Ramsay Maunder ASSOCIATES

Prepared: 23rd May 2016

Finite Element Specialists and Engineering Consultants

Raising your Game with Good Simulation Governance

Angus Ramsay

1. Introduction

As anyone who attended last year's NAFEMS World Congress (NWC2015) will have observed, the sophistication of computational simulation software and hardware is moving ahead at an extremely rapid pace. In watching presentations and speaking to delegates, a new variant on the familiar phrase 'the democratisation of knowledge' was heard for the first time by the author. This phrase was ***'the democratisation of simulation'*** and it was understood, by the author, to mean something like the acquisition and spread of simulation tools amongst the non-specialist practising engineer [1].

There are obvious significant commercial advantages for engineering companies embracing and utilise numerical simulation in their design processes. If the democratisation of simulation means the development of bespoke software tools for the design of a particular product that encodes sound simulation governance principles then this is a laudable aim. The tool will be safe for use by any practising engineer and it will preserve the specialist design knowhow of the company into the future. A more worrying trend, though, which might also be described as the democratisation of simulation, is the selling of cheap and cheerful generalist finite element software or, possibly, access to such via the cloud. The reason for concern here is that it raises the spectre of all sorts of verification issues with the finite element software and the distinct possibility of ***finite element malpractice*** by non-specialist users of the software.

Peter Bartholomew's interesting Benchmark article, [2], "NAFEMS: the early days", reminded readers of John Robinson's thoughts of some 40 years ago. John, who was one of the founders of NAFEMS, noted in the early 1970s *"… that both coding and modelling errors were commonplace and only time separated the [simulation] community from* ***computer-aided catastrophe*** *[CAC]"*. Just such an incident of CAC did occur in the early 1990s when the Sleipner Platform A sank in a Norwegian Fjord, [3]. No one was injured but the estimated cost of the incident was some $700m! The subsequent inquiry found that FE modelling local to the failure had been inadequate, under-predicting the shear forces by some 45%. This, together with inadequate reinforcement detailing in the failure region, was identified as the cause of the failure.

There is an idea, expressed very elegantly in Henry Petroski's book, [4], that major failures, at least for bridges, have been observed to be spaced at approximately thiry year intervals. The reason for this is postulated as being the result of a '***communication gap'*** between one generation of engineers and the next; the *raison d'être* why structural members or components were designed in the way they were, being lost.

As the Sleipner incident happened around a generation ago, and since there remain useful lessons for all engineers to be gleaned from a re-examination of what went wrong, the author has chosen to revisit the Sleipner failure in the sixth NAFEMS Benchmark Challenge – interested readers will find the challenge in this edition of Benchmark and the formal response in the October edition.

In the simulation industry we are facing a somewhat more severe example of Petroski's generation gap. The simulation 'experts' of today are, often as not, and as pointed out in the NAFEMS overview of democratisation [5], either early [finite element] ***pioneers or direct disciples*** of them. They have often written and verified finite element software and are familiar with the practices of good simulation governance. Such specialists are becoming rarer as the opportunities to spend time researching (academic and industrial) is declining. For this reason alone the development of bespoke software tools that encode not only sound simulation governance principles but also a company's design knowhow is certainly a safe and efficient way forward for the industry.

In the environment outlined above, the practising engineer using numerical simulation needs to exercise great care in his/her work. The engineer may have little more than a vendor's training course as his/her background in simulation. His/her manager may well not have any background in numerical simulation either and in some cases might not even be an engineer. How then can engineers protect themselves against the myriad potential mistakes that lead to finite element malpractice?

There is much that the practising engineer can and should do to protect him/herself from perpetrating a simulation *faux pas* and this article provides, in a generalist fashion, an overview of the ideas of verification and validation, error estimation and convergence that are generally applicable in all forms of simulation.

Equally, and possibly more importantly, is the attitude that the engineer has towards his/her career. Numerical simulation tools are an excellent means for developing one's understanding of both the physics of a particular problem and how it might be captured accurately by simulation. After all, everything else being equal, the numerical simulation tool should be capable of converging to the exact theoretical solution whether this is known or not. Engineering is generally a commercial activity and this means, often, that the practising engineer will be under pressure to deliver in the shortest time and at the least cost. This necessarily places a potential conflict between getting the job done and getting it done right. Satisfying these two criteria simultaneously is possible, but the ability to do this only comes with experience. The engineer early in his/her career will find this conflict particularly difficult not least when those managing might have no concept of how extremely inaccurate the results from a poorly verified numerical simulation might be. The implications or consequences of inaccurate results vary depending on the engineering venture. Inaccurate CFD simulation of a turbomachine is unlikely to impinge on anyones safety although it could, if the resulting design gets as far as manufacture, have a sever negative economic impact on the company involved. Conversely, inaccurate simulation of an engineering structure might mean that the structure fails and people are injured or killed.

Whilst the practising engineer might not be proud of his errors, he should hold up his hands and admit them. He can then be proud that he has owned up, shared his error with his peers and learnt from them. The author would like to share with the readers a really basic error that he committed as a final year undergraduate conducting a finite element simulation of a convoluted bellows expansion joint designed to be a flexible member in an exhaust system for a submarine to take up thermal expansion and mechanical vibration, [6].

(a) Original Bellows Joint (b) Unloaded Joint in Testing Machine (c) Strain Gauges

Figure 1: Evidence of the author's error in plastically deforming a bellows expansion joint

The eagle-eyed reader will note that the unloaded expansion joint sitting in the Instron testing machine is rather longer than the original member. The reason for this is that the author, in pressure testing the system and in the vigour of youth, had plastically deformed the member by over-pressurisation. From recollection he even thinks that he was sitting on the member whilst he pumped it up!! This was a rather fundamental error in the author's understanding of Newton's third law of motion – action and reaction. As the submarine manufacturer didn't want the bellows back there was no significant consequence of this error and it even made the placement of strain gauges simpler so the author got away with his mistake. It is pertinent to point out, though, that bellows members can undergo a buckling phenomenon known as 'squirm' and that was the sort of failure that caused the catastrophic accident at the Flixborough chemical plant which killed 28 people in North Lincolnshire in 1974, [7].

Whilst the ***European Court of Human Rights*** would, no doubt, baulk at such a retrograde idea, the practising engineer might best be served by adopting the Napoleonic Code of Jurisprudence for the results of a numerical simulation:

Guilty until Proven Innocent!

2. Verification and Validation

A sound engineering design can be considered as supported on the 'three pillars' of *numerical simulation*, *mathematical models* and *physical experiments*, as suggested in figure 1. The goal of good ***simulation governance,*** [8], is to match numerical simulation with physical experiments. This cannot be done in any rational manner without consideration of the mathematical models that govern the physics of the problem – if that pillar is removed in the figure then the structure falls down and engineering design is no longer supported. Whilst the two terms, at first sight, might appear to be almost synonymous, the processes of verification and validation have distinct meanings, [9], and provide the foundations for engineering design by tying together the three pillars as indicated in the figure – these cannot exist without the mathematical model.

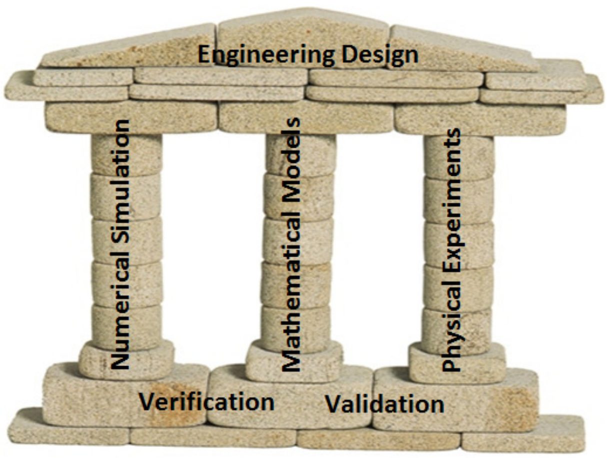

Figure 2: Three pillars of engineering design on the foundations of Verification and Validation

Validation

For a practising engineer involved with simulation, validation may not be part of his/her work. Nonetheless, he/she should be interested in validation from the point of view that it may influence the mathematical model on which the simulation is or needs to be based. There are also situations where a form of ***virtual validation*** can be undertaken by the engineer. An example of this was presented in [10] where a three-dimensional numerical model was used as a surrogate for physical experiments to obtain the real warping behaviour of an I-beam so that a model using beam type elements could be verified.

The process of validation also, necessarily, deals with the quantification of sensitivity or uncertainty in parameters that influence the mathematical model. These might include, for example, the variability in the material properties of rolled steel plates. This subject relies heavily on statistical approaches and requires thinking of the problem in a stochastic rather than a deterministic manner.

Verification

Verification, as presented in figure 1, can be considered as two distinct parts:

1) **Software Verification** – here problems with known theoretical solutions are analysed with the software to ensure that the software can recover the theoretical solution. This uncovers software bugs and educates the engineer on how the solution procedure used in the software performs, e.g., what level of mesh refinement is required to obtain good engineering accuracy.
2) **Solution Verification** – here the actual problem, which generally has no known theoretical solution, is studied to see how the solution quantities of interest converge with mesh refinement.

Assuming the mathematical models to be correct and all else being equal, then verification is a matter of controlling discretisation error.

Software Verification is a task that should primarily be conducted and publically reported by software vendors if the engineering community is to have any faith in the software tool – an exemplar of published verification can be seen in [11]. The practising engineer should be able to repeat these to ensure that the software is functioning correctly on his/her computer. It also provides invaluable experience as to how the various solution parameters need to be set in order to achieve results to a suitable level of engineering accuracy. Problems with known theoretical solutions are available in the field of elasticity, and ideas like the ***Continuum Region Element Method***, [12], have formalised the practice of this form of verification. In other fields of engineering, particularly where flow phenomenon are to be simulated, such theoretical solutions do not abound and here the software vendor, and the practising engineer, will need to make use of the ***Method of Manufactured Solutions*** (MMS), [13].

Solution Verification, essentially, involves the engineer studying, for their actual problem, how the quantity of interest converges with increasing levels of mesh refinement. A sound engineering report should contain, generally in appendices, evidence of both software and solution verification so that the reviewer is left in no doubt that this exercise has been carried out.

3. Error Estimation

The ideal situation, when conducting numerical simulation, would be able to predict the exact error. Clearly, however, this would require the exact solution and if one had that one would not need to conduct numerical simulation in the first place! Error estimation is a set of ideas and techniques, aimed at providing an estimate of the discretisation error and it is often considered as two distinct topics; ***a priori*** error estimation, which is conducted before a simulation and ***a posteriori*** error estimation, which is conducted after a simulation has taken place and when the results are available.

A Priori Error Estimation

It might seem rather paradoxical that one can even attempt to predict the error before one has even conducted an analysis but it is often possible to make some sensible predictions and some of these ideas are discussed below.

Experience gained from previous work

The experienced gained from previous work is invaluable to the practising engineer in providing an *a priori* estimate of the sort of model (level of mesh refinement, type of element etc.) that will give a solution with an appropriate level of engineering accuracy. This is particularly the case where the engineer specialises in numerical simulation of a specific type of component. This experience may, of course, be extended, with due diligence, to other similar components.

Appropriate degree of element

One can often make some prediction of the degree and number of elements required to obtain an accurate solution. When, for example, a structural engineer is faced with a frame structure where the members are loaded with uniformly distributed loads, then he will know, *a priori*, that the bending moments are going to be quadratic. If the engineer has a beam element in his FE system with quadratic shape functions then he will obtain an exact solution with a single element. If, however, he uses the same element for a frame where the loads are say linearly distributed, then he will not expect an exact solution from a single element and will need to consider mesh refinement. This is an example of 'knowing your element's capabilities' and single element tests are invaluable to the engineer in developing this knowledge and experience.

Quality of mesh

What is meant by the ‘quality’ of a mesh? Well, firstly, it might mean different things for different types of numerical simulation. For CFD formulated with a finite difference approach, the grid or mesh quality is measured in terms of orthogonality of the grid. Many years ago, the author spent a (miserable) week massaging a finite difference mesh for the flow around an aerofoil. The software provided grid quality measures and the idea was to optimise these by moving nodes around. It occurred to the author at the time that this was a rather futile exercise, which could have been automated by the software vendors (if they had been bothered). As the problem possessed no known theoretical solution, the futility of the exercise was exacerbated further by the fact that one could not assess whether or not the ‘optimised’ grid had produced better or worse results than the ‘unoptimised’ grid!

In finite element analysis, the situation is different in that grid orthogonality is not generally an issue. It was, though, in earlier times, where the number of elements in a mesh was severely limited by computational resources, a continual area of consternation for the practising analyst. The finite element model of a double Hooke’s joint, figure 3, was one of the first commercial analyses the author conducted back in 1989. The mesh took weeks to build (with the author going home each evening feeling like he has had his head in a bucket of spaghetti – no perspective in the graphics!!) and computational resources were such that the number of elements (higher-order hexahedral and tetrahedral elements in this case) was limited to the sort of number seen in the figure. Even so, the analysis took in excess of 12 hours to complete on a state of the art VAX computer. It is clear to the author that whilst this analysis did, with good spatial accuracy, predict the region of high stress that is likely to have caused the initiation of fatigue failure seen in the photograph, it is highly unlikely that the magnitude of the stress would have been accurately predicted by such a coarse model.

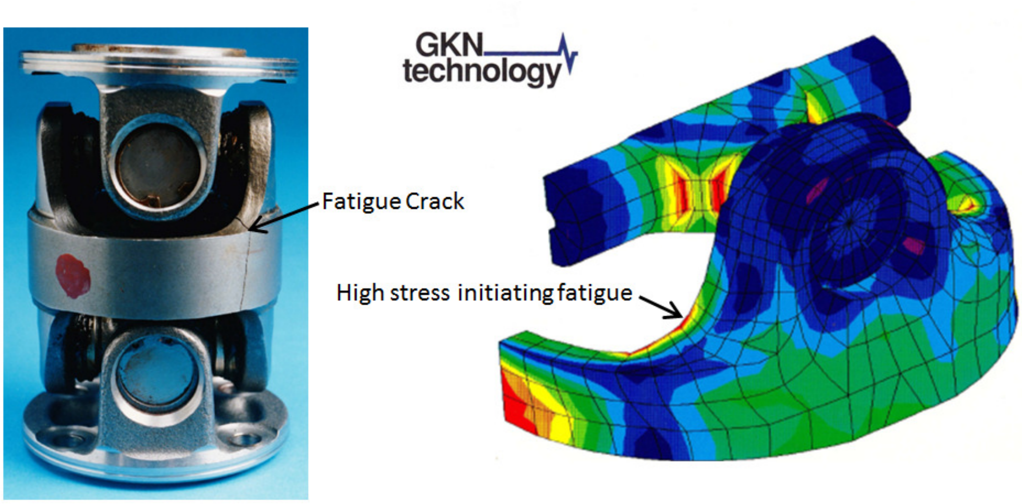

Figure 3: A 1989 FEA of a quarter model of a Double Hooke’s Joint

Much was made of how the performance of elements was affected by distortion. Physically meaningful measures of shape were developed, e.g., aspect ratio, taper, warping etc., and the performance of elements examined as they were distorted [14]. This led to a range of ‘rules-of-thumb’ which engineers were supposed to adopt and which software vendors encoded into their software. Distortion on its own, however, is no indicator of how the element will perform in any given situation; if the element is only feeling a constant stress field, as might be expected as a mesh is continually refined, then any sensible degree of distortion will not influence its ability to recover the constant stress field. For coarser meshes, then, different distorted meshes will produce different approximations but without conducting an optimisation study, e.g., by relocating the nodes to produce an optimum solution in some quantity of interest, it will not be possible to say which mesh is best, c.f., the CFD meshing problem mentioned earlier. The problem shown in figure 4 is a supersonic flow problem (derived using the method of manufactured solutions) the theoretical solution of which has a step change in density at an oblique shock. The mesh used for the solution has been adapted through nodal relocation from a uniform 16x16-element mesh, [15].

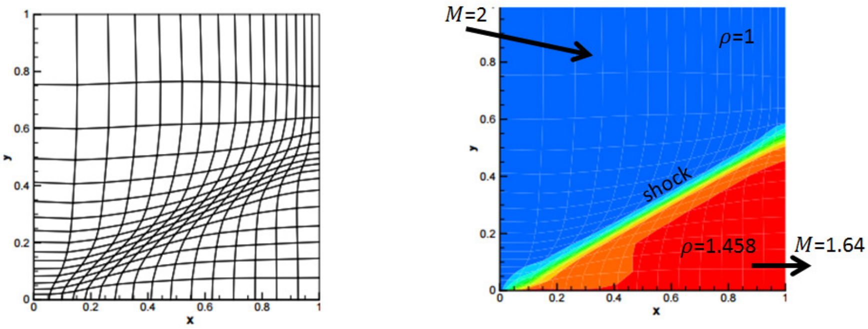

(a) Adapted mesh (b) Density contours

Figure 4: Node redistribution to optimise the solution

Rate of Convergence – Theoretical

The mathematical theory of simulation techniques enables firm statements to be made on the theoretical rate of convergence of quantities of interest. Although, in practice, the observed rate of convergence might not be identical to the theoretical rate (which is after all a limiting value as the mesh is highly refined) it is useful to know, for example, that if the theoretical rate of convergence is two then conducting a uniform mesh refinement should reduce the error by a factor of four.

Exact solutions for certain quantities

In the finite element analysis of linear-elastic continua, it can be shown that, provided consistent nodal forces are used corresponding to the applied tractions and body-loads, the finite element stresses are exact in an average sense. There are fields of mechanical engineering where this is rather useful. For example, the average hoop stress in a rotating disc, say from a turbomachine, is recovered exactly irrespective of the element type (lower or higher-order), the number of elements in the mesh and the distortion of the elements. This is extremely useful to engineers since this quantity may be compared with the strength of the material to provide a good estimate of the burst speed of the disc. This idea of exact quantities being recovered from an otherwise approximate numerical solution can be extended, with due caution and understanding, to other engineering applications. Where this extension is not valid, then it can lead to the sort of problem seen with the Sleipner incident.

A Posteriori Error Estimation

This form of error estimation occurs once the simulation has been completed and the results are thus available. The first question one might ask is, in the absence of the exact solution, how can one expect to improve on the simulation results such that a sensible error estimate is possible?

One way in which this may be done is to make use, where this is possible, of super-convergent properties exhibited in some element formulations. In the finite element analysis of linear-elastic continua, the Gauss (integration) points corresponding to the reduced integration scheme have stresses that converge at a higher rate than other points within the element, i.e., they are super-convergent. In the super-convergent patch recovery (SPR) scheme of Zienkiewicz and Zhu, [16], and adopted by more than one vendor of commercial software, a patch or stress surface is fitted to the stresses at Gauss points surrounding a node. A new nodal stress is then recovered with, hopefully, superior accuracy to the stresses at the nodes.

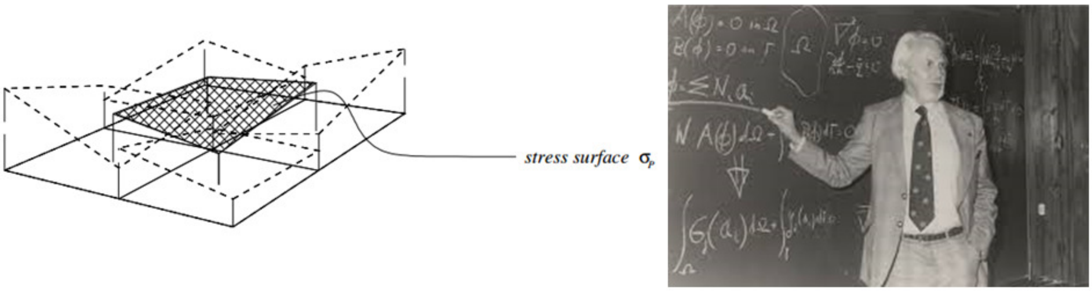

Figure 5: Zienkiewicz's Superconvergent Patch Recovery Scheme [17]

The SPR scheme has been demonstrated as effective in that it can drive an adaptive scheme to refine the mesh where the error is greatest. The adapted mesh for the supersonic flow around an ogive (nose cone) with a 10-degree incident angle is shown in figure 6, [18].

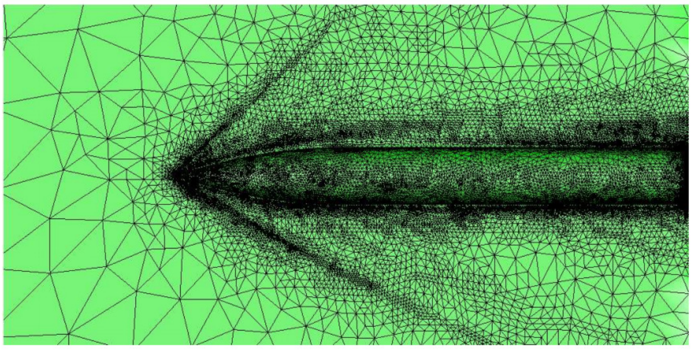

Figure 6: Adapted mesh obtained through Error Estimation and Adaptivity

However, such a 'brute-force approach' is not always of much help to the practising engineer who might be more interested in bounding the error in a local quantity of interest, e.g., a point stress or stress resultant across a section of the model. Indeed, even though the author studied *a posteriori* error estimation in some depth under John Robinson and Edward Maunder, [19], he has never actually applied it in practice preferring, instead, to use convergence techniques to ensure suitable accuracy in his models. The author's feelings on this matter are reinforced by the following paragraph presented in the conclusion to a 2004 paper, [20]:

"Hence, while the theory of error estimation has provided much valuable insight into the finite element solution process, many of the proposed techniques are at present only valuable to a limited extent in engineering practice."

Has much changed in the years following this statement? Well there has certainly been a great deal of academic work and the CFD community has realised the importance of reporting results together with an error estimate (something that the structural community might learn from here!). One trend that has been noted certainly in the structural engineering community is the development of error estimators based on better recovery of strong equilibrium from the weakly equilibrating results from conforming elements. This is a step towards the idea of dual analysis, which is discussed in section 5.

4. Convergence and Extrapolation

Error estimation, useful though it is in its current state of development, is not of much practical use to engineers because it does not provide the essential verification quantities of an accurate error bound. One technique that can help here, and which the author has found invaluable in his simulation career, is that of convergence and extrapolation. For a numerical simulation tool to be verifiable, then it must converge to the theoretical solution of the mathematical model with mesh refinement. Whether or not the mathematical model represents real physical evidence is, of course, a matter for validation.

Mathematical treatment of numerical simulation techniques will generally provide theoretical convergence rates for the various quantities involved in the simulation. The observed rate of convergence, in practical engineering problems, may differ from these theoretical rates of convergence but, in the limit, as the mesh is refined, these theoretical rates will be seen in practice. It should be noted that for very coarse meshes, where the convergence is not asymptotic (pre-asymptotic), the observed rate of convergence might not be anywhere near the theoretical rate.

Lewis Fry Richardson FRS was a polymath who worked in a wide range of scientific/engineering fields in the early 1900s – his biography is an interesting read that shows how much an inquiring mind can achieve over a (very) wide range of fields of scientific endeavour, [21]. Richardson's 1911 paper on accurate prediction of stresses in a masonry dam, [22], together with the later 1927 paper on the deferred approach to the limit, [23], provide engineers with a means of accelerating the convergence of their simulation by the technique that has become known as Richardson Extrapolation.

An example of Richardson Extrapolation is given in figure 7 for the area of regular inscribed polygons. As the number of sides increases, the area converges to the area of the circle to which the polygons inscribe. The number of sides starts with three – the triangle being the coarsest approximation that is possible. The number of sides is then progressively doubled to 6, 12, 24, 48 etc. The log-log plot of convergence in the figure uses the area of the circle to determine the relative error in each approximation. The convergence is well behaved and as the number of sides is increased, the rate of convergence tends to the theoretical value of two. Two forms of Richardson Extrapolation are used to show the idea of accelerated convergence. The first, (II) assumes the theoretical rate of convergence and uses two successive approximate values of the area to extrapolate. The second (III) does not assume the theoretical rate of convergence and therefore requires three successive approximations of the area. Both extrapolations converge at almost double (four) the theoretical rate of convergence for the unextrapolated solution. Thus, if we require an error of less than 1% in the area, this can be achieved with 12, 24 and 48 sides respectively for RE(II), RE(III) and unextrapolated solutions respectively. It is noted, with respect to the next section, that if a circumscribed regular polygon had been used then instead of converging from *below* the theoretical solution, the result would have converged from *above* the exact solution. Having both approximations available, of course, does allow one to place exact bounds on the error.

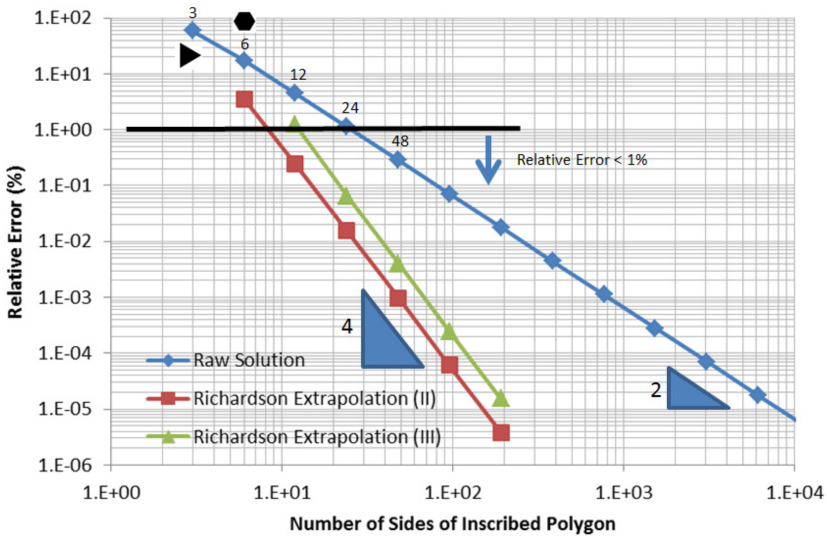

Figure 7: Convergence of the area of inscribed regular polygons

5. Dual Analysis

Whilst this may not be applicable to other fields of engineering, in structural analysis where there exists a sound theoretical basis for the existence of pure solutions that can bound the theoretical solution, the idea of dual analysis holds great prospects for providing precise bounds on quantities of interest.

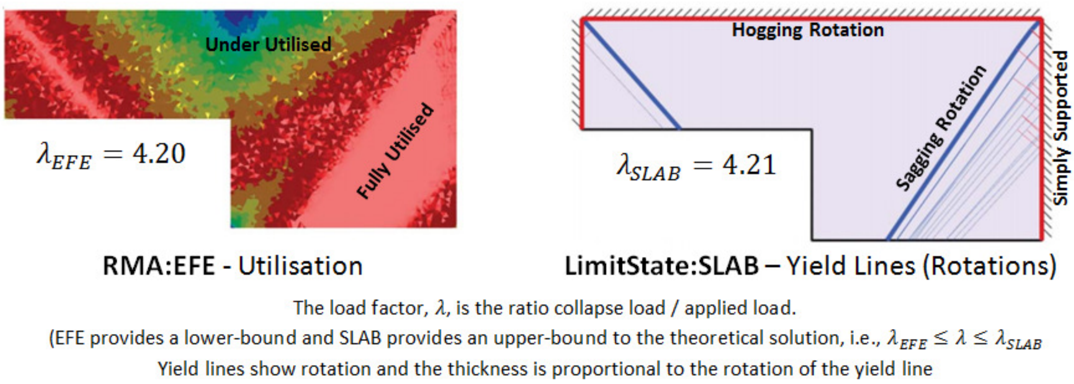

Figure 8: Limit analysis solutions for a landing slab

The results from just such a dual analysis using a pure equilibrium and a pure kinematic formulation to predict the collapse of a reinforced concrete landing slab are shown in figure 8, [24]. The collapse load is presented in terms of the load factor λ (this scales the applied load to give the collapse load). The lower and upper-bound solutions bound the theoretical solution extremely tightly, and provide a high degree of confidence for the practising engineer.

6. Closure

Good simulation governance is all about matching the results from numerical simulation with those of physical experiment and must take place with reference to a mathematical model. For the engineer working in the field of numerical simulation it is the process of verification that ensures the quality of the match between his/her simulation results and the mathematical model. Verification takes two forms in which the software first needs to be verified using problems with known exact solutions or with manufactured solutions. Once this form of verification has been demonstrated then the engineer may safely use the software to simulate his/her particular problem and conduct verification to ensure that the discretisation error is sufficiently small and the result is presented with sufficient engineering accuracy.

In this process of solution verification, the engineer is aided by error estimation and convergence techniques. There is much, too, that may be considered before the analysis has actually taken place and such *a priori* error estimation is a useful time for reflection before embarking on the analysis proper. Once the simulation has been run and the results are available, there are *a posteriori* error estimation techniques that may be employed. These techniques, together with an adaptive meshing scheme, can be used in an automated manner adaptively to drive the mesh towards an optimal solution that minimises the error. Such error estimators are, however, not yet sufficiently effective as to be able to place an accurate bound on the error for a given mesh. In such cases, the practising engineer can use convergence and extrapolation concepts, e.g., Richardson Extrapolation, to provide a more reliable estimate of the error.

In structural analysis, it is possible to perform dual analyses. Here the simulation is performed both with equilibrium finite elements (EFE) and with conforming finite elements (CFE). In this manner, it is possible to place exact bounds on certain simulation quantities of interest. The author predicts that once the engineering community becomes aware of the possibilities of dual analysis such techniques will become the norm where high quality verification is required.

One final point worthy of mention is that the vast majority of commercial finite element systems are based on libraries of low-degree displacement type elements. In practice, this means that for any practical engineering problem, significant mesh refinement is generally required to obtain reasonable accuracy. A nice example of this is using plane stress elements to capture the shear stress in a beam member. The shear stress varies parabolically across the thickness of the beam, being zero at the edges and a maximum at the centre. The numbers of elements through the thickness of the beam, required to capture the maximum shear to an accuracy of 1%, is 64 for the lower-order element and 8 for the higher-order element as indicated in figure 8. If elements with cubic displacement fields were available, then the exact solution would have been recovered with a single element. A few commercial finite element systems include elements of variable degree and these enable *p*-type refinement in addition to the conventional *h*-type refinement. The rate of convergence with *p*-type refinement is greater than with *h*-type refinement as was shown in the appendix 3 to the fourth NAFEMS Benchmark Challenge, [25].

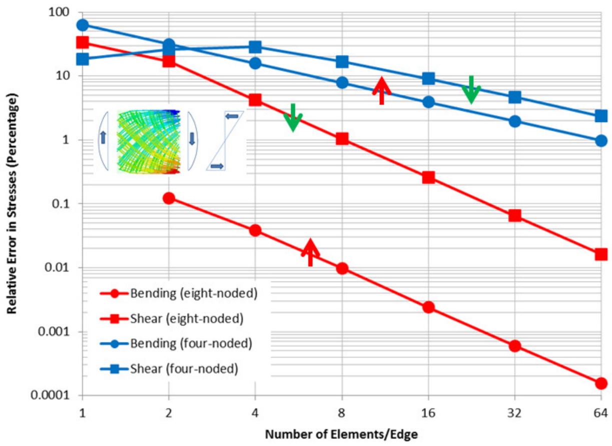

Figure 9: Plate-membrane problem with parabolic shear

The *p*-type element has been used to great advantage in the finite element system ESRD:StressCheck, [26]. This software provides the engineer with the means to conduct solution verification in an extremely straightforward manner by simply increasing the degree of the element, monitoring convergence and using Richardson extrapolation reliably to estimate the error. This can be conducted automatically by the software thereby enabling the engineer to concentrate on the engineering rather than the simulation. StressCheck has also been used to develop ESRD's Handbook and Toolbox applications. The first of these provides engineers with a repository of parameterised standard problems of the type found in texts like Roark's "Formulas for Stress and Strain", [27]. The second, Toolbox, is a tool that can be used to parameterise a company's range of components for rapid and reliable analysis by non-expert analysts. Toolbox then is an exemplar of the way in which the democratisation of simulation can be applied.

7. Acknowledgements

The author would like to acknowledge the input to this article, via discussions, of John Robinson and Edward Maunder.

8. References

[1] Ian Symington (Editor), *'NWC2015: Congress Review'*, NAFEMS, 2016.

[2] Peter Bartholomew, *'NAFEMS: the early days'* NAFEMS Benchmark Magazine, January 2016.

[3] Bernt Jakobsen, *'The Sleipner Accident and its Causes'*, Engineering Failure Analysis, Vol. 1, No. 3, pp 193-199, 1994.

[4] Henry Petroski, *'Design Paradigms: Case Histories of Error and Judgement in Engineering'*, Cambridge University Press, 1994.

[5] Democratization (NAFEMS) https://www.nafems.org/about/regional/americas/events/2020vision/democratization/

[6] Angus Ramsay, *'A Finite Element Approach to the Linear-Elastic Behaviour of Axisymmertically Loaded Corrugated Bellows Expansion Joints'*, MEng Dissertation, Liverpool Polytechnic, 1988. http://www.ramsay-maunder.co.uk/downloads/MENG_THESIS.pdf

[7] The Flixborough Disaster (1974) https://en.wikipedia.org/wiki/Flixborough_disaster

[8] Barna Szabo, *'A Case for Simulation Governance'*, Desktop Engineering, February 2015. http://www.deskeng.com/de/case-simulation-governance/

[9] William L Oberkampf & Christopher J Roy, *'Verification and Validation in Scientific Computing'*, Virginia Polytechnic Institute and State University, 2010.

[10] Angus Ramsay & Edward Maunder, *'The Influence and Modelling of Warping Restraint on Beams'*, NAFEMS Benchmark Magazine, July 2014.
http://www.ramsay-maunder.co.uk/downloads/warping_article_web.pdf

[11] Exemplar of Published Verification: http://www.limitstate.com/slab/validation

[12] John Robinson, *'CRE Method of Element Testing and the Jacobian Shape Parameters'*, Engineering Computations, Vol. 4, Issue 2, pp 113-118. 1987.

[13] Kambiz Salari & Patrick Knupp, *'Code Verification by the Method of Manufactured Solutions'*, SAND2000-1444, Sandia National Laboratories, 2000.
http://prod.sandia.gov/techlib/access-control.cgi/2000/001444.pdf

[14] John Robinson, *'Basic and Shape Sensitivity Tests for Membrane and Plate Bending Finite Elements'*, NAFEMS Publication, 1985.

[15] D. Gaston, G. Carey, R. Strogner, J. Peterson, B. Kirk and L. Branets, *'On Combining Mesh Redistribution with h-Adaptivity'*, USNCCM, July 2007.

[16] O C Zeinkiewicz & J Z Zhu, *'The Superconvergent Patch Recovery (SPR) and Adaptive Finite Element Refinement'*, Computer Methods in Applied Mechanics and Engineering, Vol. 101, Issues 1-3, pp207-224, 1992.
http://www.sciencedirect.com/science/article/pii/004578259290023D

[17] Angus Ramsay & H Sbresny, *'Evaluation of some Error Estimators for the Four-Noded Lagrangian Quadrilateral Error Estimator based on a Patch Recovery Scheme'*, Communications in Numerical Methods in Engineering, Vol. 11, Issue 6, pp497-506, 1995.
http://www.ramsay-maunder.co.uk/downloads/evaluation_of_some_error_estimators.pdf

[18] C.L. Bottasso , G. Maisano , S. Micheletti , S. Perotto, *'New recovery based a posteriori error estimators'*, COFIN Project, 2003.
https://www.mate.polimi.it/biblioteca/add/qmox/mox52.pdf

[19] Angus Ramsay, *'Finite Element Shape Sensitivity and Error Measures'*, PhD Thesis, University of Exeter and Robinson FEM Institute, 1994.
http://www.ramsay-maunder.co.uk/downloads/Exeter_Thesis.pdf

[20] Thomas Grätsch & Klaus-Jürgen Bathe, *'A Posteriori Error Estimation Techniques in Practical Finite Element Analysis'*, Computers and Structures, 83, 235-265, 2005.

[21] http://www-groups.dcs.st-and.ac.uk/history/Biographies/Richardson.html

[22] Lewis F Richardson, *'The approximate arithmetical solution by finite differences of physical problems including differential equations, with an application to the stresses in a masonry dam'*, *Philosophical Transactions of the Royal Society A* Vol. **210**, Issue: 459-470, pp307–357, January 1911.
http://rsta.royalsocietypublishing.org/content/210/459-470/307

[23] Lewis F Richardson & J Arthur Gaunt, *'The Deferred Approach to the Limit'*, Philosophical Transactions of the Royal Society A, Vol. 226, Issue: 636-646, January 1927.
http://rsta.royalsocietypublishing.org/content/226/636-646/299

[24] Angus Ramsay, Edward Maunder & Matthew Gilbert, *'Yield Line Analysis of Reinforced Concrete Slabs: is the 10% Rule Safe?'*, NAFEMS Benchmark Magazine, January 2015.
http://www.ramsay-maunder.co.uk/downloads/Yield-Line_Article_for_Benchmark_Magazine.pdf

[25] Angus Ramsay, *'Poisson's Ratio for a Steel Plate'*, NAFEMS Benchmark Challenge Number 4,
http://www.ramsay-maunder.co.uk/downloads/Benchmark%20Challenge%20004.pdf

[26] Engineering Software Research and Development (ESRD), http://www.esrd.com/

[27] Warren C Young, Richard G Budynas & Ali M Sadegh, 'Roark's Formulas for Stress and Strain', 8th Edition, McGraw-Hill, 2012.