SURVEY OF COMPUTING PLATFORMS FOR ENGINEERING SIMULATION

EESI-2 Deliverable, April 2015

Lee Margetts[1,2,3]

[1] Chairman of the NAFEMS HPC Working Group
[2] School of Mechanical, Aerospace and Civil Engineering, University of Manchester
[3] Oxford eResearch Centre, University of Oxford

1. Overview

The aim of this article is to present some preliminary results for the recent NAFEMS/EESI-2 survey on "Computing Platforms for Engineering" simulation. The survey was prepared by the NAFEMS High Performance Computing Working Group in 2014; with input from colleagues in the European Exascale Software Initiative (EESI-2), the N8 HPC Service (UK) and Teratec (France). The survey was advertised to NAFEMS members and the broader engineering simulation community in September 2014 and data collection ended in November 2014.

NAFEMS is an international trade association, established in the 1980s, that focuses on promoting best practice in the use of simulation in engineering. It has 1200+ institutional members across an engineering ecosystem that includes blue chip firms such as Boeing, Airbus and RollsRoyce and micro-firms consisting of less than 10 employees. NAFEMS organises events, provides training, certifies engineers under the industry devised "Professional Simulation Engineer" scheme and publishes guides for the practitioner. All of these activities are centrally organised by a not-for-profit SME, with input from both regional and technical working groups.

The main motivation for the survey was to find out to what extent the CAE community is making use of HPC, the Cloud and other advanced computing platforms for engineering simulation. That said, the survey was aimed at all users of engineering simulation to ensure that, as far as possible, survey responses represented the engineering simulation community as a whole. Part of the survey was aimed specifically at software vendors, to try and get some indication regarding attitudes to large scale computing, particularly in areas highlighted as issues in the Exascale community.

In this interim report, survey responses are reviewed for the following topics: (i) the largest simulations carried out for the main types of engineering analysis; (ii) the maximum number of cores used and (iii) software vendor roadmaps. Items (i) and (ii) are presented for all respondents, then separately for industry and academia.

For a more detailed analysis of the survey data and a discussion about the views of software vendors, please look out for the author's talk at the NAFEMS World Congress (2015) and a new joint NAFEMS/EESI-2 publication expected later in 2015.

2. Respondents

A total of 231 respondents started the survey. This was thought to be a reasonable sample of the targeted engineering simulation community when compared with response rates for NAFEMS surveys carried out in the past. Figure 1 shows the segmentation of respondents by geography and Figure 2 by business area.

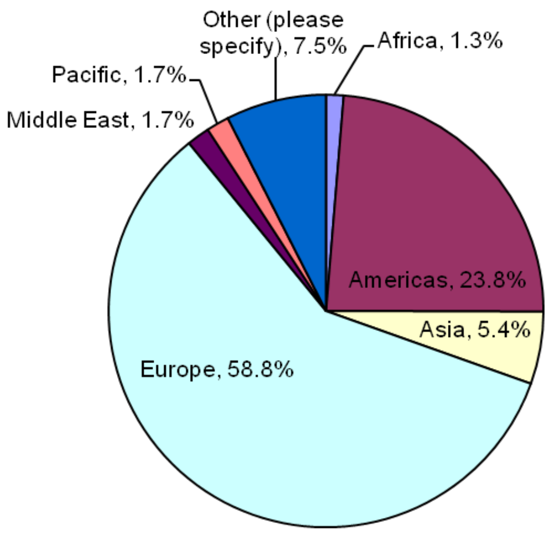

Figure 1: Respondents segmented according to location of headquarters

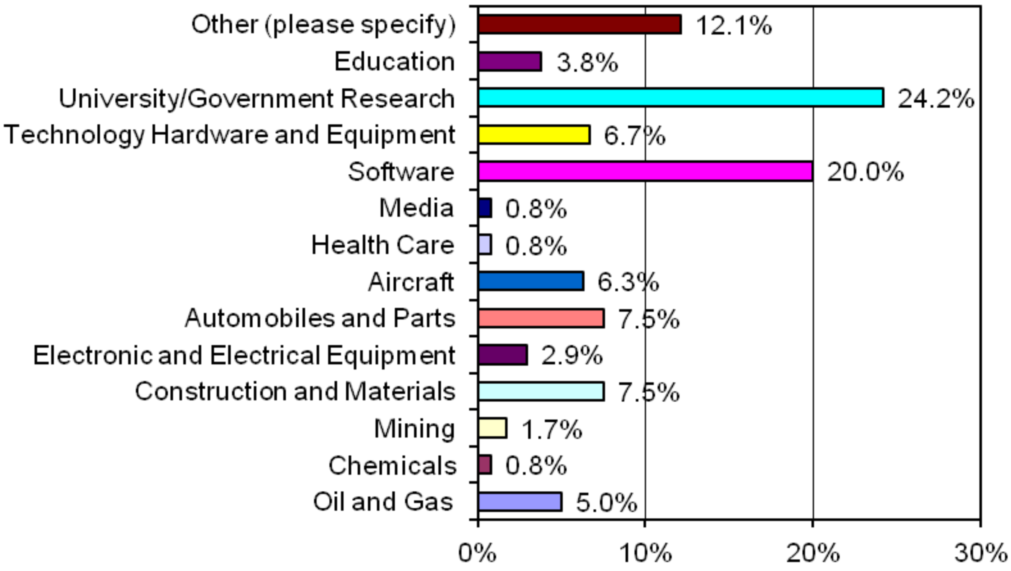

Figure 2: Respondents segmented by main business area

Responses were dominated by firms with headquarters in Europe (59%) and the Americas (24%). In terms of business area, there was a roughly even split between engineering firms (30%), the public sector (29%) and software/hardware vendors (28%). 13% of respondents marked their business area as "other".

There was an even split between small and medium enterprises (SMEs) with fewer than 250 employees (47%), and large firms with more than 250 employees (53%). The SME responses include micro-firms with fewer than 10 employees (17%), small firms with 10 to 49 employees (14%) and medium sized firms with 50 to 249 employees (16%). All percentages stated in this section are the percentages of all 231 respondents.

3. Matters of Size

One of the objectives of the survey was to get an idea of the size of the largest simulations carried out in the community and the size of the largest facilities used for those simulations. Respondents were asked to indicate the size of the largest simulation carried out in their organisation, measured in terms of the number of degrees of freedom, for a list of different types of engineering analysis. For example, in a two dimensional finite element model of a structure, displacements in the x-direction and y-direction are the degrees of freedom in the system of equations to be solved. The size of facilities is measured in terms of the number of cores in a standard multi-core processor. Respondents were asked not to include GPU cores in their estimates. This was to simplify the analysis of the survey data.

Figure 3 plots survey data for 5 bins of maximum problem size. The numbers on the scale indicate the percentage of respondents in a particular bin and add up to 100% for each simulation category. For example, adding the values along the CFD line, for each size of bin, gives 100% for the CFD data.

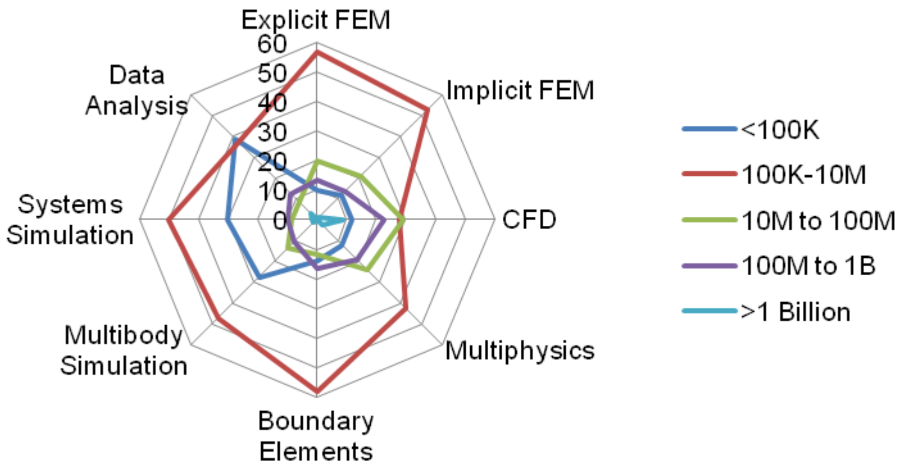

Figure 3: Maximum number of degrees of freedom by type of simulation

Figure 3 shows that the largest problems, for all types of engineering simulation, are mostly in the 100,000 to 10 million degree of freedom range. There is a strong bias towards smaller problems in data analysis, systems simulation and boundary simulation. In terms of the larger categories, there is a bias towards CFD and Multiphysics for problems >10 million degrees of freedom and only CFD has a relatively strong response (10%) for problems >1 billion degrees of freedom.

Figure 4 plots data for the maximum number of cores used in each of the categories of engineering simulation. There are four bins and as with Figure 3, the values along any line, corresponding with a particular type of simulation, add up to 100%. The figure shows a clear bias towards small core counts for Data Analysis, Systems Simulation and Multibody Simulation. The ranges 9-64 cores and 65-1024 cores are fairly similar. There is a small bias towards CFD and Multiphysics for 65-1024 cores and a stronger bias towards CFD for simulations using more than 1024 cores. The survey included a >8196 core bin, but only a handful of respondents ticked this box and therefore the data was added to the >1024 bin.

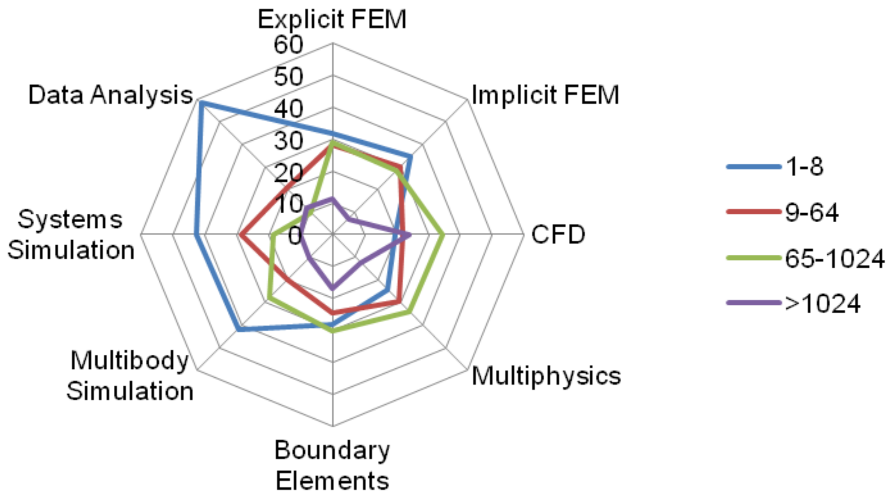

Figure 4: Number of cores used by type of simulation

4. Size in Focus - Industry versus Academia

Figures 5 and 6 segment the responses for maximum problem size and maximum number of cores, respectively, into separate plots for industry and research.

Considering responses from industry, Figure 5 shows that the largest problem sizes tackled in finite element and boundary element analyses are in the 100,000 to 10 million degree of freedom range. There is a strong bias towards smaller problems for systems simulation and data analysis and towards larger problems for CFD.

The profile of analyses carried out by respondents working in research (universities and government laboratories) is very different. Whilst problems in the 100,000 to 10 million degree of freedom range dominate, there is a higher proportion of respondents tackling larger problems.

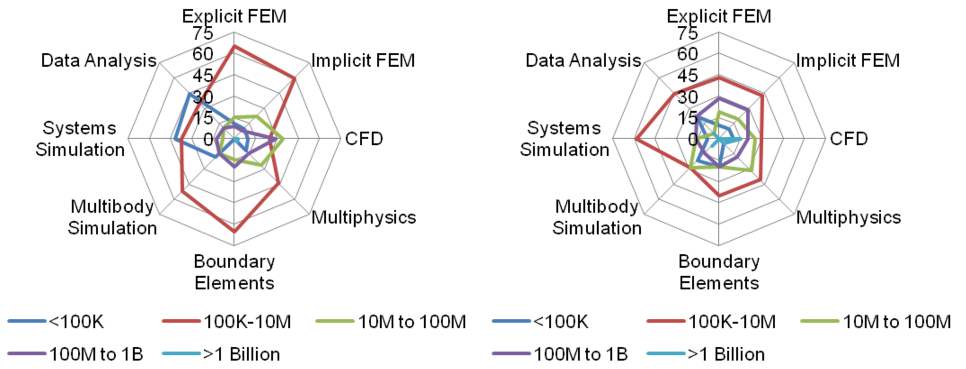

Figure 5: Maximum number of degrees of freedom solved by type of simulation. Comparison of responses from industry (left) and research (right).

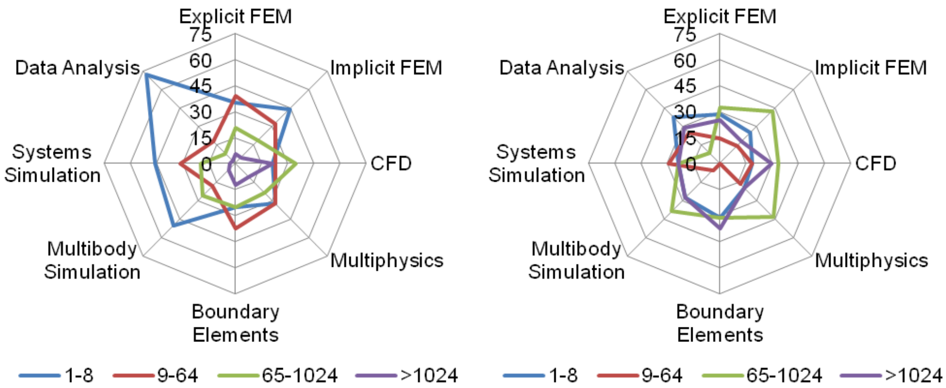

Figure 6: Maximum number of cores used by type of simulation. Comparison of responses from industry (left) and research (right).

Figure 6 is similar to Figure 5 in that it highlights strong differences in the maximum number of cores used between industry and academia. For all categories of simulation, larger core counts are used in academia compared with industry. The use of 1-8 cores in industry for data analysis, systems simulation, multibody simulation and implicit FEM is particularly striking. The differences may be due to different licensing fees for ISV software (cheaper for large core counts in academia compared with industry) or a greater tendency to use academic and/or open source software in academia.

5. Software Vendors

The survey included a section dedicated to ISV roadmaps. One of the questions asked whether the respondent's roadmap was in the public domain. An overwhelming ~90% of respondents said that it was not. The main question listed a small number of hardware and software technologies of interest to the Exascale community and asked respondents to indicate whether the technology was "Implemented", "Under development", "Planned" or "Not on the roadmap". The response "Don't know what it is" was also available.

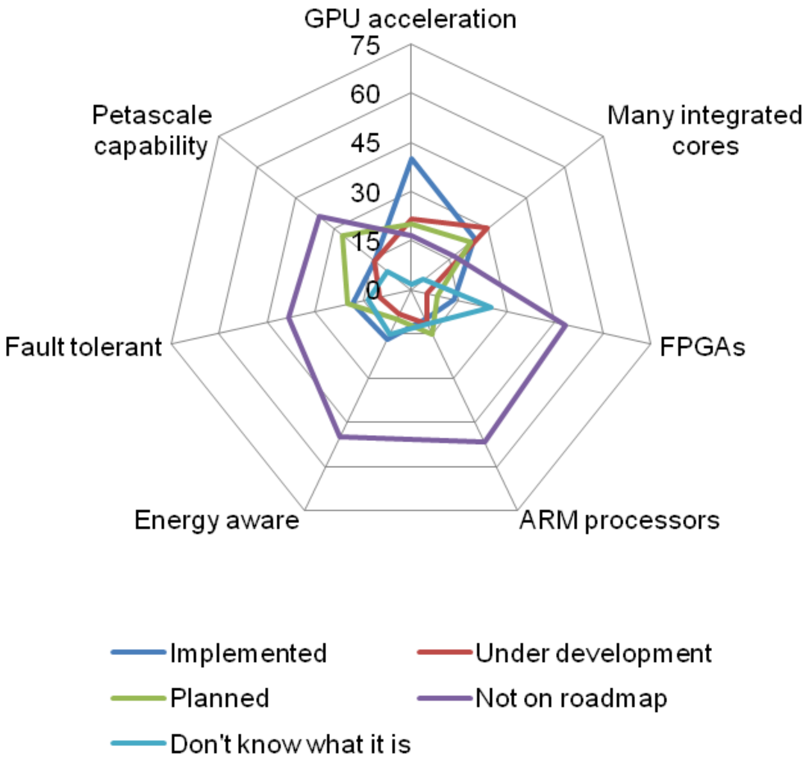

Figure 7: Roadmaps for organisations that produce software

Figure 7 shows the responses to the roadmap question. In the original question, both "Petascale capability" and "Exascale capability" were listed as technologies. As responses to both were small, they have been added together and presented under the category "Petascale capability". The response "Planned" was originally two separate responses "Planned 1-5 years" and "Planned >5 years". Very few respondents gave the latter response, so the data was added together to form a "Planned" response.

In Figure 7, there is a very strong “Not on roadmap” response for all technologies of interest to the Exascale community. It is also notable that “GPU acceleration” has a strong “Implemented” response as “Many integrated cores” has a strong “Under development” response. Both these technologies can be deployed in high end workstations; the largest end-user market for ISVs.

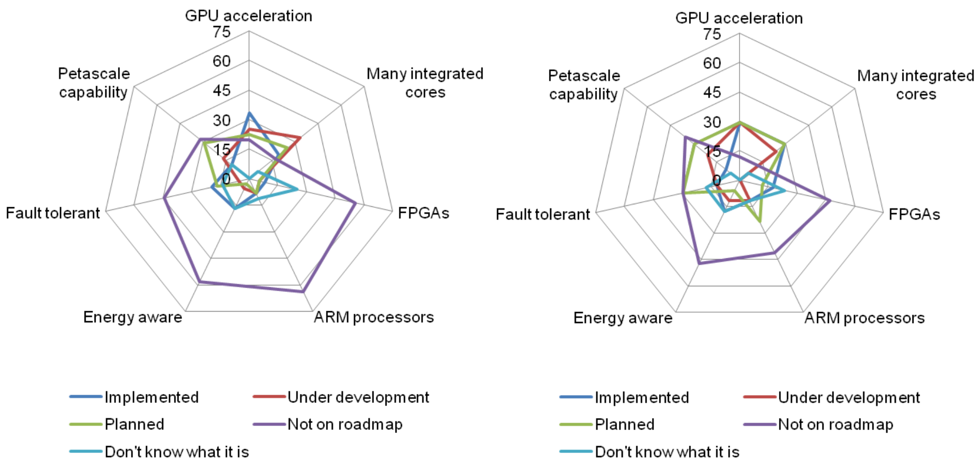

Figure 8: Roadmaps for organisations that produce software. Comparison of ISV responses (left) and OSS responses (right).

Figure 8 separates the responses for ISVs are open source software. The split is not perfect as some ISVs produce open source software and their responses are counted in both charts. The main difference between the two communities of practice is in the “Not on roadmap” response. The open source software community seems to have a roadmap that is adopting the listed technologies earlier than the ISVs.

6. Discussion

The survey provides some useful insight into the use of a range of computing platforms for engineering simulation. However, the results provide just one data point in a rapidly evolving era of computing. In the future, it would be useful to repeat the survey, perhaps every two years, to monitor trends. This particular survey asked respondents to reply on the behalf of their organisation and there was much discussion as to whether this focus was correct, particularly on social forums such as LinkedIn. Given that an individual respondent will have a much better idea of what they do than what their peers do, future surveys may focus on the individual’s usage rather than their assessment of use in their organisation.

The survey shows that problem sizes tackled by respondents are mainly in the range (up to 10 million degrees of freedom) that can be carried out on a workstation or a shared memory node on a cluster with up to ~32GB of memory. Whilst the comparison between industry and research shows that there is a shift towards larger problems in research, the problems are still not very large and around one third of respondents are running problems in the workstation range. The larger problems solved by respondents working in research organisations may be enabled by more favourable licensing terms from ISVs and may indicate a wider use of open source and/or software developed by the researcher.

The strong bias towards small problem sizes and low core counts for systems simulation is notable. Systems simulation involves running complex models comprising interconnecting components that make up a sub-system or system in a machine. Smaller models enable engineers to keep to small core counts, stay within memory constraints and run their analyses in a reasonable time. It is the author’s opinion that this area of simulation is probably artificially suppressed in terms of its potential to benefit from HPC. Scaling up each of the system components to a model resolution typically used by engineers when used separately (for finite element or CFD calculations for example) may require Petascale and/or Exascale capability to keep runtimes reasonable for engineering design.

In terms of facilities, there is a danger that the survey results will be out of date quite quickly. Future surveys may have to deal with a broader range of hardware; for example desktops, HPC clusters and cloud computing platforms comprising “standard” processors packaged together with co-processors, accelerators and other technologies that the hardware vendors may bring to market. Comparing the computational power of these systems may require a new effort from the engineering community on benchmarking; for computational speed rather than simulation accuracy.

Finally, the author acknowledges that further processing of the survey data may lead to further important insights. For example, there may also be enough data to assess whether there are significant differences between the responses received from Europe and the Americas.

7. EESI Recommendations

There are a number of recommendations that EESI would like to make on the basis of this interim report.

1. Improve engagement with trade associations.

Engagement with NAFEMS members has provided industry insights that have not been gained elsewhere by the Exascale community.

1.1. Increase efforts to engage with industrial trade associations across a broad range of business sectors.

1.2. Disseminate EESI findings and recommendations through industrial trade association publications. The responses “Not on roadmap” and “Don’t know what it is” for Exascale technologies are significant; indicating that the Exascale message might not be getting through to organisations carrying out software development.

2. Support academic and open source software.

ISVs are motivated by quarterly sales and are unlikely to invest significantly in emerging technologies, including novel hardware or new algorithms.

2.1. This report provides further evidence that ISVs and industry should be supported in the future by investing now in academic, open source software. Once the emerging technologies become mainstream, the ISVs will have solutions at the ready.

2.2. Investment in software for systems simulation is particularly recommended as it offers a unique opportunity to make use of large HPC capability by releasing suppressed capability in existing software components. Systems simulation joins together FEM, BEM, CFD and other technologies in a workflow to look at virtual machines, rather than virtual components for machines. It is clear from this survey that the current individual capabilities of FEM, BEM and CFD in terms of problem size and core counts, when used to look at components, are not used when these technologies are used to look at machines. Whatever the type of engineering simulation carried out, it seems to be mainly confined to a workstation or shared memory node. Thus systems appear to be over-simplified. R&D in larger-scale systems simulation could be carried out re-using existing software components, but adapted for emerging technologies.

This is a different philosophy to pushing a single CFD simulation, for example, to Exascale.

3. Attempt to disrupt the market.

3.1. Some ISVs have commented that they do not have access to large HPC systems, particularly for emerging technologies, for testing. Easier access mechanisms to academic facilities would help ISVs test and optimise their software.

3.2. The engineering simulation community does not have a suite of "typical" benchmarking problems for "performance". It is therefore difficult for end-users to compare/evaluate the performance of different combinations of software and hardware for their particular problem. Investment in domain specific benchmarking suites may help firms indentify ISV software with the best price vs performance vs feature. Those ISVs would out compete those that are slower in adapting to the support of HPC platforms.

Acknowledgements

The author would like to acknowledge the work carried out by the NAFEMS High Performance Computing Working Group, both in designing the survey and analysing the results. The author also acknowledges support from Teratec, France and the "European Exascale Software Initiative" project. The latter was funded by the European Commission under the 7th Framework Programme.