

WAVE-RT at Ford

When WAVE-RT was released in 2008, it was bound to cause a stir in powertrain engineering circles. Last year Ford began an evaluation programme to put Ricardo's new software to the test and find out just how effective WAVE-RT could be in accelerating new engine development programmes

Before WAVE-RT was launched in 2008, getting prompt answers to questions early in the development cycle of a new engine took time. Even then, current simulation tools were limited, especially when it came to software-in-the-loop (SiL) and hardware-in-the-loop (HiL) testing to validate new engines. The look-up maps traditionally used to model an engine in HiL tests cannot represent the engine in the detail demanded for the latest designs and the mean-value engine models used instead are unable to capture wave dynamics and do not respond well to transient events.

Ed Hernandez works in a department at Ford which is responsible for powertrain NVH and calibration. Hernandez describes himself as "working on the front line" with a supervisor, chief engineer and director at the top level. His chief engineer was very impressed with what products like WAVE had done for powertrain NVH – especially in predicting what the air induction system and tailpipe would sound like. "A lot of resource had been saved in time and materials by using analytical predictions and he soon began asking why the same techniques couldn't be applied to calibration," explains Hernandez.

Building the model from WAVE


Hernandez has almost 20 years' experience in engine design, with a background in the design and release of induction systems and camshaft

profiles; he had already used 1D codes, such as WAVE, to design engines. He talked to his colleagues in calibration to find out more about their processes and learned that these existing 1D models are not yet leveraged for use in calibration. When the team is happy with the engine performance it is mapped, and its performance regressed into a series of mathematical surfaces; it is those surfaces that are used by the calibration community for their work. This means an engine has to be built and survive the mapping process to generate the responses.

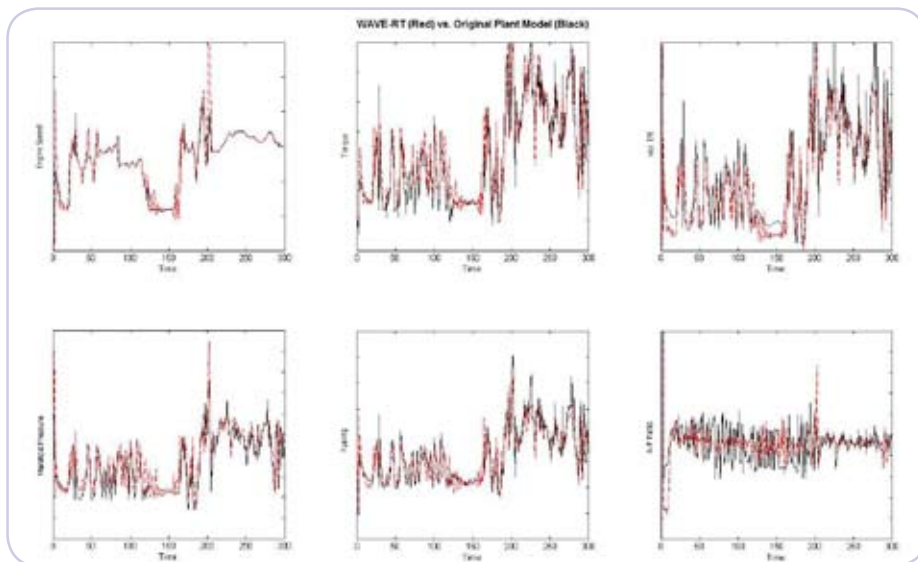
"It occurred to me that if we used the WAVE to design that engine in the first place, then WAVE-RT allows us to build a model in the absence of hardware and facilitate the mapping process. There's a lot of opportunity for work to get done ahead of time and even that working calibration can be validated back to the mapping engines we've built and run," says Hernandez.

So that's what the Ford team set out to do in 2009: evaluate RT by taking the WAVE donor

implementation in their new product development process, including for software-in-the-loop and hardware-in-the-loop validation of new engine products.



model that was used to design the engine, then use it to generate the WAVE-RT model. This would start the team on the first of the three steps they normally take in the design and build of an engine. The first is to examine the feasibility. "Does the process even work at all or does it break down straight away at a fundamental level?" asks Hernandez. Step one, feasibility, was established in 2009 and step two was to take a "deeper dive" to correlate the results. "Can I get results that are directionally correct or, better still, are absolutely correct?" If successful, the third step will be to implement the process.



Ford WAVE-RT validation results (in red) compared with the original map-based model

An important aspect of the feasibility study was that the team had never before used models like these at zero rpm and normally the engine would be running at a steady state. In contrast, this entire simulation was transient and changing continually. The first major question was whether or not the RT model could be run at zero rpm from key-off through cranking and a few hundred rpm as the engine starts. Hernandez found this was possible and set about investigating whether the model could be made to idle and respond to driver commands. Within four months this stage of the evaluation had been completed successfully.

A model at the push of a button

WAVE-RT models are generated at the press of a button within WAVE itself and Hernandez wanted to check that this basic function worked to his satisfaction. Not all the sensors in the early WAVE models exported but Hernandez found he could edit the 'C' file that WAVE produced to fine tune the final result to his requirements. The RT model was then exhaustively checked against the WAVE donor model. For example, Hernandez wanted to be sure that if he told the donor model to run at a certain engine speed and instructed the RT model to do the same thing, then it would do so. "I didn't want to do detailed correlation work," he says, "I simply said, 'here are the same set of inputs, what are the outputs?'"

Hernandez selected an engine of interest and got hold of a new WAVE model from Ford's engine division. He exported that and fine-tuned the exported RT model code as before, then replaced the mapping-based model within the simulation with the RT model. So where he previously had an engine model in his simulation based completely on hardware, he now had an RT model. This established that the existing vehicle simulation could communicate with RT and that signals were properly routed and named. With help from a simulation engineer, Hernandez made the signals generic so he could plug-and-play between the empirical model and the RT model. The vehicle simulation had in turn been built around the empirical model and had been "hand tailored for that," explains Hernandez. "We had to undo some of the 'hard wiring' and make it more flexible."

At that point Hernandez ran into a few difficulties and feared the worst. It turned out he'd made a small error in specifying the units for the

opening area of a turbocharger wastegate size which could be specified in square metres or diameter. The difference is clearly huge but Ricardo support came to the rescue. "I struggled for a couple of weeks before Ricardo figured out that I was telling it the wrong thing and the engine wouldn't survive the start cycle." Once the mistake had been discovered the engine started and ran. Hernandez could compile graphs for a presentation to his boss comparing the results from the empirical model and those of RT.

Feasibility was established using a Simulink model representing the vehicle and including details of its mass and driveline – from flywheel all the way to the point where "rubber meets the road." The model included values for inertia, details of the transmission and its calibration, powertrain calibration and the component that had been used before, the plant model. This was an empirical plant model based on the engine mapping data. To produce this an engine had been built and mapped, then the mapping data regressed into the plant model which was then "plugged into" the overall vehicle simulation.

The mapping data was produced using the US city and highway cycles. Drive cycles are performed with the car on a rolling road, the driver following a trace which has the vehicle driving along a route which includes stops and starts, flat roads and hill starts. Crucially, the cycles start from key-on. The vehicle simulation does exactly the same thing except the trace is digitised and the human driver is replaced with a virtual one, together with a controller within the model which depresses the accelerator pedal. The plant model responds with the torque which is transmitted through the transmission and driveline to produce the vehicle response.

Closed feedback loop

The drive cycle simulation itself is a closed feedback loop in that the simulation monitors the accelerator position and the vehicle speed against the drive cycle trace. "We were very happy with the way the model was able to follow that," recalls Hernandez, "which implies quite a lot of underlying success – that the torque was correct, the ECU (engine control unit) commands for spark and fuel were correct and the transmission was responding correctly to the right gearshifts. There are so many ways that could have broken down and it didn't."

Also, at around 25 seconds the engine makes the transition from open loop to closed loop fuel control as the catalytic converters have warmed up and the engine management systems begin listening to the Lambda sensors. At this stage the engine simulation has to provide the right chemistry to the Lambda sensors so that the ECU can believe it has gone into closed-loop fuel control, which was also achieved. "The graphs show how the engine is initially rich during the open-loop phase then at 20 seconds the graph snaps up to stoichiometry," Hernandez points out. For this trial he ran 300 seconds of the 1,400 second cycle in a successful demonstration to the rest of his organisation. He had learned about the export process, how to put it into the vehicle simulation, achieve driver closed-loop and finally, ECU closed-loop. "We were pretty happy with the endeavour," he concludes.

Hardware in the loop

At Ford of Dunton, Hernandez's colleague Andrew Emtage concentrated on HiL evaluation. WAVE models are used as part of the hardware development process already and there's a

team at Dunton using HiL controller software validation. "So I'm buying into a toolchain and a process which already exists, without giving that team extra work. I don't have to design my own process or develop my own team and that's my biggest benefit," explains Emtage.

"We can now respond much earlier using this process," Emtage continues. "In the past we had to wait until we had real engines and now we can have models running by the time the real engines arrive, if not before. It means I can add value earlier in the programme." Ideally, Ford would like to do even more of the calibration work without the need for expensive engine hardware. Emtage thinks it unlikely they will ever do away with engines at this stage completely "because there are other reasons why we need it."

Having worked with mean value models prior to RT arriving, Emtage is enthusiastic about being able to work with crank-angle resolved models. An example of the benefit is in simulating misfires which show up as acceleration and

deceleration of the crankshaft. This can easily be simulated by setting fuel to zero for one cycle on one cylinder leaving only friction and pumping torque. The ECU picks up the misfire as part of its normal function. Transient response of boost control is another area in which it wouldn't be possible to model as accurately using mean value models, as detail across individual cylinders might be lost. With RT, the behaviour can be scrutinised in detail on a cylinder-by-cylinder basis.

Towards desktop calibration?



All this was achieved within the evaluation licenses Ricardo provides and at that point Hernandez's chief engineer signed off the purchase of the first licenses. In early 2010, Hernandez moved on to try to demonstrate the correlation at a high enough fidelity to satisfy the calibration community "whose heads were already spinning with what we could use this for. The higher the accuracy we can achieve in RT, the more detailed the work we can do with the calibration."

The holy grail, believes Hernandez, is such that "you could one day do 100 percent of the calibration work on the desktop instead of in the vehicle. We're working our way up the ladder to discover how much fidelity is required from the various steps of the desktop calibration and how much data do we have available to achieve that correlation. Even if that goal is unattainable, the target is high enough that we can reap benefits from partial realization of that goal."

Helping shorten development programmes

When work is complete towards the middle of this year, plans are already in place to used WAVE-RT on a live engine programme. Hernandez expects the advantages could be many fold. As things stand, hundreds of hours are spent mapping a new engine, followed by a few hundred more on in-vehicle testing. Using WAVE-RT, it will be possible to create a model three to six months before a mapping engine comes into existence. Once these possibilities become a reality, the powertrain team will need to decide how best make use of them.

State-of-the-art analysis using FEARCE



Producing truly state-of-the-art product designs for complex thermal and mechanical systems – such as those of advanced powertrains and vehicles – requires an extremely diverse range of complex analytical operations. Ensuring that all aspects of performance and durability are delivered, whilst maximizing the opportunities for weight reduction, requires that operating temperatures, stresses, and fatigue lives are optimized within aggressive limits.

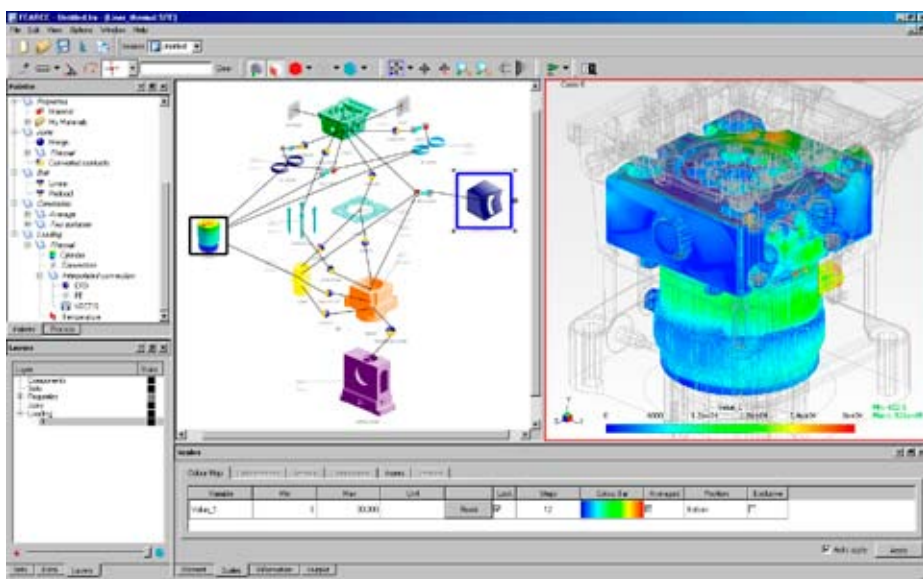
Industry CAE teams typically carry a wide range of pre- and post-processing software utilities

as well as their own preferred solvers for each type of analysis application. This complexity is compounded for structural analysis engineers – such as those at Ricardo – who need to serve multiple clients, each requiring use of their own preferred finite element solvers, mesh generation utilities and post-processors.

FEARCE was created by Ricardo as a finite element (FE) pre- and post-processing environment to support engine and vehicle analysis across a range of different simulation platforms. While its development was initially targeted at the requirements of the in-house team, expressions of interest from customers seeking a means

of automating their own analysis operations in a reliable and traceable manner led to the development of a commercial version.

The FEARCE package acts as an environment in which each stage of an FE analysis can be integrated, from the assembling of component models into larger systems, through the application of loads and boundary conditions, to the solution and post-processing of results. Automation is a crucial aspect of its design, enabling very complex analyses to be performed quickly and accurately whilst ensuring common processes are identical between iterations. As each stage of the process is clearly laid out in



Mapping point cloud data onto an FE mesh
(previous page); **complex model assembly and thermal analysis in FEARCE** (above)

the GUI and the result of each operation is fully documented with any conflicts or warnings clearly flagged, FEARCE is an ideal tool as a template for analysis procedures, allowing methods to be shared amongst engineers and groups.

In its pre-processing operations FEARCE works on a geometrically associative basis to apply thermo-mechanical boundary conditions and loads to models including, for example, applying the correct interface conditions at physical joints. By fully automating these tasks, FEARCE not only speeds up the process of model preparation but also delivers accuracy and repeatability that simply cannot be reliably matched with the alternative more conventional approach of manual analysis preparation.

FEARCE was designed from the outset to be used seamlessly in conjunction with third-party commercial FE solvers, allowing it to slot into any existing system. However it also incorporates a highly efficient linear solver, enabling many projects to be carried out entirely within the one environment.

In addition to the post-processing functionality that might be expected of any commercial FE results post processor, FEARCE also includes several specialist tools that are truly world class – for which the standalone commercial alternatives typically command high premiums. Amongst these is the FEARCE fatigue module that allows

engineers to perform durability predictions directly from FE stress analyses. With the recent release of FEARCE 8.0, users now have access to some major improvements to this module, including a number of new fatigue algorithms and the much-anticipated variance tool that enables assessment of the sensitivity of durability results to variations in material properties.

Also new for FEARCE 8.0 is the Structural Exploration (SE) solver. This post-processing tool physically removes parts of an FE model that aren't contributing work under specified load conditions. Using this tool, engineers can submit a model that represents the available volume in a design space and let the SE solver generate an optimal shape to meet the intended load conditions. In its initial release the SE solver provides results based upon displacement and frequency targets, but plans are already in place to develop this feature further to allow the optimization of structures in terms of thermal, stress and durability targets.

So from what started as an innovation aimed solely at internal needs, Ricardo's FEARCE product is now one of the most advanced commercially available FE analysis environments. And with the continued demands of an expanding user base, its capabilities look set to continue growing.

Japanese research group adopts WAVE-RT

Thus far almost without exception, every product evaluation of the WAVE-RT software has resulted in an order. As a result, the user base is expanding rapidly with some of the leading international automakers taking multiple licences in addition to truck and off-highway equipment manufacturers, motorsports teams, and various research departments. Given the new applications that real-time capability opens, this expansion includes new markets such as companies developing embedded software applications.

One of the new WAVE-RT customers is New A.C.E. Institute, a leading research group in Japan, which has chosen the software for advanced control system development. New A.C.E.'s research work focuses on diesel combustion, aiming to reduce exhaust emissions and increase thermal efficiency of future automotive diesel engines by using high boost and high EGR rates. For further information visit: www.nace.jp/E-Top.htm.

Eminox joins WAVE user base

Ricardo is pleased to announce that Eminox Ltd has joined the Ricardo WAVE/WB3D user base. Eminox is one of Europe's leading manufacturers of purpose-designed exhaust and emission control systems for commercial vehicles. The company has a range of innovative solutions, with systems designed and developed using market leading technologies, combined with applications engineering expertise. For further information see: www.eminox.com.

ISIS 1.6 released

The release of ISIS 1.6 has been announced, incorporating new features including a new R-Desk 'Project' plugin. This facility helps in the creation of project workflows, allowing engineers to graphically define a complex analysis process using multiple Ricardo Software and third party CAE tools, and helping ensure consistent and robust analysis processes.

For further information about Ricardo Software products, support services and CAE applications please contact:

CAE applications (Europe): Mark.Gregory@ricardo.com
CAE applications (US): Chi.La@ricardo.com
Software Sales: RS_Sales@ricardo.com
Software Support: RS_Support@ricardo.com
Or visit www.ricardo.com

