White paper
ANSYS® Fluent® with PRIMEFLEX for HPC: Ensuring Product Integrity for the Automotive Supply Chain

New highly detailed methods for using high performance computing now allow manufacturers to get deeper insight into the behaviour of the most complex systems, to create parts and assemblies that perform more effectively over a longer lifecycle.

Content

Introduction

Where next for engineering simulation?

Vehicle manufacturers and their suppliers have been pathfinders in the adoption of engineering simulation. Indeed, designing vehicles today is unthinkable without the use of simulation tools such as crash simulation, structural analysis and fluid dynamics. Though simulation is widely embedded into the overall vehicle engineering processes today, yet in most design areas it is still used only to perform spot checks, i.e. to evaluate the design for a few operating conditions. As a result large swathes of the operating space of a design are often left unstudied, thus overlooking potentially devastating failure modes and missing prime performance improvement opportunities. Latest advances in simulation software and hardware technologies are making High Performance Computing (HPC) readily available to automotive companies and suppliers. HPC is helping them expand exploration of a products design space and to increase product performance, quality and reliability – whether to meet further regulations or increase competitiveness.

Challenges along the supply chain

Revolutionary changes in vehicle design spread throughout the supply chain. Vehicle manufacturers who are responding to increasing legislation and changing customer preferences are looking to their supplier base to step up in their ability to apply engineering simulation. At every stage in the supply chain the performance and reliability of the part or sub-assembly is now expected to be optimised in the context of the integrated system. Computational fluid dynamics (CFD) is a key enabler for such businesses since it allows them to study the interaction of their products within the complete system without the need for a physical testbed, and equally to provide assurance to upstream integrators of the behaviour those parts under a variety of operating conditions. HPC simulation with ANSYS Fluent is delivering this capability from smaller and large organisations, for designing a single component up to a complete assembly.

An exhaust system is one example of sub-assemblies in a modern vehicle, with its own particular challenges. But by looking at the options and solutions available to manufacturers of such individual parts we can identity patterns and lessons for organisations designing many other elements of the integrated vehicle. New methods for robust design and optimisation are increasingly delivering benefits for a range of vehicle parts. This study uses a specific model of a complete exhaust sub-system to demonstrate the way in which these methods can be applied, and how accessible these methods can be on current high performance computing platforms. This accessibility is achieved through a combination of simpler user working environments, pre-programmed methods and visual tools to interact with remote data. Together these allow more users to benefit from HPC, not just the most expert; productive from first login and able to regularly incorporate these tools into project schedules.

Challenges in a Typical Automotive Sub-System: The Case of Exhaust System Design

- **Project schedules**
  The final stage in the power system of a vehicle, the exhaust assembly can be literally the last component in the overall vehicle design schedule. Timescales become compressed and suppliers have to adjust to planning changes that can be propagated from any point in the upstream process. This can sometimes leave little flexibility for the company to maintain high quality, performance-optimised, designs.
- **Vehicle geometry**
  Vehicle form is settled before the exhaust assembly is fully optimised. Subsequent change to the exhaust design therefore have to fit within a fixed external space. Creating an effective system relies on an innovative approach to virtual design.
- **Legislation**
  Increased regulations around vehicle exhaust emissions and customer demand for fuel economy have increased the importance of the exhaust system's design. Development of advanced technologies such catalytic converters, aftertreatment devices, and ever quieter mufflers has now become an essential part of exhaust system development.

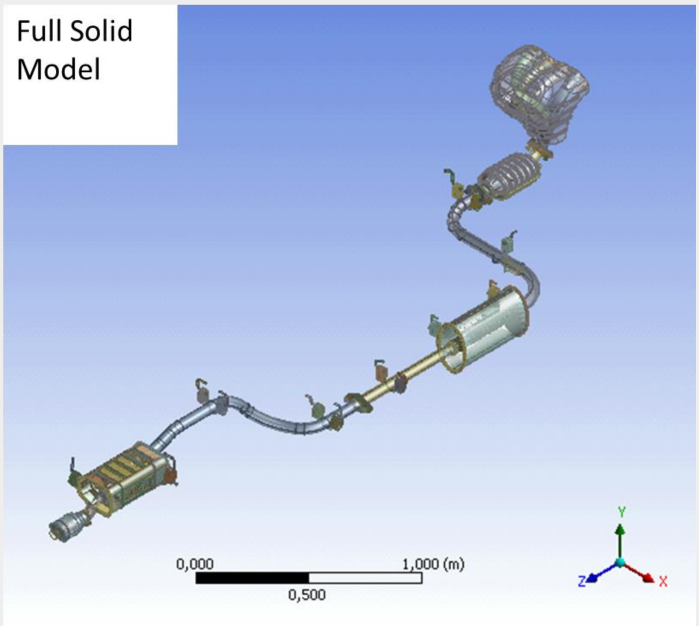

Image courtesy of ANSYS Inc.

Simulation of a fully integrated vehicle system

Across the automotive supplier industry we can expect to find a great variety of models corresponding to the different vehicle parts and simulated phenomena. Just within an exhaust system alone the level of detail is highly complex:

- Highly transient, turbulent flow
- Complex chemistry in catalytic converters and exhaust aftertreatment devices
- Mixing and multiphase flow in SCR systems
- Acoustics in mufflers

ANSYS and Fujitsu have partnered to study the particular characteristics of exhaust simulation workloads and define a set of PRIMEFLEX for HPC clusters optimized for such purposes. The baseline simulation studied the transient behaviour of the exhaust gases within a typical assembly of manifold, conduit and baffles. Optimisation was then done by applying variations around several dimensions concurrently:

- **Massflow** – as it varies over various engine operating conditions.
- **Geometry** – adapt the catalytic converter inlet to make flow completely homogeneous.
- **Catalytic converter resistivity** – this regulates the in-line pressure drop and eventually impacts flow distribution.

8 million cell exhaust model with transient analysis

Model setup

The basis of this study was a vehicle exhaust model representative of current production needs, with a geometry based on a full assembly from manifold through to muffler and outlet.

Mesh

- HEX/POLYHEDRA type
- Cells: 7,636,538 Faces: 43,707,080 Nodes: 33,387,905

Physics

- Transient simulation with explicit timestepping for engine startup cycle.
- Robust design based on geometry variations.

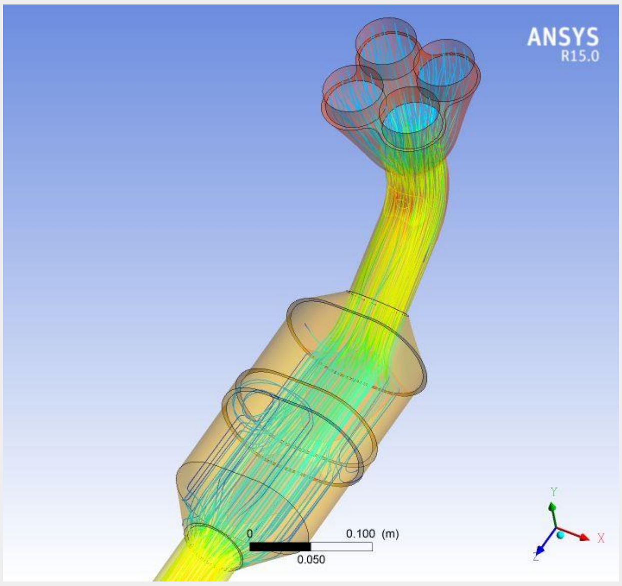

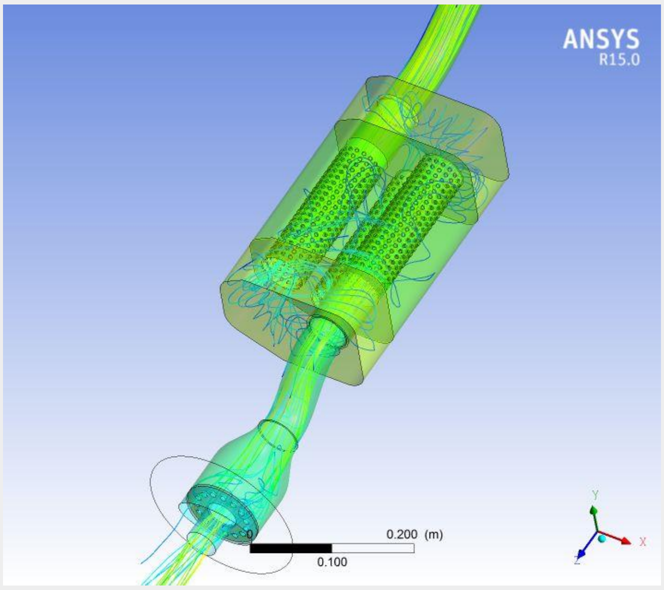

Images courtesy of ANSYS Inc.

The application used for this study was **ANSYS Fluent 15.0.**

The HPC system was a cluster of **FUJITSU Server PRIMERGY CX250** nodes equipped with dual Intel® Xeon® CPU processors. Comparisons were made with different processor types — varying frequency and core count — and interconnect. Parallelism was implemented with Intel MPI libraries.

Evaluation systems		
Node type	PRIMERGY CX250	PRIMERGY CX250
Processors per node	2	2
Processor type and frequency	E5-2690 v2 @ 3.00 GHz E5-2680 v2 @ 2.80 GHz E5-2670 v2 @ 2.50 GHz E5-2660 v2 @ 2.20 GHz	E5-2697 v2 @ 2.70 GHz E5-2695 v2 @ 2.40 GHz
Cores per processor	10	12
Interconnect	InfiniBand Gigabit Ethernet	InfiniBand Gigabit Ethernet
MPI libraries	Intel	Intel

Maximising parallel scalability

Running a single computation in parallel across multiple compute nodes is a defining characteristic of HPC workloads. More than one node becomes necessary when the turnaround time is so long that it prevents simulation of a given model within a reasonable project timeframe. There can be multiple networks that interconnect the nodes within one HPC cluster:

- A **user network** between all servers and compute nodes; uses Gigabit Ethernet technology
- A dedicated **high-speed network** to increase speed of applications when the cluster size increases; based on InfiniBand technology for low-latency, high-bandwidth for interprocess-communication between compute nodes, and sometimes for fast I/O to storage systems
- A dedicated optional network for **cluster management**; on small configurations these actions can share the user network

Measurements of ANSYS Fluent on the exhaust model demonstrated that an InfiniBand interconnect is required to continue reducing elapsed time as more nodes are used in a single parallel computation. For this exhaust model with 8 million cells and a HEX/POLYHEDRA mesh the ANSYS Fluent application demonstrates high efficiency up to 16 nodes with the 10-core processor (320 cores in total).The parallel efficiency is projected to remain high for even larger clusters, as observed in other benchmarks with ANSYS Fluent.

In comparison, scalability is much lower when just using Gigabit Ethernet, giving longer elapsed times and ineffective license utilisation. Transient simulation in particular benefits from the faster interconnect.

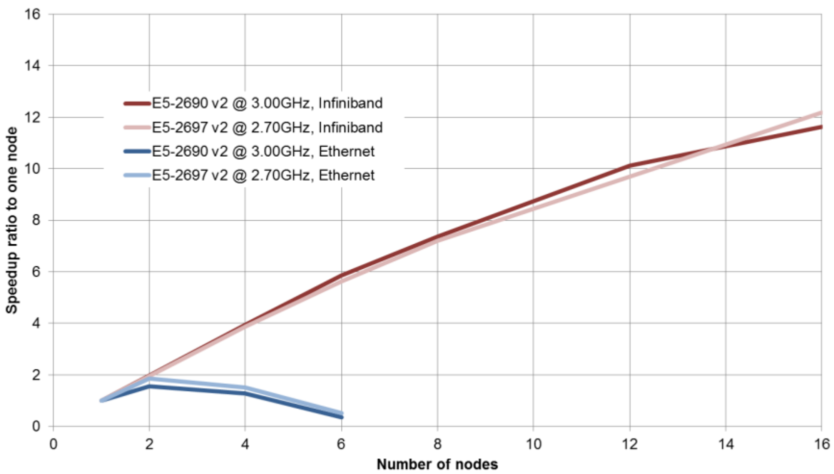

Processor types – performance versus efficiency

The choice of which processor type to use is a balance between various factors. We start with cores per job, since this is generally controlled by the number of application licenses. Conventionally, each ANSYS Fluent computation is parallelised so that all compute nodes used for the same job are fully occupied by the processes of that application execution. On todays multi-core processors, other policies are feasible that provide more balanced loading across the cores.

Across different core counts per job, the aggregate trend is for 12-core processors to deliver slightly lower throughput than 10-core processors. Moreover, the baseline 12-core price is higher, so it can be concluded that the nodes equipped with 10-core processors will generally be the preferred option in terms of aggregate price-performance.

An acceptable alternative is the E5-2650 v2 processor, with 8 cores and frequency of 2.60 GHz. Core for core, this node gives similar performance to the 10-core 2.5GHz processor. Eventually, though, more nodes will be required for a larger aggregate number of cores.

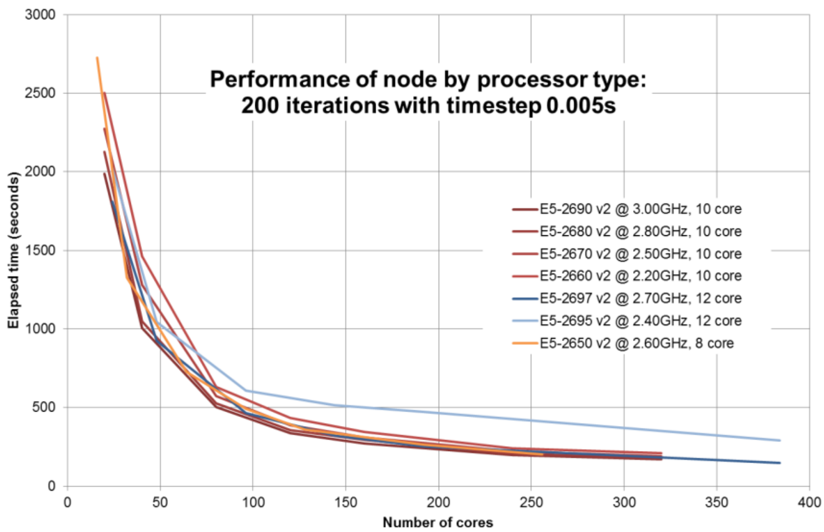

Optimal processor frequency

Evidently, the faster processor is expected to compute in the shortest elapsed time; but frequency is not the only factor. The computational algorithm of an application can depend more on other node hardware components. In many HPC cases memory bandwidth is most significant.

Higher-frequency processors are often exponentially more costly than others in the same generation. In studying the exhaust model, we compared in turn the times on 2.8 GHz, 2.5 GHz and 2.2 GHz processors with the reference Intel Xeon Processor E5-2690 v2 @ 3.00 GHz.If frequency were the only factor, then we would expect the ratio of elapsed times to equal the clock speed (inverse frequency) ratio. Instead, the elapsed ratio is actually less, meaning the node is computing relatively faster than the reference processor. So although throughput is less, all are relatively more efficient than the 3.0 GHz processor.

Still, where absolute speed is the requirement - as for full accuracy transient analysis – then the 2.8 GHz or 3.0 GHz processors remains the better option. Moreover, when application license cost is factored, then the hardware cost for higher frequency CPUs becomes relatively lower.

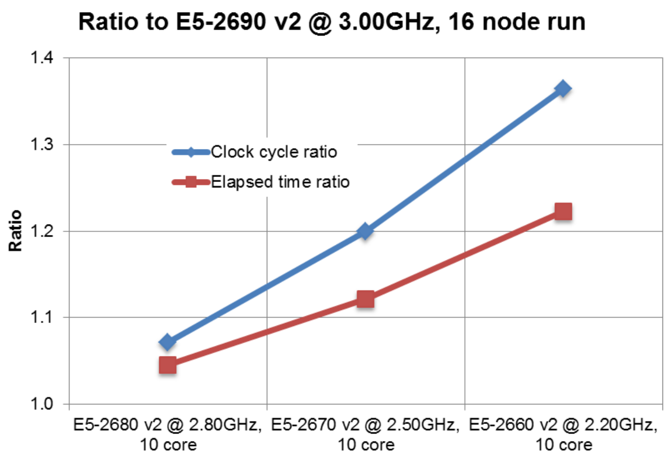

Revealing new detail – best practice in project simulation workloads

Raise quality with robust design optimisation

The first phase of a simulation project involves establishing the initial model. Both steady-state and transient methods will be applied to arrive at the model conforming to high-level design objectives. But this is just the start of the process. Designers and engineers then begin to look at a series of optimisation phases. A first phase may use a design of experiment (DoE) approach to automate the sweep across design variables for the steady state. This phase may be intended primarily to determine the ideal balance between various constraints, such as engine back pressure, noise, emission levels, material costs, and others. Then, ideally, a robust design study would assess the stability of the optimised solution to the external boundary conditions. Such optimising methods regularly generate tens to hundreds of further computational jobs for a given design, and may iterate with both steady-state and transient conditions.

In this example, an optimization study of the catalytic converter housing led to a 3.5% improvement in pressure drop across the converter – an enormous gain for fuel efficiency.

Getting full accuracy with extended transient simulation

Transient simulations reveal the detailed response of the assembly under the most critical dynamic conditions, important for exhaust system simulation where the massflow varies in direct relation with engine speed. For complete aftertreatment of engine exhaust the most critical period occurs during heating of the catalytic converter immediately after engine ignition. Conditions to obtain a fully accurate result therefore requires simulation of the transient flow behaviour for several minutes after engine ignition. For this study the following scenario was projected.

- Consider a four-stroke engine running at 2000 rpm for 5 minutes.
- The transient timestep is set to 0.005 seconds, so the total number of timesteps to run is 60,000.
- Test model was the 8 million cell exhaust case described above.

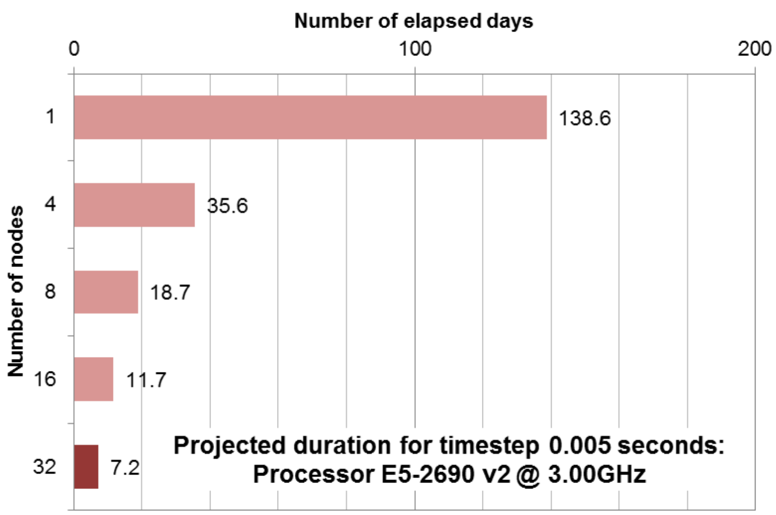

Projecting from the elapsed time time for 20 timesteps on a single node (E5-2690 v2 @ 3.00GHz, 10 core), and using an asymptotic compute rate of 20 iterations per timestep, a full accuracy transient simulation of the 8M cell model is projected to take many months of elapsed time on a single node. From the measured system we can expect that this time would fall to under 2 weeks using 16 nodes. At this point the parallel speedup efficiency of Fluent is around 75%; but we still prefer to reduce the elapsed time to a matter of 1 week. Applying a conservative efficiency of 60% we project an elapsed time of around 7 days on 32 nodes. And the HPC system can go further, indicating that 48 and more nodes would continue to reduce the elapsed time for a full accuracy transient simulation. Lower communication overhead, tuned run-time setup, and high parallel efficiency for larger models like an entire exhaust system means 64 nodes and more could still be effective at bringing down overall elapsed time for a full calculation.

Matched HPC system sizes

Different types of projects have differing simulation needs. Designing an exhaust solution for a single vehicle type may represent a simulation project that takes just a few weeks. The companies that service these projects range from engineering service companies to tier-1 automotive suppliers. From measurements on our exhaust reference model, we define matched HPC cluster systems for a normal HPC workload in common production scenarios in the following table. A speedup efficiency of 80% is assumed to encompass the mix of job types and parallelism.

Design workloads for different types of sub-assembly

Assembly type	Component (Muffler)	Sub-System (Exhaust Aftertreatment System)	Full System (Entire Exhaust System)
Overall project duration	2	3	4
Model size (number of cells)	5 Million	10 Million	30 Million
Steady-state simulation phases			**Effective number of jobs**
Problem setup (steady state)	2	5	10
Design of Experiment (steady state)	25	50	100
Optimisation	25	50	100
Robust design optimization (RDO), steady state	25	50	100
Transient simulation phases			**Effective number of jobs**
Problem setup (60 timesteps)	5	10	20
Design of Experiment (60 timesteps)	10	20	30
Full accuracy transient run (60,000 timesteps)	1	1	1
	Estimated total computational time mixing parallelism and job concurrency		
Elapsed hours on four nodes	539	1,107	3,466
Tuned cluster size – number of compute nodes	12	20	40
Total elapsed hours	**180**	**221**	**347**
	1.1 weeks	1.3 weeks	2.1 weeks

Delivering the value of HPC

Workload-based design

Fujitsu sector-designed PRIMEFLEX for HPC Clusters for the automotive industry are shaped around production workloads. For vehicle sub-assembly applications using ANSYS Fluent, the PRIMEFLEX for HPC cluster reference configurations deliver:

- HPC cluster components selected for optimal price-performance against sector workload, validated with direct application measurement and real models from the automotive environment.
- Integrated HPC architecture uniting hardware, system software and user-ready middleware: lowers acquisition risk and reduces upfront effort
- Factory-installed user environment for immediate project readiness and fast-start application usage

This sector-designed approach fully leverages expertise from low-level system tuning up to knowledge in the business layer. Benchmarking compares actual models on the widest range of system combinations to identify those most adapted to automotive sector needs, and considers how the user drives multiple jobs within an overall project. Given the large number of system combinations available today (processor type, memory, frequency, interconnect, storage), this workload-based philosophy gives confidence that the baseline system is matched to the purpose, backed with directly related expertise.

Dynamic scale

Adjustments to the baseline can be made before installation, and later during operations as the workload increases. Administration effort of adding nodes to the cluster is minimised through the graphical management desktop interface that displays an integrated view of the complete cluster. The built-in HPC Gateway interface with its application workflows means cluster dynamic changes have no impact on business end-use methods, maintaining continuity and minimising disruption.

Improving productivity and standardising optimal methods

Productivity is constantly eroded by users spending more time dealing with IT than designing — writing and constantly changing script, locating misplaced data, rerunning jobs that failed during initialisation. These issues are not solved just by providing a graphical interface for which the user still enters command lines and script extracts. Real customer value is when expertise in this application-IT layer is fully embedded in the HPC environment, so designers and engineers work with an environment that is fully relevant to their business. The HPC Gateway from Fujitsu addresses this productivity gap.

HPC Gateway – Industrialising Expertise for new and experienced uses

Running ANSYS Fluent in HPC Gateway

HPC Gateway provides a means to **industrialise expertise** in HPC usage, by capturing proven methods that can be transferred and utilised by a broad spectrum of users. Its web interface simplifies HPC cluster access from any client device. Workflow technology is used to encode optimal and expert methods to run various standardised application processes.

For the automotive supplier segment Fujitsu offers several pre-built methods to be used through the HPC Gateway, giving end-users a simplified interface to set up, submit and monitor work on the HPC cluster – from single jobs to multi-task concurrent optimisation. Packaged methods can be downloaded from the Fujitsu HPC Gateway Application Catalogue web site and self-imported into your local HPC Gateway.

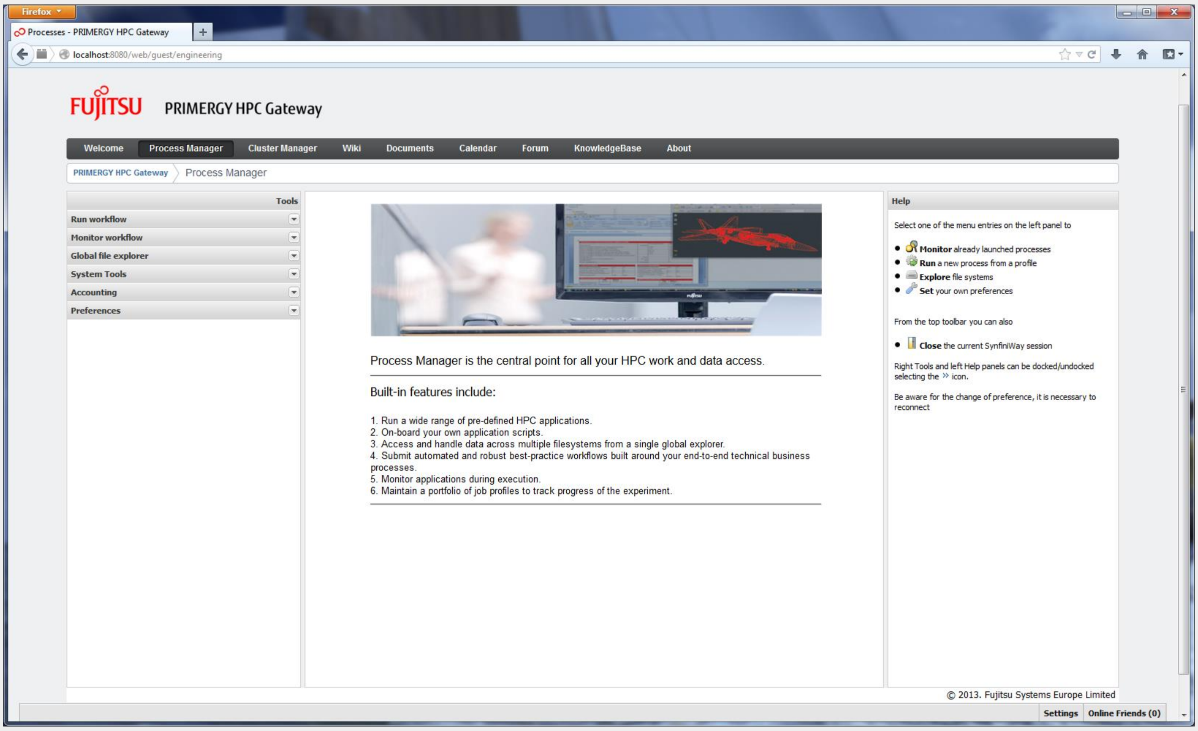

HPC Gateway provides many other benefits for both individuals and teams.

- Traceability is enhanced through the job profile table, which stores input settings for each run.
- Dynamic monitoring allows the user to graphically follow application progress through key simulation metrics.
- Intuitive web interface simplifies ease of use allowing a wider access to the power of HPC, and using robust expert methods.

Equally, with FUJITSU Software HPC Cluster Suite Advanced Edition users can develop their own workflows to encode and automate the HPC processes that form the unique competence of the organisation. Such automated workflows may connect other steps around the solver – such as input preparation, output filtering, report generation, and data archiving – and come to represent best practise or expert-based methods for raising productivity and democratising the use of HPC.

Exhaust system modelling capabilities

Transient simulation

Fully transient simulations of an exhaust system can serve three purposes:

- Validate if, for the design of the components concerning the flow optimization, an averaged steady-state approach is feasible. This requires the simulation of some engine cycles to reach stable cyclic conditions for the initial conditions. In this example, three cycles of the engine were chosen to conclude on the simulation time required.
- Perform heat-up simulations of the catalytic converter. Since the temperature range that is covered during heat-up is significant major changes in gas density (and other materials properties) occur which cannot be neglected to get accurate results. In this scenario 5 minutes real-time were assumed to be sufficient to cover the critical phase.
- Simulate aeroacoustics; an inherently transient process where pressure fluctuations create noise.

Heat-up simulation is the most challenging of these with respect to the real time that needs to be covered, so figures for this case will be shown assuming that full resolution of the engine cycle is required. This may not be mandatory, but in terms of numerical challenges gives the upper ceiling of the simulation ranges to be covered.

Automating design optimisation with ANSYS DesignXplorer – Maximise CAT volume utilisation

Among other design aspects, the catalytic converter inflow is critical for converter performance with respect to the degree of converter volume utilization. A highly uniform utilization is best with respect to catalytic material wear and converter efficiency.

As can be seen, changes in the shape of the inflow duct (in the example bumps from either top or bottom) lead to different flow patterns in the inflow, and subsequently to a different volumetric utilization of the first converter part.

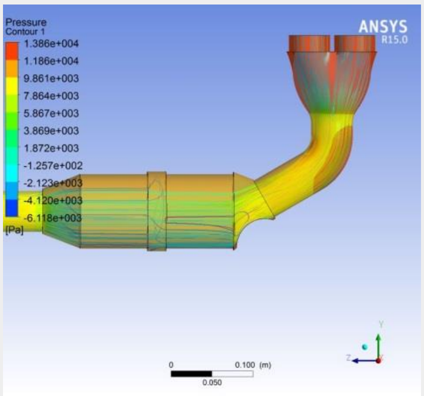

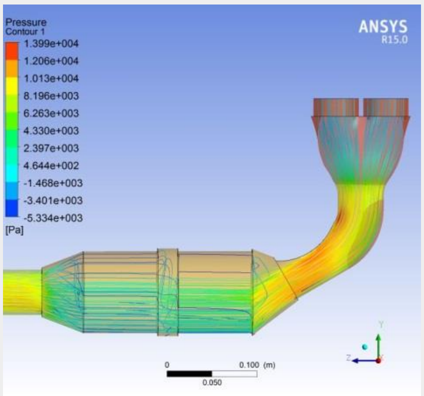

Robust design

Design exploration with the ANSYS DesignXplorer provides engineers the possibility to either perform a robust design analysis of components or the full system to ensure that the design is validated for a reasonable range of load cases. Alternatively this tool can perform automated optimization of either components or the full system by adding parameters to the geometric design of components and defining target functions to assess the performance of the modified configuration.

Ideally, when combining both approaches, robust design analysis will also give the engineer valuable information on the flexibility of the design with respect to manufacturing considerations - while sensitive changes need to be supervised and checked very carefully, less sensitive changes which give manufacturing or cost advantages could be allowed to give business advantages without significantly compromising on performance. In an industry like automotive industry where large production numbers of designed parts are common, the cost advantages of such evaluations could easily outweigh the cost related to the necessary computing resources.

Conclusion

A workload-based approach to system design clearly reveals the benefits of high performance computing (HPC) for automotive systems and components. In this joint study ANSYS and Fujitsu have architected a set of optimised integrated PRIMEFLEX for HPC clusters through direct study of a realistic example of an automotive system – the exhaust system, including all its primary components – with the aim of assisting both new and existing HPC users. This comprehensive review illustrates that today it is possible to lead in-depth investigation of automotive sub-systems including robust design optimisation for steady-state and transient cases in a time frame compatible with the length of the project.

An expansion of capacity to handle unsteady simulation will provide more accurate information on transient fluid behaviour. This is essential when studying the heat transfer and acoustic behaviour of the exhaust system. For businesses familiar with HPC the systems outlined here will ease selection of the baseline components. And with workflow automation through the HPC Gateway, these organisations can codify their established methods and best practise, and expand access to a broader range of teams and projects.

Organisations new to HPC, possibly looking to increase model resolution or add a robust design capability, will gain from the validation of application efficiency and the pre-installed system environment. But more significantly, with the integrated HPC Gateway web environment, the entry point to HPC is lowered. Prebuilt application workflows can be imported into the local HPC Gateway installation, allowing end users to launch and monitor standardised methods, reducing risk and ensuring productivity from first login.

Contact
FUJITSU TECHNOLOGY SOLUTIONS LTD

Website: www.fujitsu.com/fts/hpc
2014-05-27 EMEA EN