and increased towards boundaries to remove reflections. Small amount of numerical dissipation might be needed to stabilize the simulation when the grid is of insufficient resolution or poor quality, but the level should be kept small enough compared to the SGS dissipation to ensure the simulation quality in the region of interest.

The solver uses a second order backward time differencing with dual time stepping for unsteady simulations. A five-stage Runge-Kutta method is implemented for the pseudo time steps. This dual time stepping allows a larger stable CFL number to be used for large aspect ratio cells in the mesh of boundary layers and shear layers.

Physical Modeling. Modeling turbulence plays a key role in the simulation. LES is used to resolve the sound-generating large coherent structures in the jet plume. RANS is used near the walls to reduce the cost of resolving fine turbulent wall streaks. This thin RANS layer is blended with the LES zone using a modified wall distance [24]. The blending is achieved at the Reynolds stress τ_{ij} level:

$$\tau_{ij}^{turb} = f\tau_{ij}^{SGS} + (1 - f)\tau_{ij}^{RANS} \quad (2)$$

where f is the blending function based on the wall distance d :

$$f(d) = \min[\max(\frac{d - (1 - \beta)d_{RANS}}{\beta d_{RANS}}, 0), 1] \quad (3)$$

d_{RANS} represent the RANS layer thickness, and β defines the size of RANS-to-LES transition zone.

This sets the framework of blending RANS with LES. Freedom is left for the hybridization of any RANS and SGS models. The blending function $f(d)$ enables a quick transition from RANS to LES. Figure 5 shows the blending function near the jet nozzle.




Figure 5: Hybrid RANS-LES strategy illustration

Parallel Implementation. The parallelized flow solver developed in the Oplus frame [25]. The communication across the processors is implemented using MPI [26] in the Oplus. The unstructured mesh is partitioned using the ParMetis library [27] utilizing the k-way method with the node weight based on the number of edges attached to it. This is because the number of FLOPs is mainly proportional to the edge number for this edge-based solver. Each mesh partition is sent to the corresponding processor. The load balance is tested on 1056, 4224 and 8448 cores respectively, as shown in Fig. 6(a). The averaged edge number per core is depicted in a black line with error bars representing the standard variance around the mean value. The ideal partitioned mesh size on each core is shown in red as a reference. The real partition line is parallel to the ideal with only a small deviation (about 5-8% of the mesh count per core). This indicates that the mesh partition achieves a good load balance across cores. The halo nodes are generated for communication in each partitioned mesh to duplicate the grid nodes overlapping with the mesh on neighboring processors. Thus the average mesh size for the real partition is slightly larger than that of the ideal partition.




Figure 6: Flow solver parallelization

After mesh partitioning, the main computation is performed on the mesh owned by each individual core. The nodes of the partitioned mesh on each processor are reordered locally to make the data access as contiguous as possible to improve the cache performance. The parallelization in Oplus is fine grain where every loop is parallelized. The loop in this edge-based solver is generally categorized as: an edge loop or a node loop. Node loops are normally associated with local computation while the edge loop, e.g. flux calculation, requires halo node information that is actually stored in neighboring partitions. Communication is carefully managed to reduce latency. Computation is first performed on internal edges without halo node information. In the meantime, communication is performed to exchange halo node information among cores. After communication, computation is then performed over the outer edges. Latency is minimized by

overlapping the internal edge computation with the processor communication. The parallel scaling is shown in Fig. 6(b). The speed up in this plot is defined as

$$Speedup = \frac{T(256)}{T(N_{proc})} \quad (4)$$

225 where $T(N_{proc})$ is the wall-time spent on the given computational task using N_{proc} processors. The code shows a good speed up over 8000 processors compared to the ideal speed up line ($Speedup = N_{proc}/256$) with 256-core performance as reference. This scaling enables the large-scale simulation to be performed in a short turn-around time by making the most of the available
230 computational resources.

Jet noise simulation also needs frequent output of unsteady data, so parallel output efficiency is crucial. The conventional output is serial, gathering data from slave processors to the master processor which writes to the file system. This usually takes a large amount of time and is unsuitable for frequent unsteady
235 data output. Two approaches of parallel I/O are implemented and tested here. The first approach is called direct parallel I/O. Each processor writes out its own flow data on the partitioned mesh individually to the hard disk. These flow files are later merged into one entire flow file by a separate post-processing code. The other approach is implemented using HDF5 [28]. HDF5 is built on MPI and
240 enables the user to handle the data I/O at a higher level. It supports parallel access from all processes in a single MPI communicator to a single shared file by using MPI-I/O internally. The data is written in parallel using "collective" mode, which allows MPI to move data between processors for better file system access. The end result is a single file with no further processing required. A test
245 is done on a 150-million-cell hybrid mesh using the two parallel I/O approaches. The results in Fig. 7(a) show that the direct parallel I/O scales linearly with the number of grid nodes per processor while HDF I/O tends to saturate at a smaller number of nodes per processor. The direct parallel I/O is around four orders faster than the HDF I/O but with larger variation. This variation is
250 largely caused by the least balanced mesh partitions and the file system access latency. The output time spent on the data of one node is shown in Fig. 7(b). The normalization is used to show the relative time spent on writing data for one grid node. The direct parallel I/O shows a slow decreasing rate of relative I/O time per node with the number of nodes per processor but also a large
255 relative variation. The HDF I/O exhibits an optimal value of node number per processor to achieve best performance. This is because of the overhead related to data movement between the cores versus data writing.

2.2.2. Acoustic Solver

An acoustic solver is used in conjunction with the flow solver for far-field
260 sound prediction. It receives the unsteady data on the FW-H surfaces (illustrated in Fig. 4) from the flow solver and calculate acoustics pressure fluctuations at specified far-field locations.




Figure 7: Parallel I/O performance

Ffowcs Williams-Hawkings method. There are many existing methods to predict far-field sound based on near-field sound source information. The Ffowcs Williams-Hawkings (FW-H) integration [29] is a relatively robust and fast method compared to others, such as the Kirchhoff integration and linear Euler equations. The original formula contains a volume integration of quadrupole sources outside the closed surface. This term is computationally expensive and can be dropped by making the surface large enough to enclose the majority of quadrupole sources. When the jet is operating under the flight stream conditions, the convective FW-H formula [30] should be used, which considers the convective effects of mean flow. The flight stream velocity U_0 is included in the convective formula and the far-field acoustic pressure \tilde{p}' is calculated at physical reception coordinates \mathbf{x} and time t .

$$4\pi\tilde{p}'(\mathbf{x}, t) = \int_S [(1 - M_0\tilde{R}_1) \frac{\dot{Q}_i n_i}{R^*} - U_0 \frac{\tilde{R}_1^* Q_i n_i}{R^{*2}}] dS + \int_S [\frac{\dot{L}_{ij} n_j \tilde{R}_i}{c_0 R^*} + \frac{L_{ij} n_j \tilde{R}_i^*}{R^{*2}}] dS \quad (5)$$

where, \mathbf{x} , \mathbf{y} is observer location and source location respectively. The mass flux is $Q_i = \rho U_i - \rho_0 U_{0i}$, the momentum flux is $L_{ij} = \rho u'_i (u'_j + U_{0j}) + p' \delta_{ij}$. Subscript 0 represents flight stream properties. ρ_0 and c_0 are the flight stream density and sound velocity, M_0 is the flight mach number, and U_{0j} and U_{0j} are the flight stream and actual velocity in the x_j -direction. n_j is the FWH surface normal unit vector. The time derivative is expressed by $(*) = \partial(*)/\partial t$, while the spatial derivative is expressed by $(*)_i = \partial(*)/\partial x_i$. The acoustic distance R^* is defined as

$$R^* = \sqrt{(x_1 - y_1)^2 + \beta[(x_2 - y_2)^2 + (x_3 - y_3)^2]} \quad (6)$$

$$\beta = \sqrt{1 - M_0^2} \quad (7)$$

$$R = \frac{-M_0(x_1 - y_1) + R^*}{\beta^2} \quad (8)$$

This formula requires flow data recorded on the near-field surface when the flow solver is running. Surface integration is performed to calculate pressure fluctuations p' at the far field afterwards, however could be calculated simultaneously. On the closing end surface, filtering could be performed in the near-field and far-field to remove the pseudo sound caused by passing vortices [22], but the effects vary for different jets. However, if the surface is long enough, the end surface can also be left open with only a small amount of sound missing at low polar angles.

Parallel Implementation. The FW-H solver is usually run serially as a post-processing tool. It can be acceptable in some aeroacoustics prediction cases which contains a small number of surface elements, time samples and observer locations. However, when the FW-H solver is used for large-scale industrial use, it can be time consuming. For example, the sound prediction for the installed jets, illustrated in Fig 4(b), requires 5076 observer points (36 azimuthal angles and 141 polar angles per azimuthal angle). To compute the far-field sound at this number of observer locations can take several days if using a serial code. Therefore, the parallelization of the FW-H solver is necessary to accelerate the process.

With the advent of many-core architectures, the code can achieve a significant speed increase. The Intel Knights Landing (KNL) processor is one example of this category. It contains 64-72 cores per node and 512-bit vector units per core. This leads to a high potential for vectorization and parallelization in shared memory mode. The KNL also integrates a high bandwidth memory (MCDRAM) with the main memory (DRAM). This enables significantly improved performance for memory-bound codes. The parallelization of the FW-H solver is based on KNL nodes and the implementation of the convective FW-H equations 5-8 is outlined in Fig. 8. The code starts with the data reading. The outer most loop (L1) covers the observer locations. Since the sound prediction for each observer location is independent, it is ideal to use MPI to parallelize the observer loop. It does not need to pass any message during the computational process. Thus the scaling of MPI parallelism is almost near the ideal line. There are two types of inner loops: the time loop (L2x) and the element loop (L3x). The calculation in the element loop is localized, which does not need other element information. This loop can be vectorized by carefully masking the if-statements and aligning the critical data in memory. The time loop can be parallelized using OpenMP based on shared memory. The integration takes place using OpenMP reduction.

The sound prediction of installed jet cases was tested on the KNL nodes of the UK national high performance computing facility ARCHER. Each KNL node contains 64 cores with 16GB MCDRAM and 96GB DRAM. The memory in the tested node is in cache mode, where the high bandwidth MCDRAM serves as a layer between caches and the main memory DRAM. The vectorization test




Figure 8: Hybrid MPI+OpenMP parallelization with vectorization for KNL nodes

325 result is shown in Table 1. The element loops, L31 and L32, were vectorized
and the calculation time was measured for one observer location. The vector-
330 ization speedups for L31 and L32 are 1.61x and 8.55x respectively. L31 could
potentially be optimized for higher speed up using better memory alignment,
which requires further exploration. Overall, the current speed up 2.04x using
335 vectorization for the whole computation is acceptable. Based on these vector-
ized loops, the OpenMP parallelization is tested on one KNL node from 1 thread
to 64 threads. Figure 9 shows the OpenMP parallel scaling for L21, L22 and
the whole computation on one KNL node. L21 shows better OpenMP parallel
340 scaling, which compensates for the poor performance of L31 vectorization. L22
does not show a good scaling over 16 threads because the time derivatives calcu-
lation require unaligned memory access. The OpenMP threads are imbalanced
in L21 using the default scheduling, so the scaling can be significantly improved
by using the dynamic scheduling, which is based on a first-come-first-serve prin-
ciple. The scaling of L21 is super linear compared to the ideal scaling line.
345 The scaling performance of the whole computation begins to decrease from 32
cores, but still has a reasonable speedup of 51.88x using 64 cores compared to
the vectorized serial code. Compared to the original unvectorized serial code,
the overall speed up combining vectorization and OpenMP on 64 cores of one
KNL node can hence be approximately 103.8x. Therefore, the four and half day
sound prediction could be completed in around an hour using one KNL node
with vectorization and OpenMP and in around half an hour using two KNL
nodes using 2 MPI processes.

Table 1: Vectorization speed up per observer location

	L31	L32	Whole computation
Original	56.22s	20.40s	77.55s
Vectorized	34.87s	2.31s	38.11s
Speed up	1.61x	8.55x	2.04x




Figure 9: FW-H solver parallel scaling using OpenMP on one KNL node

2.3. Post-processing

Post-processing takes place after the simulation to analyze the results and extract the required knowledge from statistics, far-field acoustics and unsteady data.

2.3.1. Turbulence Statistics

For calculating turbulence statistics, the variables that are usually collected during the simulation are primitive flow variables ($\bar{\rho}$, \bar{u}_i and \bar{p}) and second order flow variables ($\overline{u_i u_j}$, \overline{pp} and \overline{TT}). Additional terms can also be recorded if turbulent budgets are required for analysis.

The statistics can be directly or indirectly obtained from the collected variables. The mean flow variables are $\bar{\rho}$, \bar{u}_i and \bar{p} , while the fluctuations can be computed by

$$\phi' \psi' = \overline{\phi \psi} - \bar{\phi} \bar{\psi} \quad (9)$$

where ϕ and ψ stand for primitive flow variables and an over bar represents a time mean. During the process, the statistics are also used iteratively to inform mesh refinement and FW-H surface adjustment.

2.3.2. Far-Field Acoustics

Far-field acoustic pressure fluctuations ($p' = p - p_0$) are obtained from the acoustic solver. To evaluate the sound pressure level, acoustic post-processing is performed to obtain sound spectra and the overall sound pressure level (OASPL).

Fourier transformation is used to convert the time signal of acoustic pressure fluctuations into sound power spectral density (PSD) $|\hat{p}'(\mathbf{x}, f)|^2$ at given observer points. The narrowband sound pressure level is expressed as

$$SPL(\mathbf{x}, f) = 10 \log_{10} \left(\frac{|\hat{p}'(\mathbf{x}, f)|^2}{p_{ref}^2} \right) \quad (10)$$

370 The OASPL can be obtained by integrating the sound pressure density over a given frequency range:

$$OASPL(\mathbf{x}) = 10 \log_{10} \left(\frac{\int_{f_{low}}^{f_{high}} |\hat{p}'(\mathbf{x}, f)|^2 df}{p_{ref}^2} \right) \quad (11)$$

where f_{low} and f_{high} constitute the frequency range, over which OASPL is obtained and the reference pressure p_{ref} is 2×10^{-5} Pa. OASPL is usually used for indicating the overall sound level and directivity.

375 2.3.3. Unsteady Dataset

Unsteady flow data is recorded in and around the jet plume, shown in Fig. 10. Analysis can be performed on this database to understand the sound generation process. The data is stored and managed using HDF5. As mentioned in the previous section, HDF5 allows parallel data reading and writing. It also provides
380 direct access to parts of the dataset instead of reading through the whole content. This is useful to perform large-scale unsteady data post-processing, because flow analysis can be directly performed in the specified time period at specific spatial locations in parallel. One example is space-time correlation analysis.

$$R_{\phi\psi}(\mathbf{x}, t, \delta\mathbf{x}, \delta t) = \frac{\overline{\phi(\mathbf{x}, t)\psi(\mathbf{x} + \delta\mathbf{x}, t + \delta t)}}{\overline{\phi(\mathbf{x}, t)\psi(\mathbf{x}, t)}} \quad (12)$$

where ϕ and ψ represent desired flow variables, e.g. velocity, turbulent
385 Reynolds stress and turbulent heat flux. The correlation analysis can be used to analyze turbulence scales (both time and space) and sound sources. This will be demonstrated in the next section.

Figure 10 shows that the calculation of the two-point cross correlations at the given point \mathbf{x} requires data from the adjacent spatial points $\mathbf{x} + \delta\mathbf{x}$ at
390 different time separations δt . By making use of direct access and parallel I/O features, each processor can read the required data of one given point from the same HDF file and compute the correlations in parallel. In addition to this correlation analysis, other in-depth analysis, such as modal analysis of sound-source-related instability waves, can also be performed.

395 2.4. Outlook: Automation and Data Mining

In addition to the design validation, high-fidelity simulation can be used to push the design towards more optimal levels, especially at off-design conditions. The models that are currently being used in the design system are not accurate enough to achieve this task. They do not incorporate the physical complexity
400 of today's design and suffer from a lack of reliable data. In the far future LES can be used in the final design phase to optimize the design, broadening the edge of the design envelope. In the near future, LES can be used to supplement




Figure 10: Example of unsteady dataset exploitation

costly rig testing and inform rapid design tools. The requirement of a short turnaround time is demanding in the design cycle. Automation of the whole process is necessary to reduce human error and used time. Automating the whole process will also enable optimization near the edge of flight envelope.

As this simulation generates larger volume of data than ever before, techniques to extract useful information are still in active development. Some data mining techniques might be employed in this scenario. Modal analyses (such as POD and DMD) can be used with machine learning to detect critical flow features. This can be integrated with experimental datasets to assist with quick design diagnosis.

3. Simulation examples

The parallelized jet noise prediction process has been outlined above. In this section, two examples are briefly introduced to demonstrate how the process is applied to the real jet noise problem of different levels of complexity.

3.1. Example 1: flight stream effects on heated jets

The jet flow exhausted from aircraft engines is at a higher temperature than the ambient air. The enthalpy fluctuations can produce extra sound, known as hot jet noise. Running heated jet tests is more difficult and expensive than regular isothermal jets. It becomes even more challenging by further incorporating a flight stream around heated jets in experiments. Therefore, the LES of heated jets in a flight stream is especially helpful to supplement experiment tests with more insight into jet flow field and its sound source.

The case simulated here is a single-stream heated jet with temperature ratio $\frac{T_j}{T_a} = 2.7$. The jet operates with a flight stream of $Ma_{flight} = 0.3$. In contrast, a static heated jet is simulated to show the flight stream effects. In the pre-processing, the mesh is constructed using hexahedral elements due to the relatively simple round nozzle geometry. A series of FW-H acoustic radial and closing end surfaces are placed around the jet for far-field sound prediction, see Figure 4(a). This enables a sensitivity study on the FW-H surface locations and closing strategies to obtain a reliable far-field sound prediction. The simulation is performed using hybrid RANS-LES with dual time stepping. After the flow transient phase, unsteady data is collected for $200 D/U_j$ for the post-processing of statistics, far-field sound and space-time correlations.

Statistics. The flight stream has a strong impact on jet flow in both mean flow and turbulence fluctuations. Figure 11 shows the near field statistics of heated jets with the static jet on the top and the flight stream jet on the bottom. The jet shear layers are highly stretched by the flight stream in the axial direction. A longer jet potential core occurs for the flight stream case. The turbulence fluctuations are also reduced by the flight stream because of the decreased velocity difference between the jet and ambient fluid. The simulation further confirmed the stretching theory of flight stream jets [31]. The stretching factor A is 1.7 for this heated flight stream jet [11].




Figure 11: Statistics of heated jets with and without flight stream

445 *Far-field sound.* The jet noise is projected to the far-field from the FW-H surfaces using the acoustic solver. Figure 12(a) shows the predicted sound directivity from two sets of FW-H surfaces. Surfaces 1 and 5 refer to the outer and inner most radial surfaces respectively. The surface locations are shown in Fig. 4(a) and the surface colors are the same with the lines in Fig. 12(a). Convergent far-field sound predictions are achieved for different radial surfaces and agree with the experiment measurements. The flight stream reduces the overall sound pressure level by 5-10 dB and alters the sound directivity towards downstream polar angles. Figure 12(b) shows the sound spectra at a downstream polar angle of $\phi = 30^\circ$. The predicted spectra agree with measurement up to $St = \frac{fU_j}{D} = 1.5$ at this polar angle. The spectral shape is similar for the two jets but the peak value is reduced around 10dB by the flight stream.

450

455




Figure 12: Far-field acoustics for heated jets with and without flight stream

Space-time correlations. It is difficult to measure the unsteady near-field flow data in experiments for this heated jets. The unsteady dataset provided by LES is valuable to investigate the near-field jet turbulence and sound sources. An example of space-time correlations is given here to analyze eddy convective speed and indicate time and length scales in the shear layer, as shown in Fig. 13. The

460

jet axial velocity U/U_j and eddy convective velocity U_c/U_j are plotted along the nozzle lip line. A flow measurement of another hot jet at slightly higher Mach number at test point 46 of [32] is plotted for comparison, showing the validity
 465 of this simulation. Eddy convective velocities can be calculated from the space-time trace of correlation peak values. The eddy convective velocities have the same trend with axial velocities but are much closer to the axial velocities in the flight stream case. This indicates that the shear-layer eddy convective velocity is greatly influenced by the flight stream. The correlation contours for both flight
 470 stream and static cases is shown at $x/D = 10$. The correlations decays more quickly with time and space separations in the flight stream case, indicating a smaller length and time scales at this given axial location.




Figure 13: Space-time correlations and eddy convective velocity along the nozzle lip line

This flight stream jet example demonstrates that LES can accurately predict jet noise with a flight stream and compensate experiments with high-fidelity
 475 near-field flow information. This enables in-depth flow and sound source analysis.

3.2. Example 2: Jet and wing interactions for Ultra-High Bypass-Ratio (UHBPR) engines

The aircraft engine and airframe design becomes more coupled when moving
 480 to the future generation of UHBPR engines. In order to keep a safe clearance from the ground, the engine is installed closer to the wing. This causes non-negligible aerodynamic and aeroacoustic interactions compared to conventional configurations. To understand and evaluate the interaction and its resulting noise, detailed flow fields are clearly in need. However, it is very expensive and
 485 difficult to measure the flow field around this complex configuration. The unsteadiness of the jet and wing interactions also make RANS struggle to provide reliable predictions. In this circumstance, the hybrid RANS-LES is a cost-effective method that enables designers to assess the high fidelity flow data and evaluate the acoustics.

490 There are many challenges in this simulation: complex geometry, turbulence simulation strategy and far-field sound prediction. In order to model the industrial-level complex geometry, modular structured-unstructured mesh, illustrated in Fig. 3, is employed to obtain required grid quality and resolution in different flow regions. To maintain an affordable cost, RANS is used to model
 495 the inner boundary layers on the nozzle and wing surfaces, as shown in Fig. 5. For far-field sound prediction, FW-H surfaces were set outside the active flow region using vorticity and turbulence intensity as criterion, as shown in Fig. 4(b).

The simulated cases presented here are to predict installation noise and
 500 evaluate the potential of serrated nozzles to mitigate the installation effects on flows and aeroacoustics for UHBPR engines. Figure 14(a) gives an overview of the flow and acoustic field for the installed jet. The jet flow interacts with the wing near the trailing edge (TE), generating installation noise. This is the major sound source generated by installing the engine near the wing. The jet noise is also reflected by the wing and visible beneath the wing. In addition
 505 to these, the jet noise escapes from the gap between the wing leading edge (LE) and the engine and acoustic waves get scattered around the LE. The serrations have been designed on the nacelle lip to reduce jet-wing interactions. Figure 14(b) compares the turbulent kinetic energy of the installed round and serrated nozzle jets. The serrated nozzle increases the jet mixing near the nozzle and then decreases the turbulence level before the jet reaches the wing trailing edge. The interaction has been reduced near the wing TE, as indicated by lower turbulent kinetic energy levels.




Figure 14: jet turbulent flows around the wing

The reduced interaction leads to a reduction in installation noise in the far-field. As the installation noise is known to be more pronounced at the upstream
 515 polar angles, Figure 15(a) shows the predicted far-field sound spectra at the upstream polar angle $\phi = 120^\circ$ at the sideline point. For the round baseline nozzle, the installation noise prediction agrees with the measurements. The reduction of far-field sound for the serrated nozzle jet is in the mid range of the

520 shown frequency. The sound sources are indicated using the amplitude of 4th order correlations inside the shear layer according to Goldstein acoustic analogy [33]. The reduction of sound source is seen on the major components R_{1111} , R_{1113} , R_{2222} , R_{3333} , R_{1212} and R_{1313} .




Figure 15: Installed jet noise and its sound source decomposition

525 As this high fidelity simulation can predict installation noise and be used to evaluate noise reduction strategies, this type of simulation may replace parts of rig testing in the design verification phase. The simulation turn-around time and cost is crucial for industrial use. The time scale of the simulation is compared with its counterpart rig test in Fig. 16. The time spent on each step in rig testing and numerical simulation is estimated. The estimation on numerical simulation is based on this installation noise prediction example. The mesh size is around 150 million cells and the simulation was performed using 4800 cores on CRAY XC30 cluster. The rig test takes about 2 months to get the designed product fully tested and analyzed while the simulation takes about 1 month to complete. The typical cost of this rig test is around £200,000, although it can include many test points. The simulation costs around £46,000 for this 2013-era machine. As the computing cost is decreasing with the advent of new architectures, the simulation cost is expected to be at least halved in the next 10 years. In addition to the turnaround time and cost, another key benefit comes from the detailed flow insight and the real flight condition that is difficult for rig tests to achieve. Therefore, this high-fidelity simulation has indisputable potential in the future to partly replace rig tests for validating the product design and to be used to inform low order design models.

4. Conclusions

545 Since experimental tests of jet engine exhaust acoustics and flow exhibit difficulties in turbulent flow field measurement and achieving flight conditions, high-fidelity eddy resolving simulations are proving to be a viable alternative for design validation. This paper presents the process of high-fidelity simulation of jet aerodynamics and aeroacoustics using hybrid RANS-LES with the FW-H method. The process consists of three key parts: simulation preprocessing, flow

Actions	Time scale	Actions	Time scale
Manufacturing	1 month	Pre-Processing	1.5 weeks
Assembling	1 week	Transient phase	1-1.5 weeks
Testing	2-3 weeks	Sampling	1 week
Data processing	1 week	Postprocessing	1 week
Total	~ 2 months	Total	~ 1 month

(a) Time breakdown for rig test (b) Time breakdown for simulation

Figure 16: Turn-around time estimate for installed jet test and simulation

550 and acoustic simulation, and data postprocessing. Through the three parts, critical techniques are developed for predicting industrial-level complex-geometry jet flows and acoustics accurately and efficiently in industrial design time scales.

The pre-processing part includes mesh generation and FW-H surface placement. With complex nozzle geometries, the modular hybrid structured-unstructured 555 mesh is implanted inside the jet plume in line with the shear layer, leading to better mesh resolution and grid quality. The placement of FW-H surfaces is demonstrated for both isolated and installed jets. The criteria is that all the sources are enclosed by the surface, for which turbulence intensity and vorticity magnitude can be used as an indicator. In the simulation part, hybrid RANS- 560 LES is performed to predict near-field jet turbulence and the FW-H method is used to predict far-field sound. The flow solver is an edge-based finite volume solver. LES employs a low-dissipative kinetic energy preserving scheme for better turbulence resolving. The LES zone is blended with wall-distance-based RANS for inner boundary layer near the nozzle and wing surfaces. Parallelization 565 is achieved using the Oplus framework for distributed memory machines. Good load balance is achieved using ParMetis for domain decomposition and processor communication latency is overcome by overlapping the halo node information exchange process with the internal edges/nodes computation. The solver shows a good parallel scaling over 8000 cores. Implementation of HDF 570 and direct parallel I/O achieves efficient parallel data output. For far-field sound prediction, the convective FW-H formula is used for the jet subject to flight streams. Hybrid OpenMP and MPI parallelism with vectorization is employed in the FW-H solver for parallelism on many-core computer nodes, i.e. intel KNL. This reduces the completion time of sound prediction job from 4.5 days 575 to about 1 hour using one KNL node and 0.5 hour using two KNL nodes. The efficient parallelization of flow and acoustics solvers greatly accelerates the simulation process as they takes more than half of the total process time. Finally, the postprocessing is performed on statistics, far-field sound and unsteady data. This can be used to inform the whole simulation process iteratively and analyze 580 the flow and acoustics for design validation and model improvement.

The use of the described jet noise simulation process is successfully demonstrated by two examples. The first is for heated jets with and without flight stream. The flight stream effects can be investigated in more detail using LES to compliment experiments. The predicted far-field sound agree with measurements. The LES confirms the stretching theory of flight stream effects and provides unsteady data for turbulent source analysis. The second example is the RANS-LES simulation of installed UHBPR engine flow and acoustics. The critical simulation techniques have been demonstrated through this example to make industrial-level complex-geometry simulation possible. The installation noise is predicted accurately and the potential of serrated nozzles are confirmed to mitigate installation effects by reducing the wing and jet interaction. The cost and turn-around time of this simulation are compared with the experiment counterpart. High-fidelity simulation shows indisputable advantage and a great potential to replace parts of rig tests for product designs in the future.

To boost future use of the jet noise prediction process, automation is needed to enable quick industrial design validation and off-design points optimization. As the simulation generates a large volume of data than ever before, advanced post-processing methods based on the simulation-generated unsteady database are required to be developed. Intelligent machine learning techniques could be used for knowledge extraction and process automation. This can then be integrated with experiment test data for fast flow diagnosis.

5. Acknowledgment

The work is performed as part of the EU-funded project JERONIMO (ACP2-GA-2012-314692-JERONIMO). We also acknowledge the computing time provided by the UK Turbulence Consortium under EPSRC grant EP/L000261/1 on UK national supercomputer ARCHER and the DECI resource Sisu based in Finland with support from the PRACE aisbl.

References

- [1] R. Tinseth, Current market outlook: 2017-2036, Boeing Commercial Airplanes, Seattle, USA.
- [2] D. Crighton, J. F. Williams, L. Cheeseman, The outlook for simulation of forward flight effects on aircraft noise, *Journal of Aircraft* 14 (11) (1977) 1117–1125.
- [3] J. Slotnick, A. Khodadoust, J. Alonso, D. Darmofal, W. Gropp, E. Lurie, D. Mavriplis, CFD vision 2030 study: a path to revolutionary computational aerosciences NASA CR-2014-218178.
- [4] D. J. Bodony, S. K. Lele, Current status of jet noise predictions using large-eddy simulation, *AIAA Journal* 46 (2) (2008) 364.

- 620 [5] J. R. DeBonis, Progress toward large-eddy simulations for prediction of realistic nozzle systems, *Journal of Propulsion and Power* 23 (5) (2007) 971.
- [6] J. Tyacke, I. Naqavi, Z.-N. Wang, P. Tucker, P. Boehning, Predictive large eddy simulation for jet aeroacoustics—current approach and industrial application, *Journal of Turbomachinery* 139 (8) (2017) 081003.
- 625 [7] S. Karabasov, M. Afsar, T. Hynes, A. Dowling, W. McMullan, C. Pokora, G. Page, J. McGuirk, Jet noise: acoustic analogy informed by large eddy simulation, *AIAA Journal* 48 (7) (2010) 1312.
- [8] D. R. Chapman, H. Mark, M. W. Pirtle, Computers vs. wind tunnels for aerodynamic flow simulations, *Astronautics and Aeronautics* 13 (1975) 22–30.
- 630 [9] P. G. Tucker, Novel miles computations for jet flows and noise, *International Journal of Heat and Fluid Flow* 25 (4) (2004) 625–635.
- [10] S. Eastwood, P. Tucker, H. Xia, High-performance computing in jet aerodynamics, in: *Parallel Scientific Computing and Optimization*, Springer, 2009, pp. 193–206.
- 635 [11] Z.-N. Wang, I. Z. Naqavi, M. Mahak, P. Tucker, X. Yuan, P. Strange, Far field noise prediction of subsonic hot and cold jets using large-eddy simulation, *Proceedings of ASME Turbo Expo* (2014) GT2014–25928.
- [12] Z.-N. Wang, P. Tucker, P. Boehning, Large-eddy simulation of the flight stream effects on single stream heated jets, *AIAA paper* (2017) 2017–0457.
- 640 [13] H. Xia, P. G. Tucker, S. Eastwood, Large-eddy simulations of chevron jet flows with noise predictions, *International Journal of Heat and Fluid Flow* 30 (6) (2009) 1067–1079.
- [14] H. Xia, P. G. Tucker, Numerical simulation of single-stream jets from a serrated nozzle, *Flow, Turbulence and Combustion* 88 (1) (2012) 3–18.
- 645 [15] S. Eastwood, H. Xia, P. G. Tucker, Large-eddy simulation of complex geometry jets, *Journal of Propulsion and Power* 28 (2) (2012) 235–245.
- [16] S. Eastwood, P. Tucker, H. Xia, P. Dunkley, P. Carpenter, Large-eddy simulations and measurements of a small-scale high-speed coflowing jet, *AIAA Journal* 48 (5) (2010) 963.
- 650 [17] Z.-N. Wang, J. Tyacke, P. Tucker, Hybrid LES/RANS predictions of flows and acoustics from an ultra-highbypass-ratio serrated nozzle, *Note on Numerical Fluid Mechanics and Multidisciplinary Design: Progress in Hybrid RANS-LES Modeling* (2018) 1–12(in press).
- 655 [18] H. Xia, P. Tucker, S. Eastwood, M. Mahak, The influence of geometry on jet plume development, *Progress in Aerospace Sciences* 52 (2012) 56–66.

- [19] J. C. Tyacke, M. Mahak, P. G. Tucker, Large-scale multifidelity, multi-physics, hybrid reynolds-averaged navier–stokes/large-eddy simulation of an installed aeroengine, *Journal of Propulsion and Power* (2016) 997–1008.
- 660 [20] J. C. Tyacke, Z.-N. Wang, P. G. Tucker, LES-RANS of installed ultra-high bypass-ratio coaxial jet aeroacoustics with a finite span wing-flap geometry and flight stream-part 1: round nozzle, *AIAA paper* (2017) 2017–3854.
- [21] Z.-N. Wang, J. Tyacke, P. Tucker, LES-RANS of installed ultra-high bypass-ratio coaxial jet aeroacoustics with a finite span wing-flap geometry and flight stream–part 2: chevron nozzles, 2017, pp. 2017–3855.
- 665 [22] I. Z. Naqavi, Z.-N. Wang, P. G. Tucker, M. Mahak, P. Strange, Far-field noise prediction for jets using large-eddy simulation and flowcs williams–hawkins method, *International Journal of Aeroacoustics* 15 (8) (2016) 757–780.
- 670 [23] R. Watson, P. Tucker, Z.-N. Wang, X. Yuan, Towards robust unstructured turbomachinery large eddy simulation, *Computers & Fluids* 118 (2015) 245–254.
- [24] P. Tucker, Differential equation-based wall distance computation for DES and RANS, *Journal of Computational Physics* 190 (1) (2003) 229–248.
- 675 [25] D. Burgess, P. Crumpton, M. Giles, A parallel framework for unstructured grid solvers, *Programming Environments for Massively Parallel Distributed Systems* (1994) 97–106.
- [26] W. Gropp, E. Lusk, N. Doss, A. Skjellum, A high-performance, portable implementation of the *mpi* message passing interface standard, *Parallel computing* 22 (6) (1996) 789–828.
- 680 [27] G. Karypis, METIS and ParMETIS, in: *Encyclopedia of parallel computing*, Springer, 2011, pp. 1117–1124.
- [28] M. Folk, A. Cheng, K. Yates, HDF5: A file format and I/O library for high performance computing applications, in: *Proceedings of Supercomputing*, Vol. 99, 1999, pp. 5–33.
- 685 [29] J. F. Williams, D. L. Hawkings, Sound generation by turbulence and surfaces in arbitrary motion, *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences* 264 (1151) (1969) 321–342.
- 690 [30] A. Najafi-Yazdi, G. A. Brès, L. Mongeau, An acoustic analogy formulation for moving sources in uniformly moving media, in: *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, Vol. 467, The Royal Society, 2011, pp. 144–165.

- 695
- [31] A. Michalke, U. Michel, Prediction of jet noise in flight from static tests, *Journal of Sound and Vibration* 67 (3) (1979) 341–367.
 - [32] J. Bridges, M. P. Werner, Establishing consensus turbulence statistics for hot subsonic jets, *AIAA paper* (2010) 2010–3751.
 - [33] M. E. Goldstein, A generalized acoustic analogy, *Journal of Fluid Mechanics* 488 (2003) 315–333.