Title	Aviva Stadium
Subtitle	Use of Parametric Modelling in Structural Design
Author	Paul Shepherd MA, PhD, CSci, CMath, FIMA
Affiliation	University of Bath
Postal Address	Department of Architecture & Civil Engineering
	University of Bath
	Bath BA2 7AY
	UK
EMail	P.Shepherd@bath.ac.uk
Tel	01225 384406
Fax	01225 386691

Contents

Synopsis

The new 50,000 seater Aviva Stadium being build at Lansdowne Road and due to be completed early in 2010 will provide a world class sporting facility for the city of Dublin, Ireland. Due to the nature of the site and its surroundings a responsive architectural form was required which could be manipulated to optimise the design for a wide variety of design variables and constraints by accommodating these design changes quickly and easily. This naturally led project architects, Populous, to adopt a parametric approach to generate the form of the building envelope, whereby CAD models were defined using rules and relationships between objects rather than absolute coordinates in space. This architectural parametric model was then passed directly to the structural engineers, Buro Happold, who extended it to include the geometry of the structural members.

This paper adds to the body of scientific knowledge by documenting the details of the parametric model created by the author for the structural engineers, including not only the generation of geometry, which incorporates design constraints and embraces structural engineering logic, but is also linked to structural analysis software to provide rapid feedback on the effects of design decisions on structural performance.

Furthermore, since this is the first example of a project being designed and built using a single parametric model from architectural design right through to structural analysis, the advantages and challenges of adopting such an approach are discussed. Suggestions are made, based on the lessons learned, to assist those using parametric methods on live projects in the future.

1   Introduction

Ever since the introduction of the desktop computer into the workplace, parametric design methods have been used to help architects and structural engineers develop models of buildings with complex geometries. With the recent introduction of such methods into mainstream architectural CAD packages, it is clear that these methods will become more popular in the future.

The subject of this paper is the Lansdowne Road Stadium project, which, in the view of the author, is the first project to use a single parametric model jointly developed and used by both architect and engineer, to encapsulate and communicate design intent throughout the whole design process.

1.1   Project Background

The original Lansdowne Road stadium in Dublin, Ireland was the oldest international rugby ground in the world, hosting international matches since 1878. However, by the twenty-first century its lack of hospitality suites and the presence of standing-only viewing areas meant that it no longer met the expectations of a world-class sporting venue. In 2003 a feasibility study concluded that the best solution would be to build a new stadium on the existing site. Demolition of the old stadium began in May 2007 and the new home of Ireland's international soccer and rugby teams is currently under construction, due to be completed in early 2010.

The site is in a residential area of Dublin, sandwiched between housing to the north, a river to the east, a railway-line to the west and a road to the south. Additional issues such as planning height restrictions, rights-to-light and an underground sewer further complicate the design. To preclude overshadowing to the northern end of the stadium the envelope needed to be much lower at the north than the south, whilst the requirement for 50,000 seat capacity pushed the height of the seating bowl up to the east and west. This led to the final form shown in Fig. 1.

1.2   Parametric Approach

Historically, building designs were conceived and communicated through ink drawings on paper. Minor changes to a design could be scratched out and re-drawn but major changes required a whole new drawing.

Nowadays, desktop computers and printers are common-place in the design office. CAD models have directly replaced paper drawings as a means to communicate design intent, since changes can be performed relatively quickly by deleting and redrawing lines on the screen and a new copy of the updated drawing printed out or emailed. Such CAD models are created by the user specifying the position of lines in a drawing by defining the coordinates of their start- and end-points. A graphics tablet or computer mouse can be combined with tools such as snapping to speed up the definition of these coordinates. But internally the computer is simply recording the absolute position of each line, arc or curve. So, as shown in Fig 2, if lines representing secondary beams are drawn to span between two primary beams, and then the primary beams are moved further apart, each of the secondary beams must be manually redrawn (or extended) to maintain their connectivity with the primary beams.

In contrast, Parametric CAD models store the underlying relationships between objects, rather than their absolute coordinates. Visually they appear just the same, but in the above example a parametric secondary beam would be defined by the relative positions of its ends along the primary beams. As Fig 3 shows, a secondary beam which is defined as spanning from 1m along one primary beam to 1m along another does not need any further editing if the primary beams are moved apart. So long as the rule defining its position does not change, the CAD system will automatically update the geometry of an object in response to any change.

Parametric modelling techniques were first introduced into the produce-design and aerospace industries, which didn't have such a long history of using paper to communicate design and were therefore quicker to reap the full benefits of computer aided design. In 1994 the Boeing 777 became the first commercial aircraft to be designed entirely on the computer using parametric modelling software. It is especially the parallel nature of the aerospace design process which means that, for example, if aerodynamic test results require a plane's wing positions to move, its window positions, electrical wiring paths, etc can be automatically updated simple by re-applying their embedded design rules.

At first the building construction industry was more conservative and resisted such changes to the established method of using the computer as nothing more than electronic paper. Whilst some of the more geometrically challenging building projects such as the completion of Gaudi's Sagrada Familia and the work of Frank Gehry allowed those few designers with programming skills to develop bespoke parametric software in order to achieve a working design, the extra effort required up-front to define all the design parameters and relationships meant that parametric methods were not used on the vast majority of projects. Since buildings are generally one-off designs which do not need to be changed over time, and often any up-front effort is wasted when a design is not taken to construction, the take-up of parametric modelling in the construction industry was limited.

However, as computers and software have evolved, the up-front effort required has become less onerous, especially as a new generation of designers become more familiar with the concept. For decades now, Frank Gehry's office has used the same software (CATIA) that was used to design the Boeing 777. And having seen first hand how the benefits of using this technology could be easily transferred into building design, its sister company Gehry Technologies have developed a version of CATIA especially tailored for the construction industry (Digital Project) and more established CAD vendors such as AutoDesk and Bentley each now have parametric modelling capabilities built-in to their products.

For the Aviva Stadium project, due to the tight site constraints and the expected, sometimes conflicting, requirements of any large-scale multi-disciplinary project, it was identified early on that the design would go through many iterations before converging on a solution. A parametric approach would therefore allow the design to be adapted and refined without the usual impact on engineering analysis or drawing production. Furthermore, some of the site constraints could be introduced directly as constraints on the parametric model, therefore embedding compliance into any generated design solution.

1.3   Architects' Model

The starting point for the development of the structural model of the Lansdowne Stadium roof was an existing parametric model developed by the architects to define the doubly-curved skin geometry.

At an initial design meeting, the building envelope was separated at the point where cladding meets structure. It was agreed that the architects would take responsibility for defining everything outside this point, including the cladding support system and cladding panels, whereas the structural team would be responsible for defining the structural support members inside this dividing line, from purlins and bracing members to main structural trusses.

The architects chose Bentley's Generative Components (GC) as the software platform for their parametric model, and GC was therefore also adopted by the structures team as their platform. This sharing of a common platform allowed continuity in the information workflow and any changes by the architects could be easily communicated to, and incorporated by, the structural team.

Detailed discussion of the architectural parametric model is outside the scope of this paper and more information can be found in papers by Hudson[1,2,3].

1.4   Engineers' Model

This paper discusses the parametric modelling undertaken by the structural engineers to develop the design of supporting structural members for the stadium roof. In this context, the term "roof" refers to the external skin of the building, both where it is horizontal in nature in the traditional sense of the term "roof" and where the skin becomes more vertical and acts as what would traditionally be called a "wall". The purpose of the roof is to provide rain protection to the spectators below and therefore constructed from polycarbonate cladding panels. In contrast, the wall area is constructed from louvers, providing a semi-permeable rain-screen whilst at the same time allowing air to enter the stadium for cooling and ventilation. Since there is a smooth form to the geometry, there is a transition area where the solid cladding panels of the roof curve round to become the top of the louvered wall.

The engineering solution for the structural members required to support this "roof" is a hierarchy of tubular steel truss members. The main feature is a large primary truss in the shape of an inclined horseshoe, the ends of the horseshoe being at the much lower north-end of the stadium allowing them to be supported directly by thrusting into the ground. Around the rest of the horseshoe, support is provided by a series of radial secondary steel trusses which take the load from the horseshoe and transfer it out to the edge of the roof and down columns. A series of smaller, tertiary steel trusses provide lateral support to the primary truss as well as supporting the roof panels, and these span between the primary truss and an outer edge ring truss which itself spans between columns. This arrangement is shown diagrammatically in Fig. 4.

The architectural parametric model defined the surface representing the interface between engineer- and architect-responsibilities by creating a surface which passes through a series of radial curves (lofting). These curves are made from arcs and lines which join together smoothly without a crease (tangent-continuous). It was therefore decided to use this same radial grid as the basis for defining the structural geometry. This would result in planar

tertiary trusses with top-chords made from faceted arcs and straight lines. Any other grid layout would not result in arc geometries and may therefore be more complicated to fabricate.

Another key consideration in choosing to create a parametric structural model was the definition of structural loading. Since the geometry of each radial truss was different, so too was its effective width at any particular point along its length and therefore the corresponding wind- and cladding-loads were also different at each position. By incorporating the calculation of load magnitude and direction into the parametric model, and by creating a direct link from this extended information into the structural analysis software, vast amounts of time were saved in preparing the geometric model for structural analysis.

The details of the parametric geometry model are outlined in the next section, with the integration of structural analysis described in section 3. The final section contains a discussion of how this work was carried out in the context of a live project and suggests some lessons learned which should be considered in future projects of this kind.

2 Generation of Structural Members

The starting point for the generation of the structural geometry was the parametric model generated by the architects. Since the actual construction logic of this model was not required, many of the underlying geometrical entities were hidden graphically by the architects and not referenced by the engineers. The main part of the architectural model required is the series of radial curves from which the architectural inner-cladding surface is constructed by lofting. These curves are comprised of two arcs and a straight line. The first arc represents the outer near-vertical wall part of the skin surface from floor level upwards, the second arc represents the fillet corner where vertical wall transforms into more horizontal roof and the straight line represents the more central part of the roof surface reaching towards the pitch. These curves are shown in Fig. 5 and the structural geometry is defined relative to them.

2.1 Primary Structure

The basic on-plan geometry of the primary truss is formed from a simple faceted horseshoe. This horseshoe was generated by parametrically defining points on a subset of the radial grid-lines (known as key grids) contained in the architect's parametric model. These key-grids are shown in Fig. 4 as thicker radial lines every two or three bays, this subset being defined parametrically. Points were defined on the floor-plane by specifying their distance along the key gridlines. To define the height of these points for the primary truss top-chord they were projected vertically to intersect with the architectural surface and then offset vertically downwards by a chord tube radius plus a tolerance offset, so that they are guaranteed not to clash with the cladding.

The primary truss gains its structural stiffness from its depth and there is therefore a benefit in having it as deep as possible (within reason). However, another important design consideration was that the primary truss should not interrupt the sightlines of the spectators and should therefore not be too deep. The depth of the truss at each key gridline was therefore not specified explicitly through parameters from a spreadsheet, but calculated algorithmically to be as

deep as possible without interfering with the view. The position most likely to experience sightline interference from the truss is the spectator furthest from the front, who is not only far away from the action, but also seated highest. The position of the eyeball of such a spectator was calculated for each gridline by using a constant offset from the back-of-bowl position. This back-of-bowl position was inherited by the engineers from the architect's parametric model, and the offset was defined by the distance horizontally and vertically that a person's eyeball would be from the back-of-bowl, assuming a standard seating arrangement. The object this spectator is most likely to have problems seeing is a ball when kicked high into the air, and the design requirements defined a point 18m above the centre-spot of the pitch as a general position which should be viewable. The calculated eyeball positions were therefore joined by a straight line to the point 18m above the centre-spot and lofted together to produce a sightline surface, a no-go-zone, below which no structure would be allowed lest it interfere with the spectators' view. This is shown diagrammatically in Fig. 5.

To define the bottom chord positions, the floor-plane key gridline points were therefore projected vertically to intersect with this sightline surface, and then offset upwards by a chord tube radius plus a tolerance, to guarantee that they would not interfere with the view.

These key-grid points were then joined together with straight lines to define the top and bottom chord geometries. These straight top- and bottom-chord segments were then split by planes formed from the non-key gridlines extruded vertically to define the top and bottom positions of the vertical and diagonal truss lacing members.

2.2   Secondary Structure

The primary truss can thrust straight into the ground at the north end of the building where the roof is much lower. However around the rest of the truss, its self-weight and the weight of the roof it supports must somehow be supported. Since any columns would interfere with the seating below, both physically and in terms of sightlines, radial secondary trusses were introduced. These secondary trusses span from the primary truss to column positions around the outside of the bowl behind the spectators. The secondary trusses lie in a vertical plane and are triangular in elevation, with top- and bottom-chords defined as straight lines from the top- or bottom-chords of the primary truss to a single column position. Since a lateral force is introduced at each point where the primary truss is kinked, the secondary trusses were chosen to meet the primary at exactly these kink-points, in effect "peeling-off" some of the force as the horseshoe changes direction.

Through the parametric nature of the modelling process, a simple parametric study of the secondary trusses was possible. A set of schematic structural models were quickly generated containing primary and secondary trusses only. The secondary trusses met the primary at different angles, in effect sharing out the force between primary and secondary trusses in differing proportions. From simple finite element analyses of the different layouts under simplified loading conditions, it was decided that the secondary trusses should leave the primary truss at a key-gridline (kink point) and land on a column one bay round from the other end of the same key-gridline (see Fig. 4). This therefore defined the

column positions, and the top- and bottom-chords of the secondary trusses could be created accordingly.

A number of different systems of secondary-truss lacing were explored, such as dividing the chords up into a number of equal divisions, or a number of equal-length bays. Different schemes of lacing such as N-, K- or W-shaped bracing were also investigated. Practical considerations of the angles that diagonal members made with the vertical- and chord-members were also studied. In each case, the parametric nature of the model helped to speed up this process, allowing many different variations to be tested against structural and constructional performance. The final geometry spaced the bays unevenly such that the purlins would land exactly on a vertical truss member.

2.3   Tertiary Structure

Since the large primary truss was designed to transfer the load of the roof into the two abutments at the north, and the secondary trusses to transfer the load back onto the supporting columns, tertiary trusses were required at every radial grid-line only to transfer the load of the roof onto the primary truss.

The shape of the top chords of the tertiary trusses was defined by a constant normal-offset from the roof geometry. The size of this offset was defined to be equal to the top-chord radius plus sufficient room to accommodate the purlins and cladding support structure.

The bottom chord geometry was less constrained and more open for experimentation. They would be highly visible to the spectators and so the architects had requested that on aesthetic grounds they be made from two arcs meeting at a tangent. From structural considerations, the point of maximum depth should be located near to the point of maximum bending moment. The tertiary trusses are supported at their outer end and where they pass through the primary truss. Since the outer end meets a triangular ring-truss (see section 2.5) they could be said to be heavily rotationally restrained at this point. The truss was therefore idealised as a propped cantilever under uniformly distributed load. This allowed a relatively simple equation to be used (the derivation of which is beyond the scope of this paper) to calculate the position of maximum moment. Fig. 6 shows this equation, the idealised cantilever and bending moment, along with the resulting shape of the tertiary trusses. The equation was embedded within the parametric model, so that each individual tertiary truss had its bottom chord made from two tangential arcs whose maximum depth was exactly in the position required for the structural efficiency of that truss. The size of this maximum depth was also calculated individually for each truss, set to be equal to a quarter of its span. A maximum limit of 4.4m was also placed on this depth to ensure that the trusses could be fabricated off-site and would fit inside a standard lorry for transportation to site.

The lacing of the tertiary trusses was also the subject of detailed investigation. The parametric model allowed many different arrangements to be compared and the exact spacing of vertical lacing members to be optimised (see Fig. 7). The final design spaced the vertical members 8m apart and used diagonal bracing. Where the 8m spacing rule would introduce a member close to the ends of the truss (within 4m), these members were automatically removed, making a slightly larger bay but removing unnecessary material.

2.4   North Stand

At the time the design was being performed, GC could only access a maximum of around 1GB of memory, irrespective of the memory available. Unfortunately the parametric model was so complex that it required more than this and could not be stored in memory all in one go. The model was therefore constructed in two parts, with a natural split where the north part of the stadium roof joins the primary truss. The majority of the stadium (solid lines in Fig. 4) was generated, stored, and then deleted to free up memory, to allow the rest of the structure, the North Stand (broken lines in Fig. 4), to be modelled.

The roof of the north stand is much lower than the rest of the roof structure to ensure the right-to-light is maintained to neighbouring properties to the north of the site. There is also no primary truss to span onto in this area therefore a different structural system was used. The concept of planar tertiary trusses lying on each grid-line was maintained, but these trusses were extended down to the ground behind the spectators (see Fig 8). Since the north stand structure is lower and the extent of the roof towards the pitch smaller, this was shown to be sufficient with the addition of significant bracing members (shown as broken red lines in Fig 8) to provide out-of-plane stability. Again, the top chord of the truss was defined by an offset from the cladding surface and the depth of each truss was specified as a proportion of its span.

2.5   Other Structure

In addition to these major structural elements, more detailed structural members were also included in the model so that absolutely every member required for structural analysis would be present in the model when eventually exported.

A triangular truss was run around the outer edge of the entire roof, encompassing both north and south geometry. This provided out of plane stability as well as rotational restraint, and was also chosen as the transition point between vertical and horizontal cladding systems. Single column elements were also added to transfer the load down from the roof above the spectators' heads to the concrete column position at floor level. Purlins and diagonal bracing members were also included where necessary.

3   Integration with Structural Analysis Software

A key influence on the engineers' decision to adopt a parametric approach to the modelling of the structure was the ability to harness the additional intelligence embedded within the model to automatically generate a structural analysis model. Traditionally, any CAD representation of a structure is transferred to structural analysis through CAD file formats, which deal only with geometrical information and not structural sections or analysis parameters. The additional information required to perform a structural analysis, such as section size and loading, is then added to the analysis model manually. Although techniques such as the careful use of CAD layers can be used to make this process quicker, it still represents a significant time commitment and also leaves the process open to human error. For the Lansdowne Road stadium design it was therefore decided to extend the parametric approach to include generation of a structural analysis model. Apart from the obvious time saving from automatically generating the geometry, further significant savings were

made by also calculating and applying the many complicated load patterns to each member.

3.1 Software Framework

The parametric modelling was performed using Generative Components (GC) for reasons discussed above. GC has ways of interacting with other programs, especially Excel, however it also has a programming interface with Microsoft's C#.NET which makes it extremely flexible and adaptable.

The software chosen by the engineers for their structural analysis was Robot Millennium (Robot). A Robot model can be defined most easily through its open text format input file, which allows every aspect of a Robot model to be defined and the resulting file, once opened in Robot, is ready for an analysis to be performed without any further user interaction. Robot also has a COM interface which allows the program to be driven directly by another application, and a C# programming "wrapper" is available to facilitate this. However, the main obstacle to using this route to populate, and even control, a structural analysis was memory and in particular the way that the parametric model was generated in two parts to be merged together. This made populating Robot with a single model based on two separate GC structures much more complicated. There was also concern over whether any resulting single analysis model would then be able to function with the remaining memory that GC had not used.

It was therefore decided to use the text input file as the bridge between parametric model and analysis. This had the additional benefit of making the programming more transparent so it was simpler to debug and a non-expert user could more easily follow what was happening if need be.

A new GC object was created based on the GC "parameter" class, which took in the relevant geometry objects from the parametric model, and returned a simple true/false indication of success. This was then turned into a component in GC, which automatically created the equivalent C# code. This C# code was then manually edited to create a new text file and populate it with the relevant information (geometry, sections, materials, supports, etc.) required by Robot to generate its analysis model. This meant that data was automatically extracted from the parametric model by the software and used to create a structural analysis model. This removed the time-consuming step of manually creating an analysis model and eliminated the possibility of introducing errors in translation.

3.2 Geometry

Using a sensible system of numbering to keep track of the different parts of the model, the entire geometry of the stadium roof structure was replicated from GC to Robot. This involved the creation of nodes and bar elements to represent all the steel members including Primary, Secondary and Tertiary trusses, purlins, columns and bracing members. The cross-section profile of each element was defined as a steel tube of specific size and the local axis of each element was orientated to be normal to the cladding surface.

Since this was all performed by the software program and took the GC entities as input, the parametric nature of the process was maintained. If the design of the stadium was altered at any stage, either by architect or engineer, an

updated analysis model would automatically be generated reflecting these changes as an output of the GC component.

3.3 Loading

Due to the complex nature of the geometry, an accurate representation of the loading on the structure would require every different finite element to be allocated a unique triangular loading pattern for each loadcase. With nearly 3000 bar elements and 23 base load cases this would be prohibitively time consuming to define manually even just once, let alone many times in response to design changes, and conservative assumptions would have had to have been made.

However, since the parametric model contains all the geometric information required to calculate the effective area of each bar element, the C# component was extended to generate the loading information automatically. For each bar, its effective width was calculated, based on half the distance to the equivalent bar one bay around in either direction. Since on plan each bay was triangular in shape, the effective width of the start and the end of each bar was generally different. Therefore the load pattern applied to each bar was a trapezoidal distributed load, representing the change in effective width along the bar, as shown in Fig 9. The distributed load was applied either normal to the cladding surface to represent wind-induced loading, or vertically to represent dead, live and snow-induced load, in which case it was calculated based on vertical projections of the effective width.

The roof was divided into the fourteen zones shown in Fig. 10, and a different load factor, derived from the results of wind tunnel tests, was applied in each zone to build each base load case representing the worst cases of wind from the north, south, east, west, etc. as well as live loads and snow. These base load cases were then used in different ways by the 74 load combinations to investigate the envelope of forces present in the structure. Clearly, the dependence of all these load cases on the geometry of the structural skin would have made such an accurate representation of the loads by manual methods impractical. It is only through integrating the analysis model generation into the parametric modelling software that such detail could be included and such an efficient structural design could result.

3.4 Other Details

Since the aim was to automatically generate a finite element analysis model ready for calculation without manual intervention, every detail of the model had to be incorporated. This meant that all the boundary condition supports were included; in particular the sliding bearings on the ends of the primary horseshoe were automatically rotated to align with the truss. Pin-type released were also included on the ends of all the bracing members as well as sliding-type releases to represent the thermal movement joints.

Since memory limitations meant that the model was exported to Robot through two separate data files, considerable care was taken to make sure the interface between the two models was consistent. This meant that Robot could then automatically combine the second file with the first, merging coincident nodes, and matching the corresponding load cases. By simply running the GC component to export the data and importing the two files, the engineers had a

fully functioning analysis model ready for calculation. A number of spreadsheets were generated to facilitate structural section design directly from the Robot output and a traditional CAD model was constructed which would directly import the 3D analysis model and cut 2D sections to produce the drawings for issue to the steelwork contractor. The steel contractors were therefore using traditional 2D drawings familiar to them, rather than the parametric model. However there is no reason why the fabricators could not have used the parametric model directly, and benefitted from its added intelligence, had they been familiar with such technology.

4   Conclusions and Lessons Learned

4.1   Workflow

Sharing a single parametric model between architecture and engineering teams was of great benefit in terms of allowing changes to permeate down through the design teams. The model itself was hierarchical in nature, and followed the design process in that it was initiated by the architects and passed on to the engineers for further development. By conforming to the flow of design data in this way, design changes initiated by the architects were communicated quickly and effectively to the engineers, who were able to assess the effects of the changes by performing updated structural analysis and were able to produce updated shop-drawings with minimal manual intervention. This allowed the engineers to respond quickly to changes very late in the design process. It is arguable as to whether this is always a benefit. Obviously an ability to efficiently change the design can help when problems are discovered late-on in the design process. But this capability can be abused, and design decisions deferred when the design team as a whole would benefit from more certainty earlier on.

By agreeing a clear, physical division between each team's areas of responsibility, the engineers were able to concentrate on the design of the structural members, within the boundaries and parameters set by the architects, with confidence that their design decisions would be compatible with those of the architects. However, this hierarchical nature made the design very much a one-way process. Any time the engineering analysis suggested a change was required to the overall geometry of the structure, this had to be communicated to the architects using traditional methods and the architects would then have to update their top-level parametric model accordingly. Whilst imposing an asymmetry to the architect-engineer conversation, which seems likely to stifle innovation, the author believes this was a sensible strategy, since the architects were able to act as overall model managers and were able to think carefully about the consequences of any design change suggestions before implementing them in their parametric model.

The main issue to come out of a critical assessment of the project is the over-reliance on a limited number of skilled personnel. At the time of the design, parametric modelling on live projects was extremely rare, and as such, access to sufficiently expert staff was limited. Both the architectural and engineering teams allocated the parametric modelling task to individual specialists, which meant that the whole design had to pass through these people before appearing in more usual forms such as 3D CAD or Finite Element Analysis

models. Similarly, any suggested changes to the design had to be implemented by these individuals.

Historically, the same problems were encountered during the infancy of CAD, where the majority of a design office would be trained in the use of drawing-board and pen, and specialist CAD modellers were creating electronic versions of pen and paper. Their specialist knowledge placed undue reliance on them and if they were suddenly ill, the rest of the design team would have trouble to print off or make changes to drawings. As CAD became a more popular tool in the design office, the number of CAD-trained personnel increased, so that this bottle-neck was no longer an issue. However, this only functioned due to the parallel development of CAD-standards, an agreed method of representing drawings in CAD, which enabled one CAD user to pick up a colleague's model and immediately understand it and be able to make changes without a personal hand-over.

Currently, with parametric modelling systems fast becoming ubiquitous in, at least for now, the larger design offices, the issue of limited numbers of trained operators is likely to vanish similarly to the way it did with CAD technicians. But without any sign of a parametric equivalent to CAD standards on the horizon, it is difficult to see how the risk of undue reliance on a single individual can be mitigated. Since the very nature of a parametric approach makes commonality between models for different projects limited, the scope for standardising the models themselves is also limited. However, it is proposed that a standard method of documenting the construction of a parametric model is desirable, possibly even imposed by the software, such that models are easier to transfer between users without a steep learning curve. This would also mean that they could more easily be developed in a team environment, and would be more open to checking, validating and quality assessment. The exact nature such a "CAD++ standard" should take is open for discussion and is the subject of future research by the author.

There is no reason why the benefits of taking a parametric approach to building design cannot be equally well applied to smaller, more day-to-day building projects. The benefits seen on the Aviva Stadium in terms of fast reaction to design change and the embedding of design constraints directly into the model to ensure compliance are just as applicable to the design of a small office building or domestic extension as they are to the complex geometry associated with a stadium roof. It is simply a case of weighing up such benefits against the extra overhead involved in setting up the parametric model in the first place. As parametric software becomes commonplace in the design office, and as more engineers become familiar with it as a tool, this associated overhead will reduce and parametric modelling will become the norm as it has become in other engineering disciplines.

4.2 Design Optimisation

A clear benefit of taking a parametric approach to the design of the Lansdowne Road stadium was in the scope for structural design optimisation. Tasks such as altering the depth profile of a truss or generating various different truss-lacing arrangements could be performed in a matter of seconds or minutes, rather than the hours or days it would take using non-parametric methods. This allowed detailed optimisation of the structure to be performed quickly and

efficiently, making real cost savings. Furthermore, by building-in real practical constraints into the model, such as maintaining sightlines and ensuring sufficient drainage, the designers were freed up to explore the variable aspects of the geometry without fear of inadvertently violating more rigid design constraints.

By far the most innovative aspect of the parametric model was its link to structural analysis software. Significant resource was invested in developing this link, much more than was anticipated at the outset, but the time invested was certainly recouped in terms of having instant creation of analysis models. The main benefits were in terms of automatically generating the load-cases for analysis. There were other methods, such as transfer by DXF file, which would have allowed the transfer of geometry between parametric model and analysis package. With creative use of colours, layering, and combinations of multiple files, workarounds could have been developed to transfer section properties as well. However, the calculation and application of multiple geometry-dependant load-cases was only possible through scripting within the parametric modelling environment. And the transfer of load information into the structural analysis package required use of native data formats not traditionally supported in CAD.

More recent moves by the CAD software industry mean that geometric and parametric modelling software is becoming more closely integrated with engineering analysis. However, the importance of including the less-physical aspects such as structural loading in their software should be recognised. When investing in supposedly integrated modelling and analysis applications, engineers should not underestimate the advantages an effective approach to structural loading can bring.

4.3 Acknowledgements

5   References

1   Shepherd, P. & Hudson, R. (2007), "Parametric Definition of Lansdowne Road Stadium", in: International Association of Shell and Spatial Structures, Venice, Italy: CD-ROM.

2   Hudson, R., 2007, "Frameworks for practical parametric design in architecture", in Pottman, H., Hofer, M. & Kilian, A. (eds), "Advances in architectural geometry", pp17–20.

3   Hudson, R., 2008, "Knowledge Acquisition in Parametric Model Development", International Journal of Architectural Computing, 6(4), pp435-451.
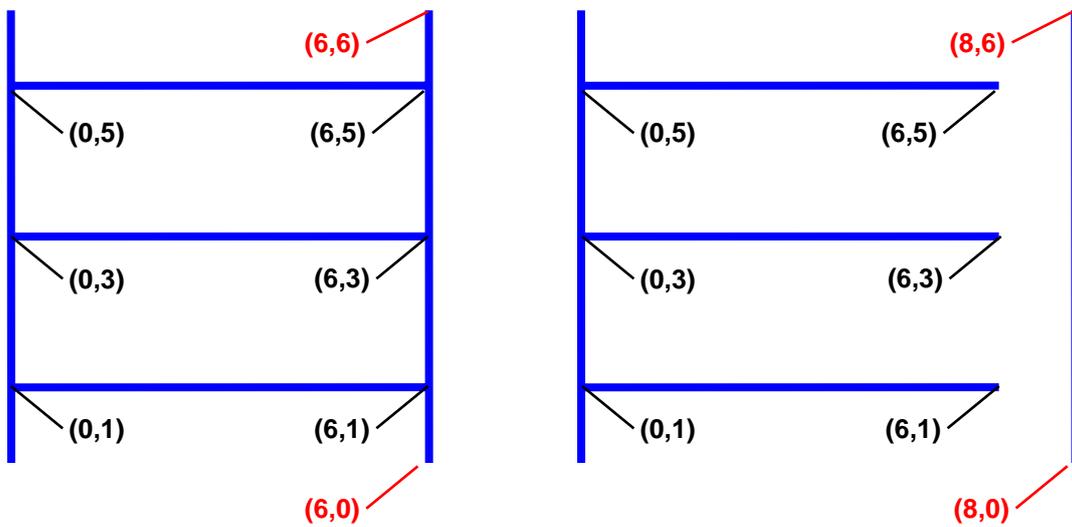
Fig. 1 New stadium near completion

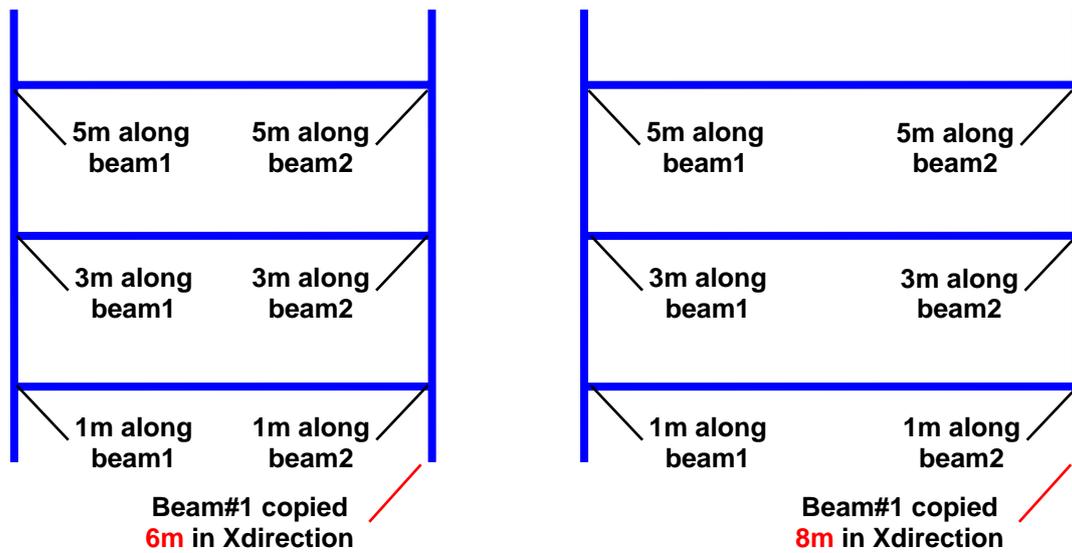Fig 2 Coordinate definition of beam locations before (left) and after (right) a change.

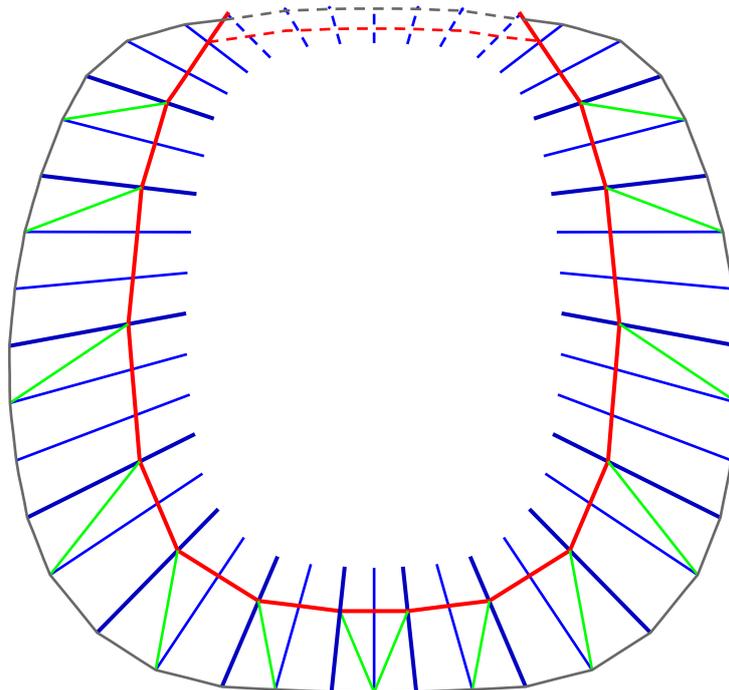Fig 3 Parametric definition of beam locations before (left) and after (right) a change.

Fig. 4 Plan view of structure showing Primary (red), Secondary (green), Tertiary (blue) and Edge (grey) truss locations.
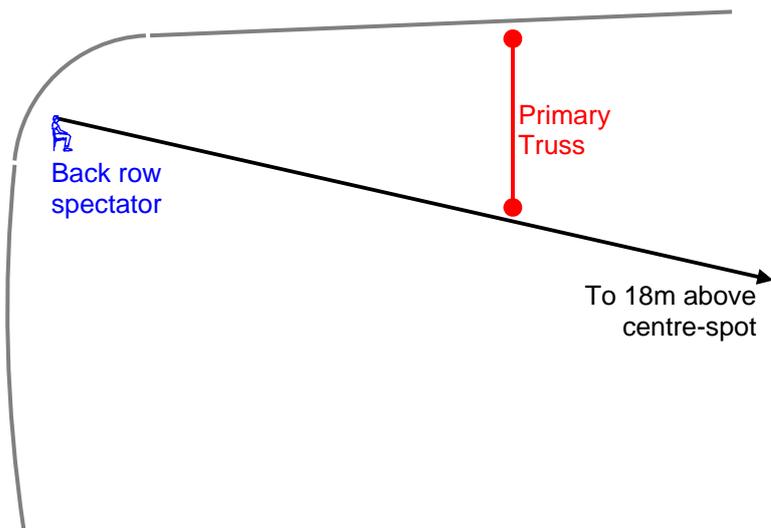
Fig. 5 Elevation view of structure showing Primary (red) truss location defined by offset to structure and sightline surface.
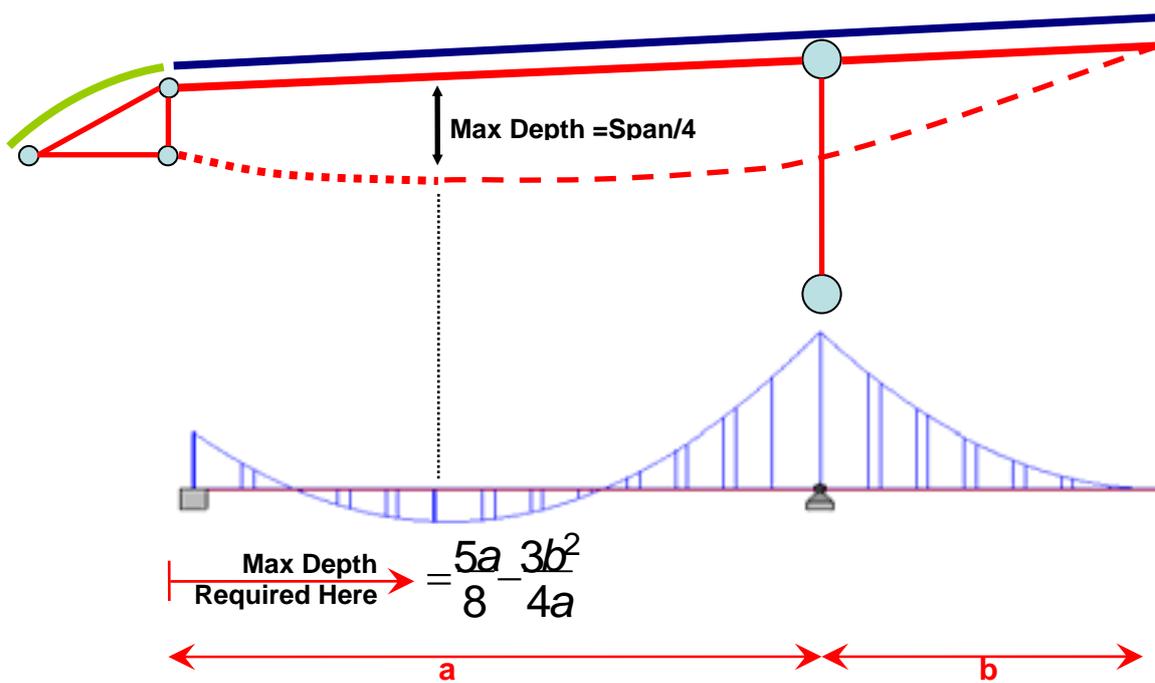

$$\text{Max Depth Required Here} = \frac{5a}{8} - \frac{3b^2}{4a}$$

Fig. 6 Elevation view of structure showing Tertiary truss chord geometry.
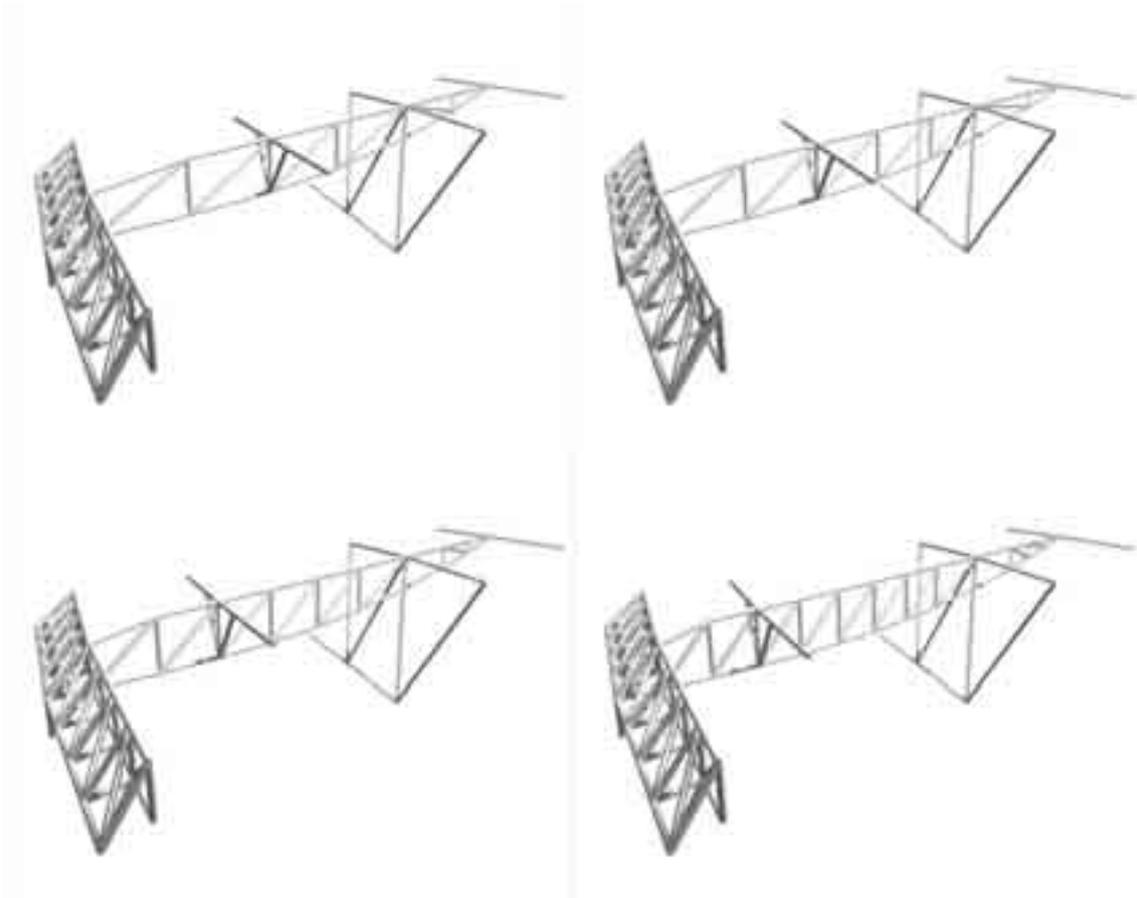
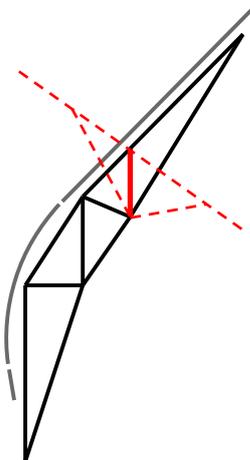Fig. 7 Tertiary truss options with varying spacing of vertical members.

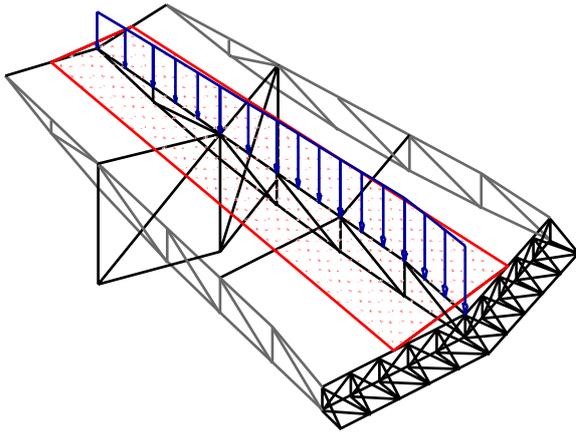Fig. 8 Tertiary truss in the North Stand.

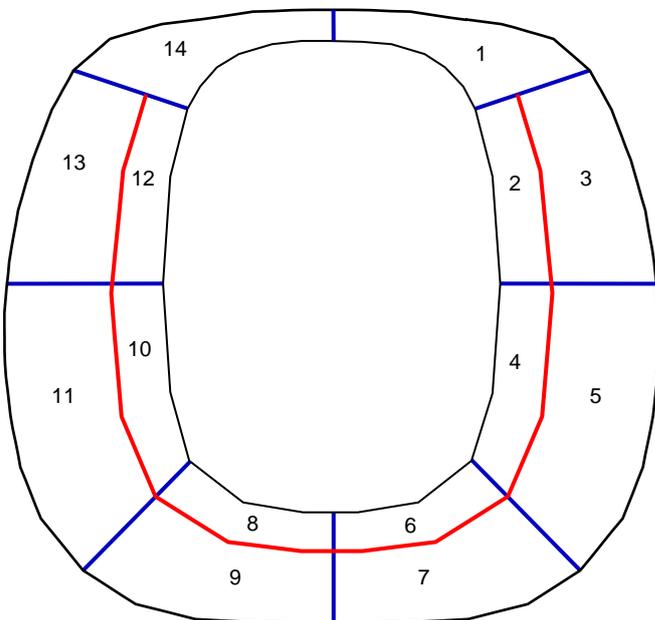Fig. 9 Effective width of tertiary truss (red) and applied triangular loading (blue).

Fig. 10 Idealised loading zones from wind-tunnel tests results