Observations and Points of Friction When Working With: Complex Models in ESP-r

Dr. Jon William Hand[1],
[1] Energy Systems Research Unit, University of Strathclyde, Glasgow Scotland

Abstract

General purpose simulation environments such as ESP-r are typically used in projects of moderate complexity (e.g. a few dozen thermal zones and control loops and mass flow network nodes). Complexity beyond this norm has costs, but the specific nature of the frustrations and extra resources associated with complex models has not been documented and the causal factors investigated.

This paper uses a complex building, as a case study to identify instances where the ESP-r interface, data structures, model reporting or user working methods introduced friction into the work flow.

Points of friction within the model creation process were documented, alternatives (either a software patch, data structure change or a provision for more feedback to the user) implemented and evaluated. Roughly 200 patches were made made coping with mind-boggling levels of detail more tractable. It also discusses where the information included and excluded from CAD drawings and specification impacts deliverable in large projects.

1 Introduction

ESP-r's roots are as a research tool and thus it is over-engineered in terms of the thermophysical assessments it supports. It follows the remit that 'functionality follows description'. A minimal description supports a few high level issues, add more descriptions and more performance information becomes available until issues such as contaminate dispersion and mold growth can be assessed. One must, of course, be passionate when dealing with topics which require extreme attribution of model entities.

Unlike the myriad of constraints which commercial tools put in place, the ESP-r interface and underlying data model allow those with strong opinions to select the level of descriptive granularity needed to create models fit for specific purposes. For example, if you are truly curious and want to represent a thermostat enclosure in extreme detail you can represent the case and internal components explicitly as a thermal zone and surfaces as well as tracking air movement within it via a transient 3D CFD domain. A facade integrated PV embedded in a phase change layer suspended in a double skin facade can also be explicitly represented in terms of its transient optical, thermophysical, air movement and electrical characteristics.

Within the ESP-r community there is a need for alternatives to *walled-garden* commercial approaches and projects which fall outside the mainstream of performance issues. ESP-r has often been employed by practitioners because it supports combinations of analysis facilities or the arbitrary thermophysical complexity observed in cutting edge designs.

The point is, allowing users choose the scale and granularity of the description introduces friction in the model creation processes. The need to manage a rich data model constrains the

scale of models which simulation teams can work with. If we want to support a step-change in model complexity a close look at friction and management issues is a good place to start. As ESP-r is not the only tool to allow user control of model resolution it is hoped that the observations in this paper might be broadly applicable.

2 Goals of the investigation

Much has been written about new knowledge which has resulted from deploying simulation within a research context and about simulation use in cutting edge design projects. There have been many high-level discourses about barriers between simulation and the design process. Much less has been written about the scores of specific points of friction which together form barriers that limit the use of simulation in the design process.

More than a dozen years ago a range of case studies in the UK (Hand 1998) looked at barriers to simulation support for the design process. The findings were:

- Users were required to maintain a mental model of much of the model detail and draw on that information store at various points in the model creation process
- Much information about assumptions made in models, the intent of models and differences between model variants were recorded on slips of paper and easily lost
- Simulation tools did not take into account the level of multi-tasking and the intermittent nature of simulation work, or indeed the often distributed and iterative nature of projects
- Simulation training focused on individual skills and did not take into account team working issues, work-flows and the isolation of managers from the process
- Non-keyboard skills (model planning, model quality checking, recognition of performance patterns) are essential within simulation teams
- Simulation models are often opaque to others in the design process due to poor naming conventions and documentation

A number of these issues were taken into account in the design and implementation of the ESP-r Project Manager. This supported a step change (~one magnitude) in the level of model complexity proficient simulation teams could work with on a sustainable basis at that time. It was hoped that this would allow simulation to be deployed in a broader range of projects.

A follow-on project called the Scottish Energy Systems Group (McElroy 1997, 1999, 2007) focused on embedding simulation within the design process via mentoring and short term deployment of trained staff in small to medium sized enterprises. These deployments provided further evidence to enable the evolution of simulation tools. It also highlighted how simulation tools absorb information from many disciplines within design teams and subsequently deliver performance data from which each of the professions can extract useful patterns.

The range of projects as well as the range of design goals allowed different facets of simulation tools and user skills to be observed. This resulted in incremental improvements in tool functionality and training approaches. Over the duration of the SESG project simulation entered the mainstream of the design process in the UK and in-house skills grew to the point where SESG largely achieved its goal.

One key observation was that simulation staff have a comfort range vis-a-vis model complexity within which they work productively, consistently and within the resource constraints of projects. It was observed that within this comfort range there was scope for practitioners and researchers noticing building and environmental system performance patterns where improvements could be made and sufficient time to test out ideas that could be presented as value-added deliverables to clients. Beyond this point it became increasingly difficult to deliver

planned work tasks, it was less likely that value-added deliverables would be found, risks of errors increased as did user frustration.

3 Context of the investigation

To make another step change in (in this case to double model complexity) this paper focuses on specific points of friction which together form barriers that limit the use of simulation in larger projects. Irritants and distractions which can be accommodated in the work-flow of models of several hundred surfaces can form significant barriers in models of several thousand surfaces. Oft-repeated tasks which require (or assume) user attention and/or confirmation can lead to significant mental fatigue.

The approach taken was to create a model of an administrative building in a UK University to identify instances where the ESP-r interface, data structures, model reporting or user working methods were problematic. The building shown in Figures 1 & 2, has a different floor plate on each level, vertically interconnected spaces and often non-rectilinear layout. The scale of the building often confronts design teams using simulation for building code compliance issues and for which commercial tools are often employed.

Figure 1: Administrative building façade.

Figure 2: Administrative building internal view.

The goals of the project was to study of the influence of changing occupancy patterns on environmental conditions within the building, as well as supporting local comfort assessments and visual assessments. It was also zoned so that in the future detailed studies of air flow within the inter-connected spaces, raised floor system and ceiling returns could be supported.

As with many modern buildings, some sections of the facade were made more complex by their depth and inter-linked internal air spaces. Careful parsing of the plans and sections revealed that the building, which looks like it is composed of a small pallet of constructions, is riddled with exceptions and thus it was necessary to pay particular attention to surface attribution and subdividing the facade in order to capture the different performance characteristics.

Zoning and level of detail

The zoning of the building and the level of detail was defined based on the above goals. This required that the raised floors be explicitly represented as were many ceiling voids and the visually massive end facades. The open plan office spaces were subdivided to capture changes in local occupancy. Some adjacent cellular offices were conflated (internal partitions were retained). The internal light wells and atria were explicitly represented.

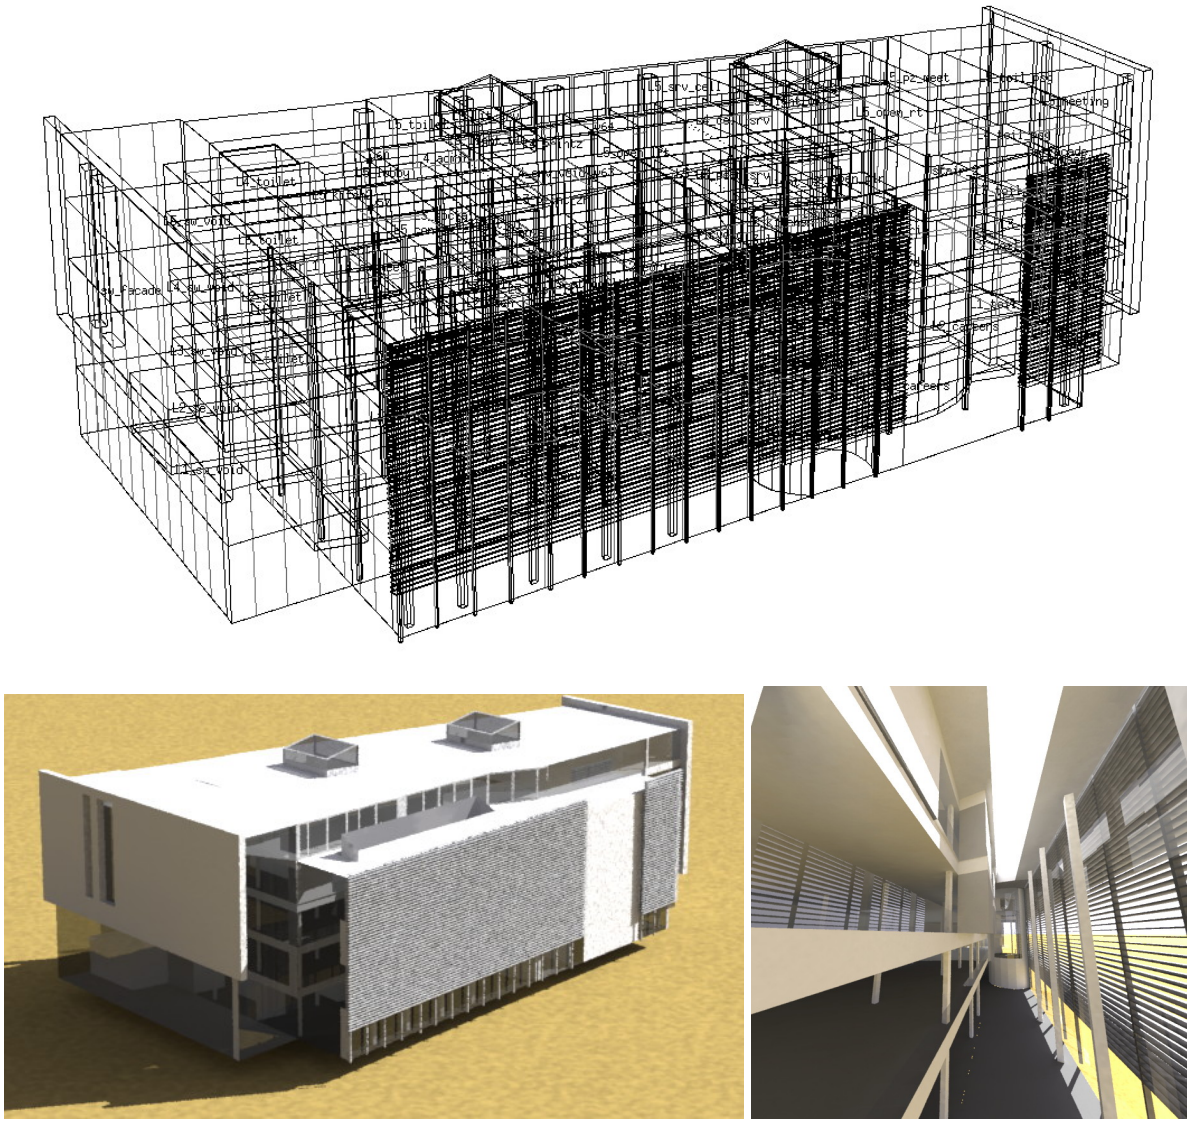

Figure 3: View from south-east and looking along the front façade atria.

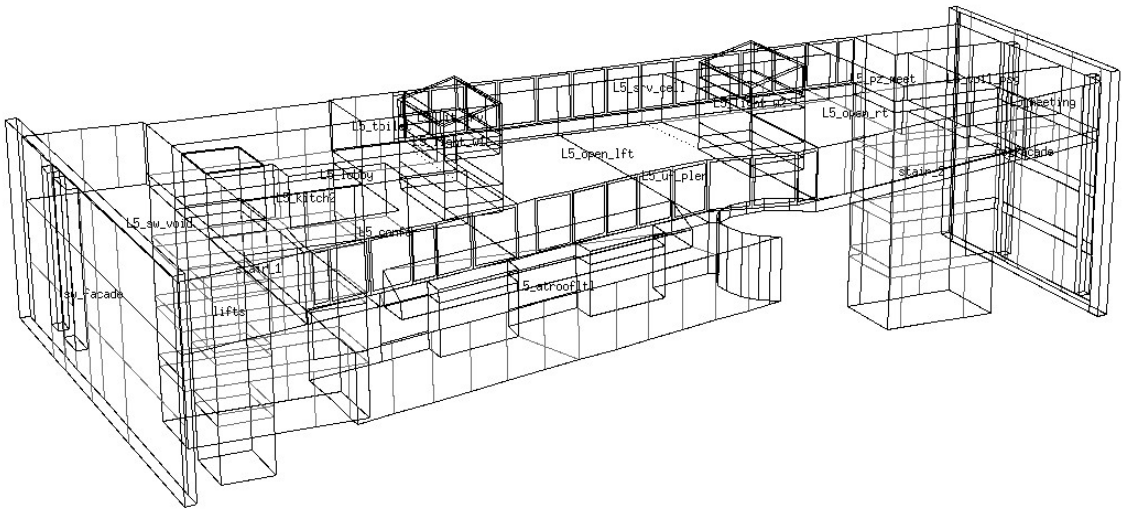

Figure 4: Explicit facades light wells & floor voids.

This is very much in keeping with the user directed detail associated with general simulation tools but which require over 80 thermal zones and roughly 2200 surfaces. This places it beyond the standard complexity limits for ESP-r and well beyond the comfort range of the majority of ESP-r users. It also setup the challenge of how one might subsequently apply the constraints for building code compliance.

4 Identifying friction within working practices

To capture as broad a range of issues as possible the model creation process was undertaken with a very low irritation trigger and a remit to identify the source, investigate the underlying issue and, if possible, to test alternative and then go through the process to test, document, and commit this into the source code control system before proceeding further. In parallel with the evolution in the interface and simulation data structure the evolution of the model itself was also captured in the source code control so that it can be 'replayed'. In the context of this paper irritant triggers include:

- Noticing that information relevant to a task was not available or was located more than a few mouse clicks away
- Not noticing an inconsistency in a model at the point it occurred
- Menu hopping, especially in repetitive tasks or delay in list selection tasks
- Evolving the model and finding that the model contents report did not fully reflect the change
- Elevated attention required to locate and address portions of the model which had not been fully attributed
- Instances where knowledge of adjacent portions of the model would clarify tasks
- Unresolved dependencies when changes were made to a portion of the model
- Geometric issues (e.g. clusters of nearly-identical vertices, misalignments between surfaces in adjacent spaces)
- Visual clutter obscuring relevant details and resulting in error or delay
- Points of interaction confusion leading to incorrect selections or elevated attention requirements
- Using a calculator to derive a value which the interface should have provided

From the start of the investigation until the initial simulation were run over 200 commits had been made to the repository. The space in this paper is not sufficient to fully describe each intervention. Instead some of the notable issue will be discussed.

Information availability

Much of the evolution of ESP-r during the 1990's case studies and during the SESG project centered on ensuring that model attributes were included in the model contents reports and available at some point within the feedback within the interface. In less complex models the need to navigate within the interface or consult an external document for data to support a specific task was a bit of background noise in the work-flow.

As the administrative building began to take form the triggers points for information availability were set-off. Many users attempt to save time in complex models by accepting default names for model entities. This saves a couple of seconds for each surface but results in a model that is much less descriptive (see Figure 5). This *opacity* slows down selection tasks. A drill-down to a specific surface can easily take twice the time. In complex models the additional mental effort to clearly identify the intent of surfaces impacts work-flows. The additional effort needed to review complex models with default entity names is nothing short of debilitating.

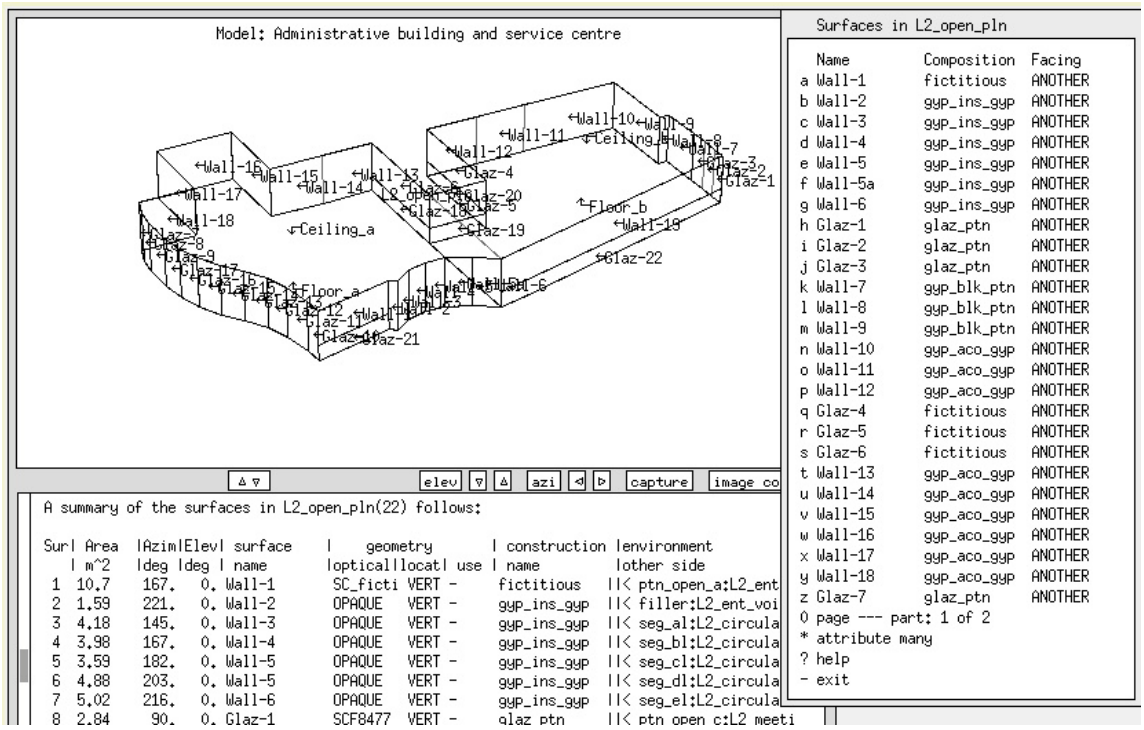

Figure 5: Default surface attributes and selection.

ESP-r has traditionally provided a number of self-reinforcing feedback mechanisms. For example, Figure 6 is focused on a surface within a zone. The majority of the attributes of the surface are included in the menu along with the attributes of the surface in the adjacent zone. In conjunction with the wire-frame view and the attribute shown in the feedback window below it, information availability should not have been an issue. And it was still necessary to jump out of this view, for exempt to gather information on the coordinates associated with the surface and nearby surfaces.

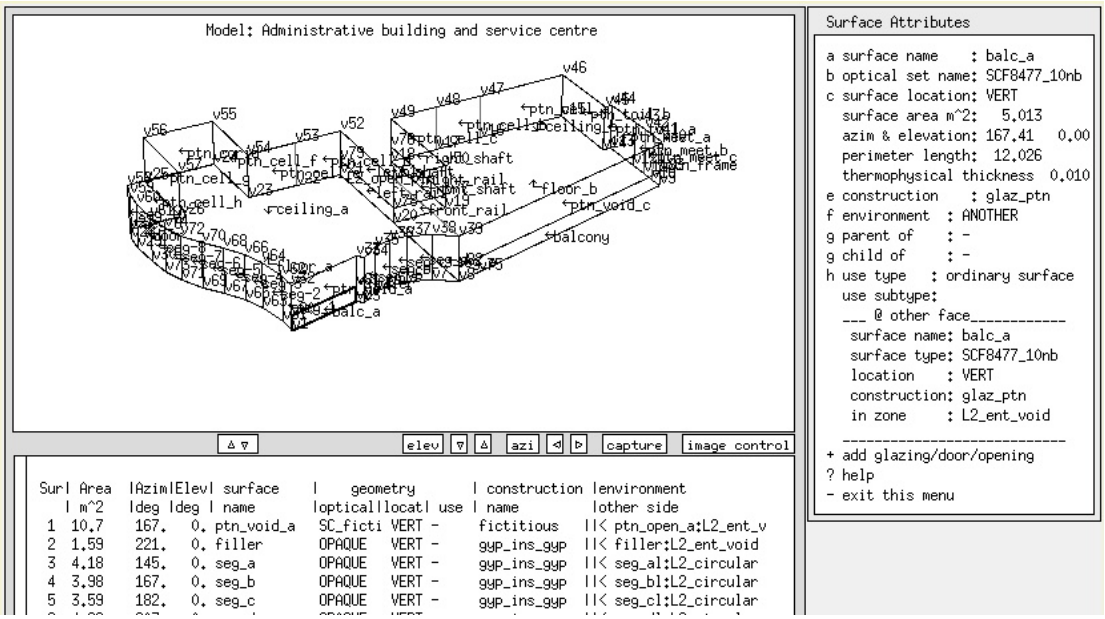

Figure 6: Traditional self-reinforcing attributes.

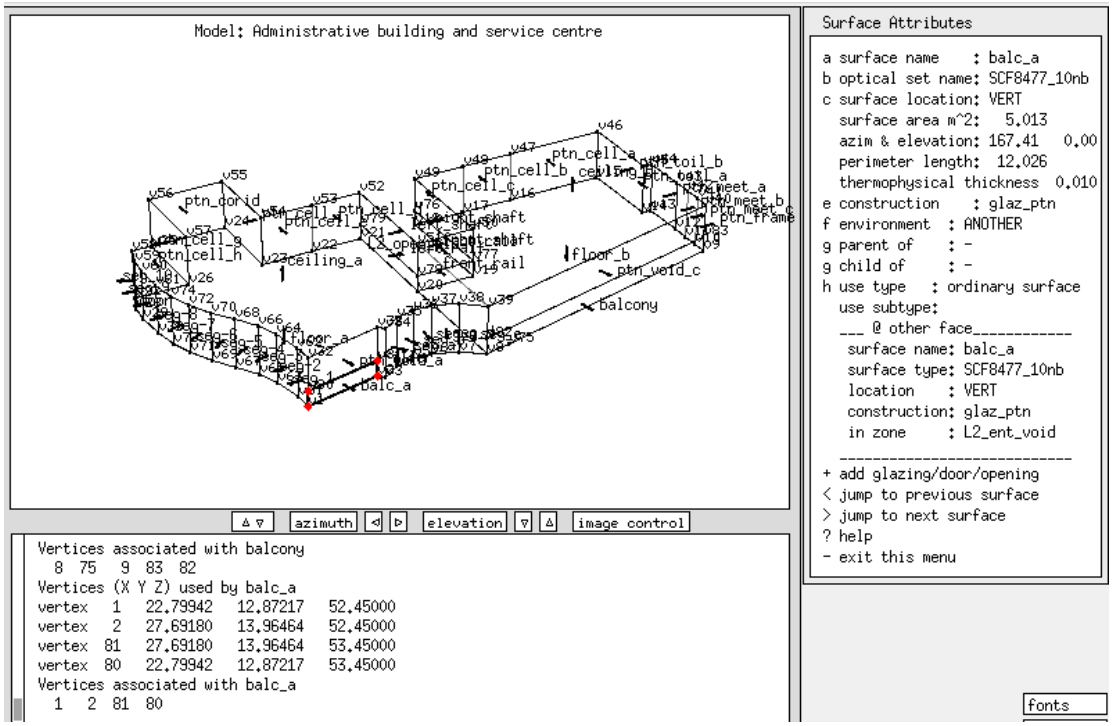

Figure 7: Revised interface with color, extra reporting and prev:next.

Opening up a second instance of the project manager with a replica of the model to allow adjacent zone information to be viewed (either side by side on a large monitor or on adjacent monitors) was helpful as was a hardcopy of the model contents report. Clearly, what was needed was to ensure that data, such as the coordinates and edges associated with the surface, were echoed to the feedback on entering the surface menu. As shortcomings were noted and the feedback extended much of the need for traversals of the interface was eliminated and the revised interface is shown in Figure 7.

The model contents report in ESP-r evolved during the SESG project to provide critical support for checking model quality. It is, in many ways more complete that the offerings of many commercial tools (where the raw model files or screen captures are the only facilities

available to document models). Figure 8 is a portion of a model contents report where the geometry of the zone is reported in a verbose rather than a standard level of detail.

```
Zone L2_open_pln (22) is composed of 46 surfaces and 83 vertices.
 It encloses a volume of 1350.m^3 of space, with a total surface
 area of 1220.m^2 & approx floor area of 437.m^2
 L2_open_pln describes main open plan office area at level two

 A summary of the surfaces in L2_open_pln(22) follows:

 Sur| Area  |Azim|Elev| surface    |    geometry       | construction |environment
    | m^2   |deg |deg | name       |optical|locat| use | name         |other side
  1  10.7    167.   0. ptn_void_a   SC_ficti VERT -     fictitious    ||< ptn_open_a:L2_ent_vd
  2  1.59    221.   0. filler       OPAQUE   VERT -     gyp_ins_gyp   ||< filler:L2_ent_void
  3  4.18    145.   0. seg_a        OPAQUE   VERT -     gyp_ins_gyp   ||< seg_al:L2_circular
  4  3.98    167.   0. seg_b        OPAQUE   VERT -     gyp_ins_gyp   ||< seg_bl:L2_circular
  5  3.59    182.   0. seg_c        OPAQUE   VERT -     gyp_ins_gyp   ||< seg_cl:L2_circular
  6  4.88    203.   0. seg_d        OPAQUE   VERT -     gyp_ins_gyp   ||< seg_dl:L2_circular
  7  5.02    216.   0. seg_e        OPAQUE   VERT -     gyp_ins_gyp   ||< seg_el:L2_circular
  8  2.84     90.   0. ptn_frame    SCF8477_ VERT -     glaz_ptn      ||< ptn_open_c:L2_meeting
 . . .
 45  5.01    167.   0. balc_a       SCF8477_ VERT -     glaz_ptn      ||< balc_a:L2_ent_void
 46  17.9    166.  -0. balcony      SCF8477_ VERT -     glaz_ptn      ||< balcony:L2_careers

 Vertices ( index X Y Z) in the zone...
vertex   1   22.79942   12.87217   52.45000   vertex  42   51.60000   19.23701   55.59300
vertex   2   27.69180   13.96464   52.45000   vertex  43   51.60000   20.42448   55.59300
vertex   3   28.07178   13.63215   52.45000   vertex  44   51.60000   23.17941   55.59300
vertex   4   29.16426   14.39213   52.45000   vertex  45   51.00000   23.17941   55.59300
. . .
vertex  39   34.25000   13.06217   55.59300   vertex  80   22.79942   12.87217   53.45000
vertex  40   51.60000   17.28956   55.59300   vertex  81   27.69180   13.96464   53.45000
vertex  41   51.60000   18.19204   55.59300   vertex  82   34.25000   13.06217   53.45000
                                              vertex  83   51.60000   17.28956   53.45000
 Number of vertices followed by list of associated vert
*edges,4,80,81,33,32   ptn_void_a
*edges,5,2,3,34,33,81   filler
*edges,4,3,4,35,34   seg_a
*edges,4,4,5,36,35   seg_b
. . .
*edges,14,20,19,18,17,16,15,14,13,12,11,10,9,75,8   floor_b
*edges,13,39,40,41,42,43,44,45,46,47,48,49,50,51   ceiling_b
*edges,4,82,83,40,39   ptn_void_c
. . .

 All surfaces will receive diffuse insolation (if shading not calculated).
 No shading analysis requested.
 No insolation analysis requested.
```

Figure 8: Portion of model contents report.

In the context of the administrative building the trigger point was the need to view the raw model files for clarification. The standard report contained all of the surface attributes and many derived values (e.g. perimeter length, parent-child relationships) but did not include the vertex coordinates or the surface edge lists or a report on linear thermal bridges.

For example, it was found that some of the facade and structural coordinates at lower levels diverged from those in the upper levels and this error was not picked up during the initial coordinate capture phase. Aligning coordinates over several levels proved to be tedious. Use of side-by-side monitors helped greatly as did the inclusion of additional information in the model contents report seen Figure 8.

Menu hopping and list navigation

ESP-r uses a distributed file system wherein different classes of information are held in different model files (e.g. environmental controls are separated from definitions of room activities and the composition of each zone is held in separate files. The layout of the interface

tends to reflect this separation of information types. High-level separation of topics was generally not found to be problematic.

As the number of entities (e.g. database entries, zones, surfaces etc) grows the burden of navigation became noticeable. This is in part a result of the ad hoc evolution of models. It is necessary to change the focus within the interface to instantiate entity attributes that were previously treated as placeholders, resolve dependencies or to adjust model resolution.

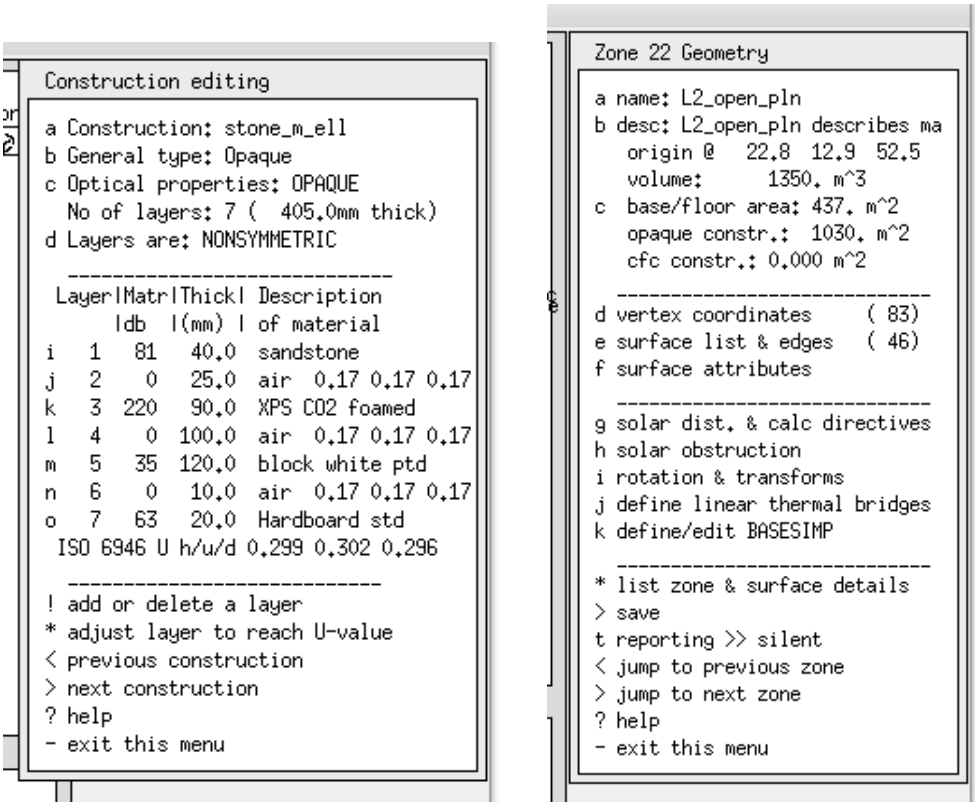

Figure 9: Jump to previous and next entity.

Another observation in large models which use explicit representation of rooms is that the number of internal partitions grows in relation to facade surfaces. In retrospect, the initial sketches and model planning could have been more detailed – it was found that some of the complex spaces, such as the stair towers, which involved lots of connections to adjacent spaces and structural elements approached the limits of individual zone complexity (later relaxed) and that additional attention was required not to 'paint oneself into a corner.

In a hierarchal interface where views such as Figure 6 are accessed via exiting back to the list of entities and selecting another entity is hardly worth mention is a small model where the whole of the list is within view. As lists grow the utterly trivial task of changing focus to another item in the list was no longer trivial. A quick experiment allowing a jump to the next or previous entity in the list without returning to the list was shockingly effective in supporting rapid traversal of the model. In the days that followed jump to options were enabled for several other entity types and lists as in Figure 7 and Figure 9.

Visual clutter

As we move from highly abstracted representation of rooms and attempt to explicitly represent commercial facades or the internal clutter within rooms the wire-frame views become obscured by visual clutter. Names or indices in the wireframe may not be clear and this has

many consequences: difficulty in identifying the focus surface in the wireframe and uncertainty or errors when editing surface edge lists.

A number of approaches were tested and the most effective was to use color in addition to bold lines when drawing the focus surface, draw the associated vertices index in color and list details in the feedback area below (as seen in Figure 6). In combination with the prior:next jump the effort required to select the correct entity dropped dramatically.

User attention

One of the productivity and model quality issues from the 1990's case studies was the definition of surface boundary condition attributes. Setting the boundary conditions on the other-face of each surface was initially required a rigidly sequential specification and considerable user frustration. This augmented by the ability to set boundary attributes at an individual surface basis as well as an automated scan of the model looking for matching surfaces with user confirmation of boundary conditions which did not meet the matching criteria. Visual feedback was included as each surface was considered. During SESG additional logic was included to reconcile the thermophysical composition of matching surfaces and to ensure that subsequent attribution changes were reflected in the matched surface.

In the initial stages of the administrative building use of the facility triggered action because it was noticed that matching surfaces within a room (in this case the partitions used for thermal mass and visual assessments within the conflated cellular offices) were not recognized. Normally one would re-run the facility as the model evolved. The visual pause and confirmation was particularly useful for identifying errors in the initial digitizing of points and the surface composition matching reduced some of the burden of surface attribution.

However, as the surface count passed ~1200 the time required for a scan was 15-20 minutes with 35-40 minutes for ~2300 surfaces (accounting for the visual pause at each surface and the need for user confirmations of facade surfaces). The level of user attention required was increasingly debilitating and occasional confirmation errors were detected. Turning off user confirmation speeded up the process but some borderline matches were incorrectly resolved and treated as facade surfaces.

The core issues were how to reconcile the addition of new zones into an already large list of existing surfaces and how to resolve dependencies in adjacent zones when existing zones were modified without the need for an exhaustive search. This lead to a 'what is close' option in the surface attributes menu which scanned the model for likely matching surfaces which were then reported to the user for confirmation. This allowed the exhaustive search to be delayed until the model was complete.

Digitizing and CAD data

In most cases, simulation teams have little or no influence on the nature and resolution of the CAD data used in a project. In larger buildings one is confronted by tens of thousands of 2D lines (which tend to be of little use) surfaces (which might be useful) or block entities (which might turn out to be plumbing fixtures, vases or façade elements). The choice is to spend time filtering the CAD file for useful entities or to use manual selection of points based on project criteria for zoning and level of detail. In this project manual selection was used (see Figure 10). For buildings of this scale it is definitely a plus to be able to work on the highest resolution monitor available so that the image coordinates resolve to a useful resolution.

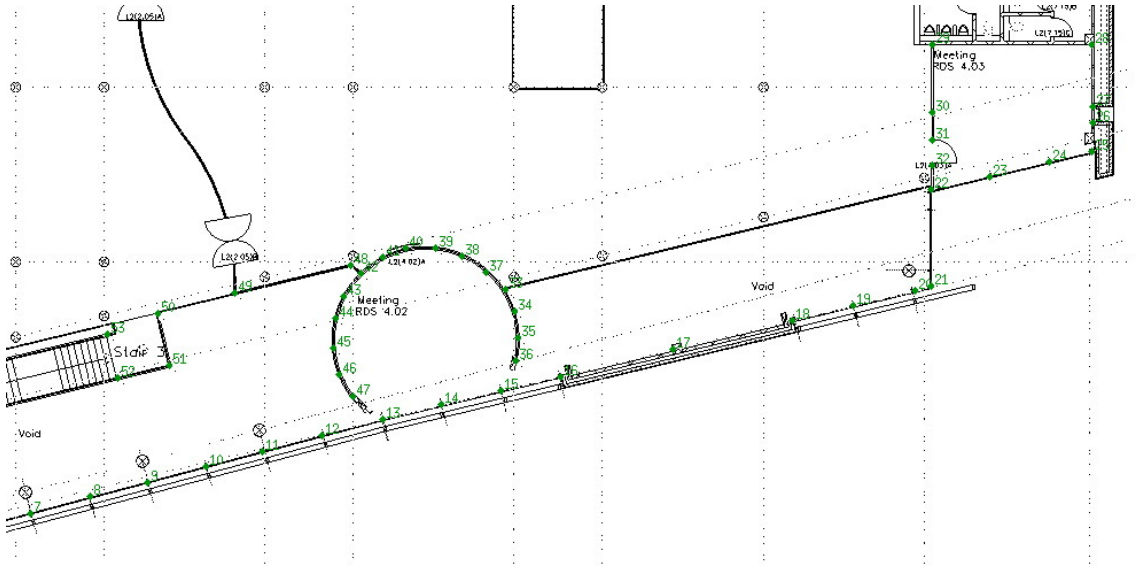

Figure 10: Collecting coordinates from plans.

One must also decide whether to click-to a user defined grid or take the raw coordinates. In the case of non-orthogonal rooms and façade elements this resulted in a lack of precision. Because the floor plate differed on each level it was necessary to capture a fresh set of points and, inevitably, there were mis-alignments that needed to be sorted out in subsequent steps. The most problematic tended to be points which were only mis-aligned by ~5-15mm. These are not quite close enough to be considered a duplicate and very difficult to identify in the wire-frame views.

With hind-sight it would have been worth accessing a higher resolution monitor as well as using a plan image of ~2400 pixels wide rather than 1400 pixels. It would also to have isolated the vertically aligned spaces such as stair-wells, toilet blocks and light wells and built them from the ground level coordinates rather than from points on each level.

5 Discussion

One might argue that some of the navigation-within-the-model issues highlighted in the paper are also problematic for commercial tools. Long lists are ubiquitous as is the risk of mis-selection and getting lost. Colour in geometric views to denote constructions are problematic when there are scores of construction variants in a façade. Real buildings are overflowing with exceptions to nominal façade cross-sections.

Clarity is also an issue in commercial tools. What should be a straightforward task of identifying a surface and viewing its attributes is often not straightforward. For example, IES-VE has a great wire-frame view but to find out what you are looking at requires a traverse of the zone:surface tree structure. Why isn't a surface clickable like it is in Sketchup-based interfaces? Why doesn't that action immediately report surface attributes, derived values and a name that means something to me?

In some commercial tools what you see if not necessarily a full representation of the model. For example, commercial facades are often characterized by their framing elements. IES-VE treats frames as an attribute of the glass and does not display framing. It is tedious to check and is confusing to clients viewing the model on screen. Older versions of TAS merged surfaces composed of the same construction. Users expecting to find surface areas matching their documentation find something else in the tool. A toggle to show explicit or simplified views would be most helpful.

The need for frequent drilling-down into the interface when checking a model with thousands of surfaces is a friction in workflow for many commercial tools. Generation of a report of the contents of a model and the attributes of entities within the model is a rudimentary concept in some commercial tools. Finding out the differences between two models or between today's model and yesterday's model is a task that most tools could do better at.

During the decade that SESG was active there emerged two trends which impact user defined model resolution. Firstly, the emergence of building code performance standards which encouraged the use of simulation has also perturbed the nature of simulation models and constrained the range of performance metrics which are typically assessed within the design process. Whereas in the early years of SESG design teams tended to use simulation for a variety of what-if enquiries (which could often be addressed with highly focused models), the latter period was increasingly dominated by constrained models that were used for code compliance or identification of a performance certificate levels.

When considering how the administrative model might be adapted for code compliance the methodically valid zoning and surface resolution decisions taken to address the project goals need to be altered. For example, there is no UK code concept for explicit treatment of massive hollow façade elements as zones or for that matter explicitly defined ceiling and floor voids or even stacked light well zones. It is certain that these spaces will diverge from the temperature of conditioned spaces and are critical to ensuring that the massing and volumes of spaces are preserved. Other conversions are mechanistic, for example, facades which are composed of frameless glazing must be placed in frames. Some thermophysical rigour must be removed e.g. numerical solution of air movement between rooms must be disabled. And whereas the initial model did not include thermal bridges, this topic is mandated for code compliance.

The other major issue, which partly stems from compliance tasks is the tendency to assume that one must model every room of a project, no matter how large or repetitious. Such brute force approaches involving hundreds of rooms tend to result in large, but low resolution models which are utterly tedious to create, horrific in terms of model quality control and of little use for investigating other design issues.

If the initial goals of the case study were primarily to support compliance then the model would generally have been simpler to compose. It would be an abstract representation diverging somewhat from the details in the plans and sections vis-à-vis the thick façade elements and the lumping of the floor and ceiling voids would require the vertical voids and atria to collapse and the environmental system being controlled on room rather than floor-void temperature. The model would, of course require much adjustment to make it fit for some of the performance issues mentioned in section 2. It is likely, though that the investigation into friction points based on a compliance model would have detected many of the same faults.

Using models that are not fit for purpose introduces risk and uncertainty into the design process. The effort needed to re-purpose a model could need to be considered on a case-by-case basis. From this one might conclude that opinionated users of simulation tools will opt for creating separate models for design exploration and compliance purposes.

6 Conclusions

This paper described an investigation of friction within a project involving a building that pushed the complexity limits of ESP-r. It identified a number of issues which were seen to obscure user tasks, delay work-flow and occasionally lead to mental fatigue. Over a one month period over 200 commits were made to the repository. Some of the points of friction were discussed in detail along with alternative interface or data model approaches that were seen to mitigate at least some of the friction and improve the clarity of models.

Unfortunately, it has not been possible to start from scratch with an evolved interface and data structure in order to benchmark shifts in the resources required and in users complexity comfort range for the administrative building. Certainly, the time to create, attribute and check the latter portions of the model were markedly improved (~30-40%) over the initial models. Navigation tasks were often a fraction of the time after interventions.

Proficient users of ESP-r will tend to find the interventions alter the resources and attention required for small and medium scale projects as well as allowing them to attempt projects of greater complexity. In recent workshops where novices were exposed to updated versions of ESP-r they were seen to stumble less on some classic confusion points and the additional feedback made it easier to explain relationships within their models.

ESP-r has traditionally been extreme in its use of entity attribution. It supports the idea that all entities in a model should be named by the user. In models approaching the complexity of the administrative building clarity is essential. One of the constraints imposed by some simulation tools and most CAD tools is to dis-allow the naming of surfaces. It also supports explicit representations of thermal clutter within building (which is also useful for creating models that are recognizable to clients) as well as explicit representations of façade framing (some tools force the treatment of façade framing as a non-geometric attribute of glazing).

Few tools allowing tagging of surfaces as to their 'use' within the building – for example a glazed surface can be marked as requiring compliance with façade regulations as well as whether it is fixed or operable and if operable the type of opening and the existence of linear thermal bridges at each of its edges. parent child relationships, acoustic and phase change attributes.

This super-set of attributes allows ESP-r to export to other tool and other domains and in the future to make further inferences about models as they evolve. The down-side is that getting users involved in confirmation of attributes and allowing them the option to be explicit in their representations of building elements takes time. Attribution can add clarity and it can require additional attention from users to check for correct attribution which was one of the mental frictions observed in the study. Are their ways to mitigate this? Further work is required.

The evidence of this project is that common assumptions that a simple pallet of constructions can be automatically applied are not justified. Real building sections are often revealed as more complex and varied than we initially assume. ESP-r has a terse naming regime for construction types and this is less and less appropriate given observed complexity of building sections. This remains an loose end in the investigation – there was not time to revise the data structure to support a different construction naming regime and associated notes.

7 Acknowledgements

The author wishes to thank staff of the Sustainable Technology Research Team of Samsung Construction in Seoul for occasional exposure to these new facilities as a benchmark against the author's observations.

8 References

Hand, J., 1998. *Removing Barriers to the use of simulation in the building design professions*. PhD Thesis, University of Strathclyde, Glasgow, Scotland.

McElroy L., Cockroft J., Hand, J., 2007. Business success through process based application of simulation. In Building Simulation 2007 Conference. Bejing, China.

McElroy L., Hand, J., Strachan P., 1997. Experience from a design advice service using simulation. In Building Simulation 1997 Conference. Prague.

McElroy L., Clarke J., 1999. Embedding simulation within Energy sector Businesse. In Building Simulation 1999 Conference. Kyoto.