EDG 91 NP • October 2017 (Edited 07.10.2018)
acumen.architecture.com.au/environment
Reviewed and approved for currency in October 2018

ISSN 2207 - 466X

Optimising environmental performance using building performance simulation

Nicki Parker, Paul Bannister, Quentin Jackson, Paul Stoller

Cover image: South Australian Health and Medical Research Institute (SAHMRI), Adelaide, by Woods Bagot (Image: Peter Clarke).

Abstract

Computer simulation of a building's environmental performance has been available as a technology for over 30 years, during which time the accuracy, depth and speed of simulation have all significantly improved. However, it is arguable that industry's use of computer simulation has not kept up with the potential contribution this technology can make to the design and construction of buildings.

This article provides an outline of how dynamic thermal simulation (energy modelling) and daylight simulation methods can be optimally applied to the design and construction process for new buildings, with an exploration of how Atelier Ten applied this to the South Australian Health and Medical Research Institute (SAHMRI).

Introduction

What is building performance simulation?

Computer simulation of a building's environmental performance is a generic categorisation typically used to cover three related but separate applications:

1. **Dynamic thermal simulation of a building**
   This is the assessment of heat flows, internal thermal loads, solar loads through the fabric and technologies of a building to calculate achieved space temperatures and energy use. Simulations are typically calculated on an hourly basis using real weather data. A subset of this category is that of peak heating and cooling load-calculation programs, which are typically non-dynamic (ie do not account for the effect of thermal mass and associated lag effects) for the purpose of sizing mechanical plant.
2. **Daylight simulation**
   Refers to the calculation of daylight levels under various conditions representative of the site. This may include tracking direct sun paths through the building and the calculation of daylight levels under diffuse skies. Simulation of glare can also be categorised as daylight simulation, and is increasingly being used to understand the risk of glare occurring for particular built-form and facade properties.
3. **Computational fluid dynamics (CFD)**
   Calculates airflow and temperature patterns within individual spaces or in the external environment. Applications can include analysis of mechanical ventilation strategies to assess whether comfort conditions can be maintained within a space, through to external microclimate analysis, where wind flows around buildings or pollution dispersion can be simulated to understand what might happen under certain weather scenarios.

Some computer simulation tools can use a single building spatial model to provide a degree of integration between these three applications, but as a whole they are conducted separately and can be considered independently.

There are many other types of simulation tools that are used in the built environment such as structural performance and moisture transfer simulation; however, this article will focus on dynamic thermal simulation (or energy modelling as it's more commonly known in the industry) and daylight simulation methods and how they apply to the environmental performance of buildings. While CFD is frequently used to demonstrate and improve building performance, the capabilities of this tool are wide ranging and beyond the scope of this paper.

How building performance simulation is used

Simulation is most commonly used for National Construction Code (NCC) and green rating tool compliance. For NCC compliance, this could be in terms of a NCC JV3 model, demonstrating that minimum energy performance of the building facade and services are at least equivalent to those defined using J1 to J8 prescriptive measures. Green rating tool compliance could include testing achievement under Green Star daylighting credits.

In theory, simulations of this type can provide significant value to projects by freeing the designer from the constraints of Deemed-to-Satisfy (DtS) processes within codes and standards; in this manner significant capital cost benefits can be realised. However, in practice much of this potential value is lost because the analyses are undertaken well after the major design decisions have been made by other team members. As such, the simulation work becomes merely a validation of existing design rather than a tool for innovation or optimisation. This fails to capture the full value of simulation.

Optimising simulation for different stages

The best value is obtained from simulation when it is used to inform all stages of building design and construction from concept through to post-construction (Figure 1).

Concept Design – setting goals, choosing systems or approaches

The largest opportunity for built-environment designers is to use simulation to compare design decisions and opportunities during concept design. This requires buy in from the whole design team, not just architects, in order to reach a common optimised goal. Early phase design decisions have the largest impact on performance, yet are those made with the least effort (Figure 2). When simulation is undertaken in these early phases, the simulation work itself can be simple and agile, due to the limited amount of design detail. This contrasts to simulation work undertaken later in the project when models will be inherently more complex, slower to run, and less capable of informing the design process in a timely manner.

Preliminary analyses can be used to inform key design decisions, such as:

- Building siting, massing, layout and orientation for energy, daylight and external wind flows
- Glazing size, shading and type for energy and daylight
- Comparing different HVAC system or central plant types for relative energy efficiency
- Performance goal setting (e.g. NABERS, Green Star, Net Zero) or preliminary feasibility assessment.

There are many software tools that can be used in early modelling, ranging from full simulation tool suites such as EnergyPlus or IES through to dedicated preliminary design tools such as Sefaira, Grasshopper and Ladybug or Honeybee (Refer EDG note: Sustainability and Building Information Modelling (BIM) for an overview of these tools).

The key to successful early phase modelling is to create abstract, simplified models of the building that allow quick testing of design ideas at a pace matching that of early design progress; in some ways presenting the antithesis of traditional, engineering-led whole-building modelling based on confirmed design decisions rather than reasonable assumptions. It is also important to understand that preliminary design models are just that – a tool that can be used to compare high level options and variants, and should not be used to give definitive and absolute outcomes, such as final NABERS ratings or JV3 performance.

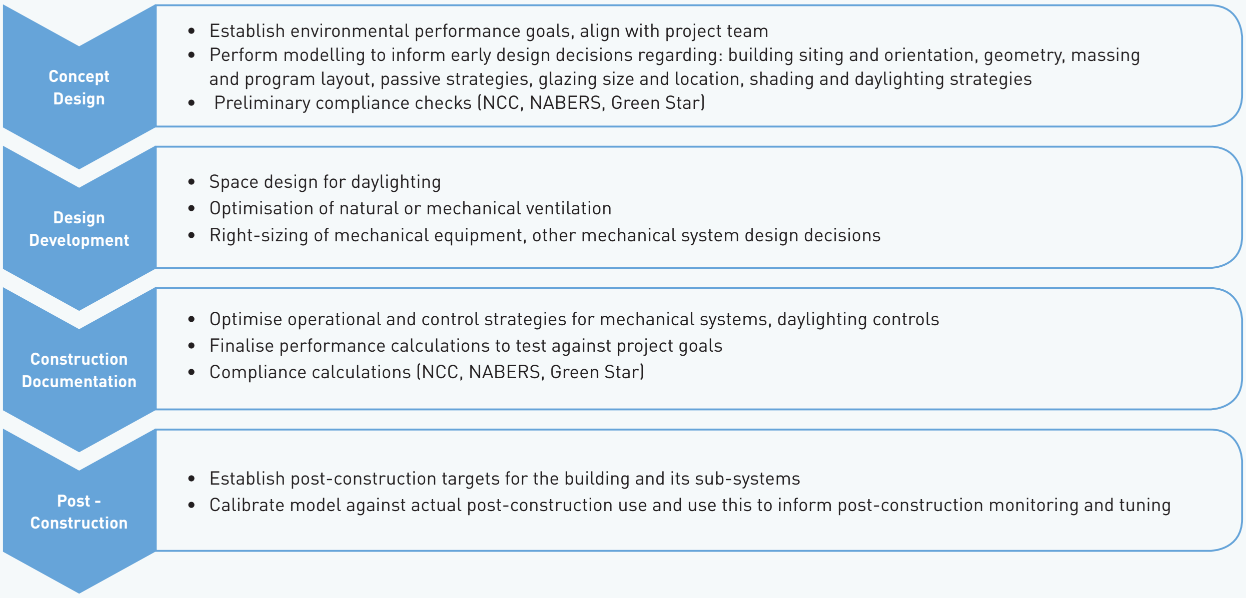

Concept Design

- Establish environmental performance goals, align with project team
- Perform modelling to inform early design decisions regarding: building siting and orientation, geometry, massing and program layout, passive strategies, glazing size and location, shading and daylighting strategies
- Preliminary compliance checks (NCC, NABERS, Green Star)

Design Development

- Space design for daylighting
- Optimisation of natural or mechanical ventilation
- Right-sizing of mechanical equipment, other mechanical system design decisions

Construction Documentation

- Optimise operational and control strategies for mechanical systems, daylighting controls
- Finalise performance calculations to test against project goals
- Compliance calculations (NCC, NABERS, Green Star)

Post - Construction

- Establish post-construction targets for the building and its sub-systems
- Calibrate model against actual post-construction use and use this to inform post-construction monitoring and tuning

Figure 1. Opportunities for the use of simulation throughout the design and construction process

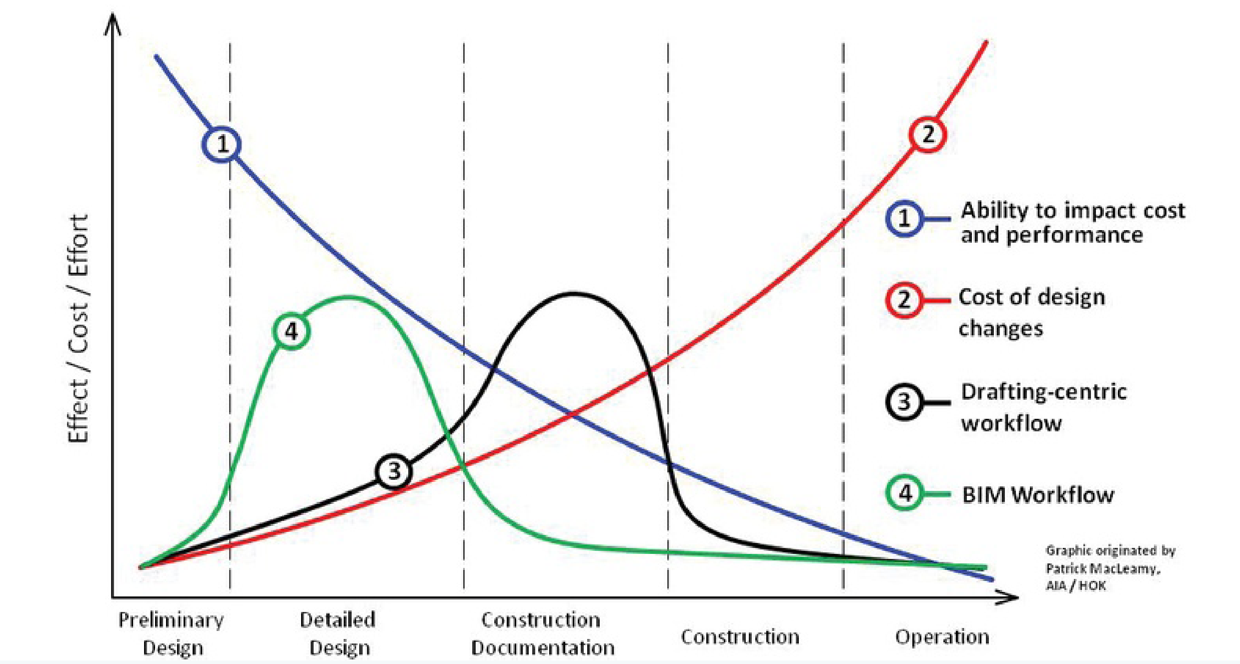

Figure 2. Projects commonly only use simulation in the latter phases of design in order to validate earlier decisions. Greater value can be obtained by earlier use of simulation to shape those decisions. (Image: MacLeamy, P. HOK, 2009)

Design Development – optimising systems

In design development, simulation has a myriad of uses in informing the decisions of architects and engineers. Key opportunities include:

- Detailed development and testing of glazing and shading design
- Optimisation of insulation
- Space design for daylight optimisation
- Right-sizing of mechanical plant to meet loads without excessive overcapacity
- Design of more complex mechanical airflows such as underfloor air distribution (UFAD) or for more challenging spaces such as theatres and atria
- Natural ventilation airflows through buildings for cooling in temperate periods and overnight
- Testing of efficiency options such as alternative mechanical systems and heat recovery
- Optimisation of central plant sizing to match load profiles.

Of the above, the question of right-sizing (or ensuring the mechanical plant is correctly sized for the building with respect to the external environment, building envelope and operational use) serves to illustrate the value of simulation. Despite the availability of high quality simulation tools, there is extensive use of rule-of-thumb and spreadsheet calculations for preliminary equipment sizing. Such methods tend to routinely oversize plant items, causing unnecessary capital costs and poorer efficiency. The extent of this oversizing has arguably increased in recent years, as rule-of-thumb calculations are often based on experience that predates Section J of the NCC, and the impacts this has had on facade performance, a major determinant of building load. Load-calculation tools even fail in this respect as they generally do not allow for the moderating impacts of thermal mass on building performance and thus also tend to oversize. In the authors' experience, the additional costs of oversized plant will typically far outweigh the costs of simulation.

Similar considerations apply to other opportunities: it is not unusual to find buildings designed with shading systems that don't work effectively; natural ventilation that doesn't provide adequate air movement for effective cooling; lost opportunities for major efficiency improvement and equally poorly directed investments in efficiency that will not deliver the desired results. Timely use of building simulation can highlight these issues before design is finalised, saving time and money as well as optimising building performance.

Construction Documentation – optimising controls, documenting anticipated results

In construction documentation, the opportunities for changing design have largely passed, but simulation can continue to add value through:

- **Testing and optimising control**
  The energy efficiency and comfort of modern buildings is highly dependent upon control. High-end simulation packages enable control strategies to be tested and optimised so that these can be documented for implementation. The alternative is for these strategies to be put together by the controls trades, often with little visibility of the project goals.
- **Testing the full design against project goals**
  The project may have diverse goals relating to occupant comfort, energy efficiency, on-floor daylight or many other aspects of the environment that the building will provide to occupants. The construction documentation phase is the last chance to test the performance against these goals using simulation; where not met, late-stage design adjustments can be made.
- **Testing compliance**
  Many buildings will have NABERS, Green Star or NCC JV3 compliance requirements. While these should all have been progressively tested through design development, the later stages of construction documentation are when the final compliance models need to be produced to ensure that the final design delivers the required outcomes.

Post-construction – verifying results

While simulation has historically been viewed as a component of the building design process, its use in post-construction verification is increasing. Given the widespread use of simulation to assess performance potential, it is a natural progression to compare the simulation to the post-construction performance to understand the differences, which may represent opportunities for improvement in control commissioning, tuning and modelling.

There are some complexities in comparing measured performance to simulation estimates, many of which reflect a simple reality: the building will not be operated the same way as it was modelled, or data may not be collected reliably. At a basic level, such as a NABERS rating, this may not matter too much if the differences in operation are not substantial, such as minor changes and differences in operating hours. However, for a more detailed view it is best to undertake some calibration of the simulation to match the building.

Calibration is the process of replacing assumptions made in the simulation model with actual data (such as energy use by lifts, car parks and occupant equipment). It is also possible to source up-to-date weather files so that the simulation uses real weather data. At a subtler level, there may be details of building design operation in the working building that may also require adjustment in a calibration process.

With the calibrated model, it is possible to use the simulation to predict details of building operation such as energy use by fans, chillers and boilers as separate subcategories in order to compare the simulation against sub-metered information from the operating building. This comparison can lead to insights as to which systems within the building are performing to expectation versus those that need further tuning and adjustment to achieve their energy efficiency potential.

Case Study: SAHMRI

Figure 3. SAHMRI completed facade showing hood geometry (Photo: Peter Clarke)

Background

Although not commonplace in how buildings are currently designed, the practice of simulation to improve and optimise design using a performance-based approach in the early design stage has had impressive results. When using this approach, it is necessary to prioritise the performance of certain metrics depending on the environment and functionality of the building as a whole, as well as individual space types; as when design teams pursue multiple performance outcomes, sometimes individual optimisation strategies diverge. Figure 4 summarises the metrics and outcomes that can be impacted by the facade alone.

The following section outlines some of the simulations undertaken to optimise the environmental performance of the South Australian Health and Medical Research Institute (SAHMRI) (Cover image and Figure 3). Located in Adelaide's health and biomedical precinct on North Terrace, and opened at the end of 2013, SAHMRI is a 25,000m2 research facility that houses over 600 local and international researchers. Its design team included Woods Bagot, Atelier Ten, Cundall, NDY and Aurecon.

SAHMRI, and in particular its striking facade solution, is a product of performance-based design that balanced aesthetics, cost, and buildability, as well as environmental variables such as solar control, daylight availability, glare and thermal comfort. Key to the success of this solution was developing quantifiable performance metrics up-front to allow robust testing of options between the environmental design consultant, Atelier Ten, and architect, Woods Bagot.

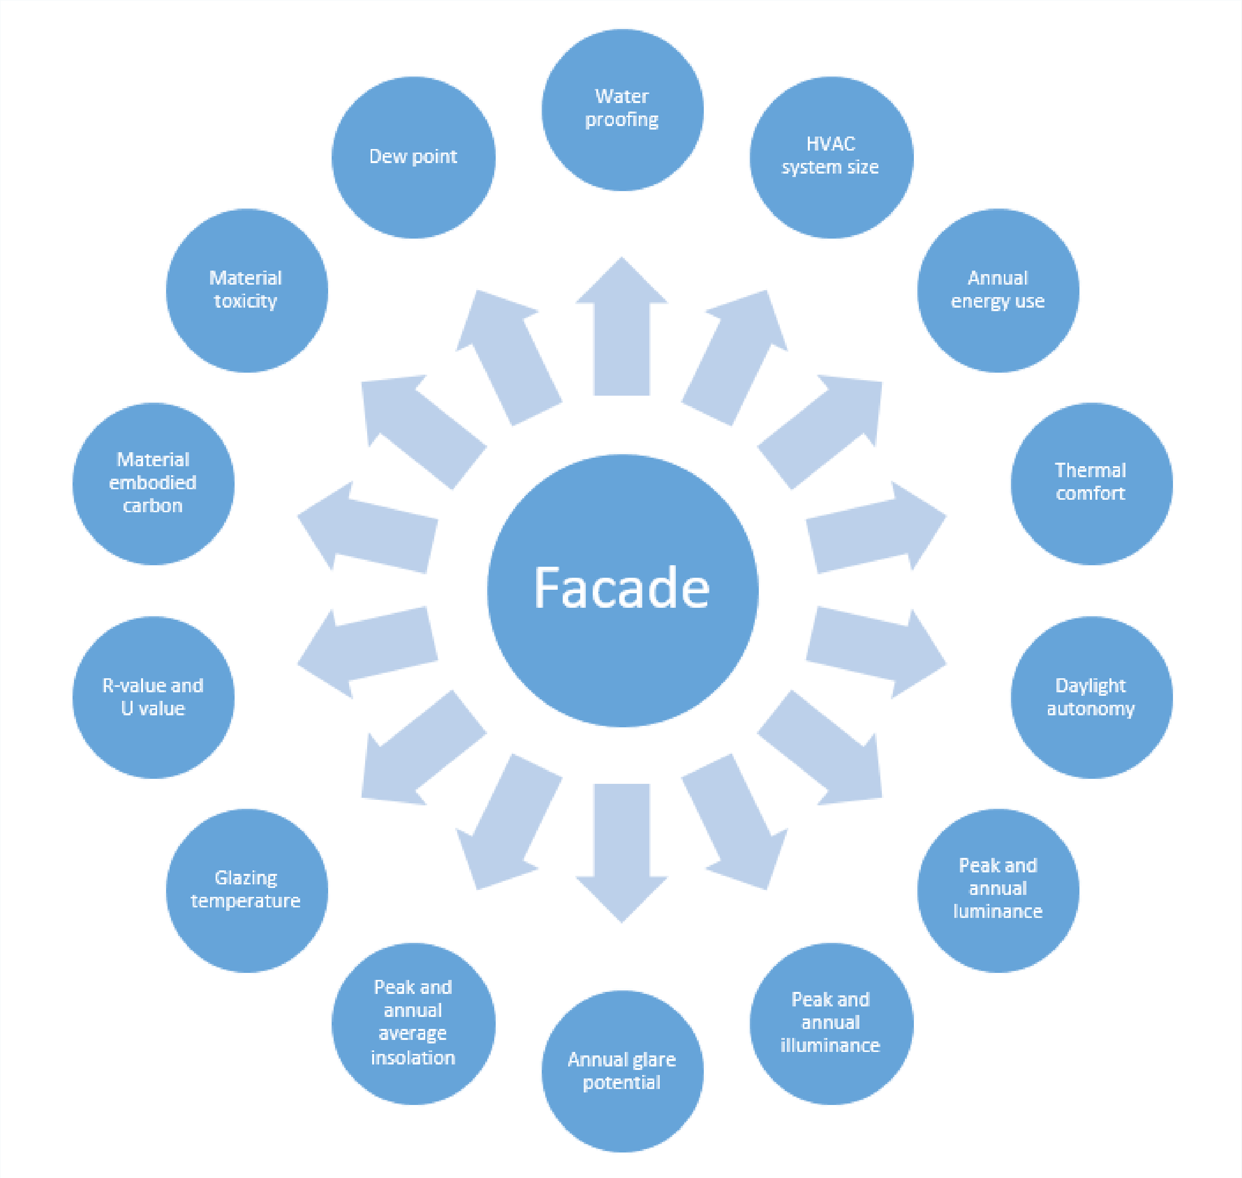

Figure 4. Impact of facade on building performance (Image: Authors)

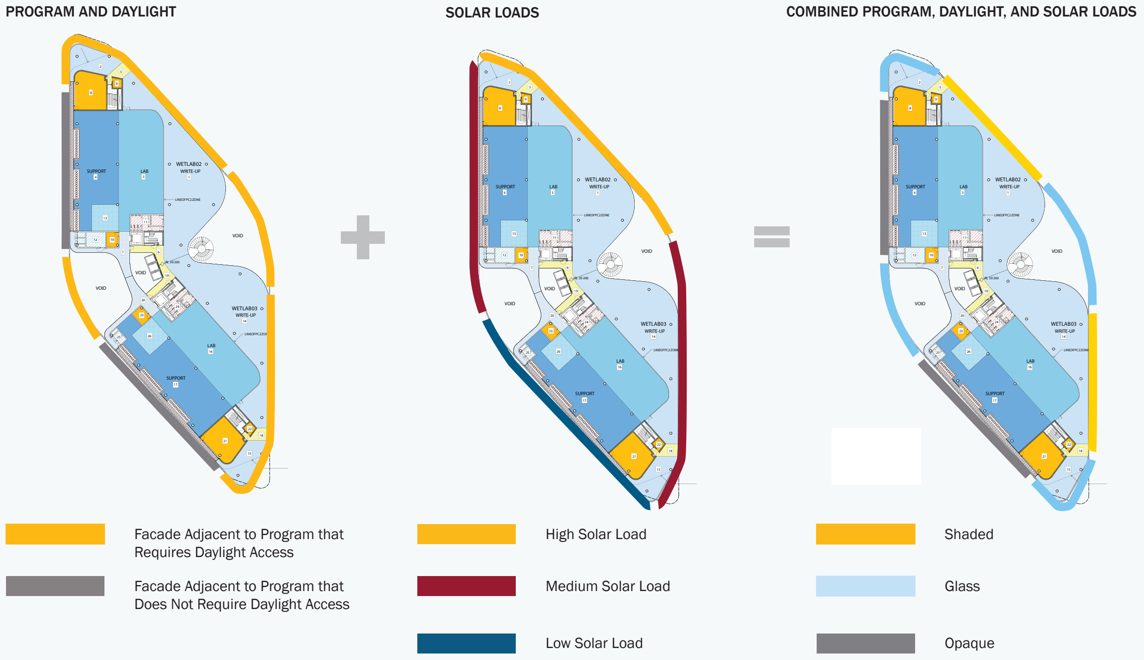

Figure 5. Facade solar loads (Source: Atelier Ten)

Space planning

With the building's overall form established by an approved Development Application, the first step of the facade optimisation process was to assess which space type required access to daylight and which spaces lent themselves to opaque elements. This was mapped with the magnitude of solar load for each facade, as well as other functional relationships. The result (Figure 5) provided the basis of iterating the facade design to its current constructed form.

With the three-dimensionally curved form of the building and significant facade structural spans across multi-storey atria, a triangular grid-shell solution was selected as the most appropriate facade system.

The next challenge was the type of solar control that was incorporated. Regardless of the space type selected to sit behind a particular piece of facade, limiting direct solar gains was crucial in reducing HVAC plant size, energy consumption, risk of glare and improving thermal comfort. A number of options were explored including: the angle of the facade, high performance glazing, fritting and external shade (hoods, louvres or extended mullions) (Figure 6). Variations of these options were also explored with peak solar load, annual solar exposure, solar energy transmitted into the building, as well as assessment of daylight and visual comfort, to understand the performance of each of the options.

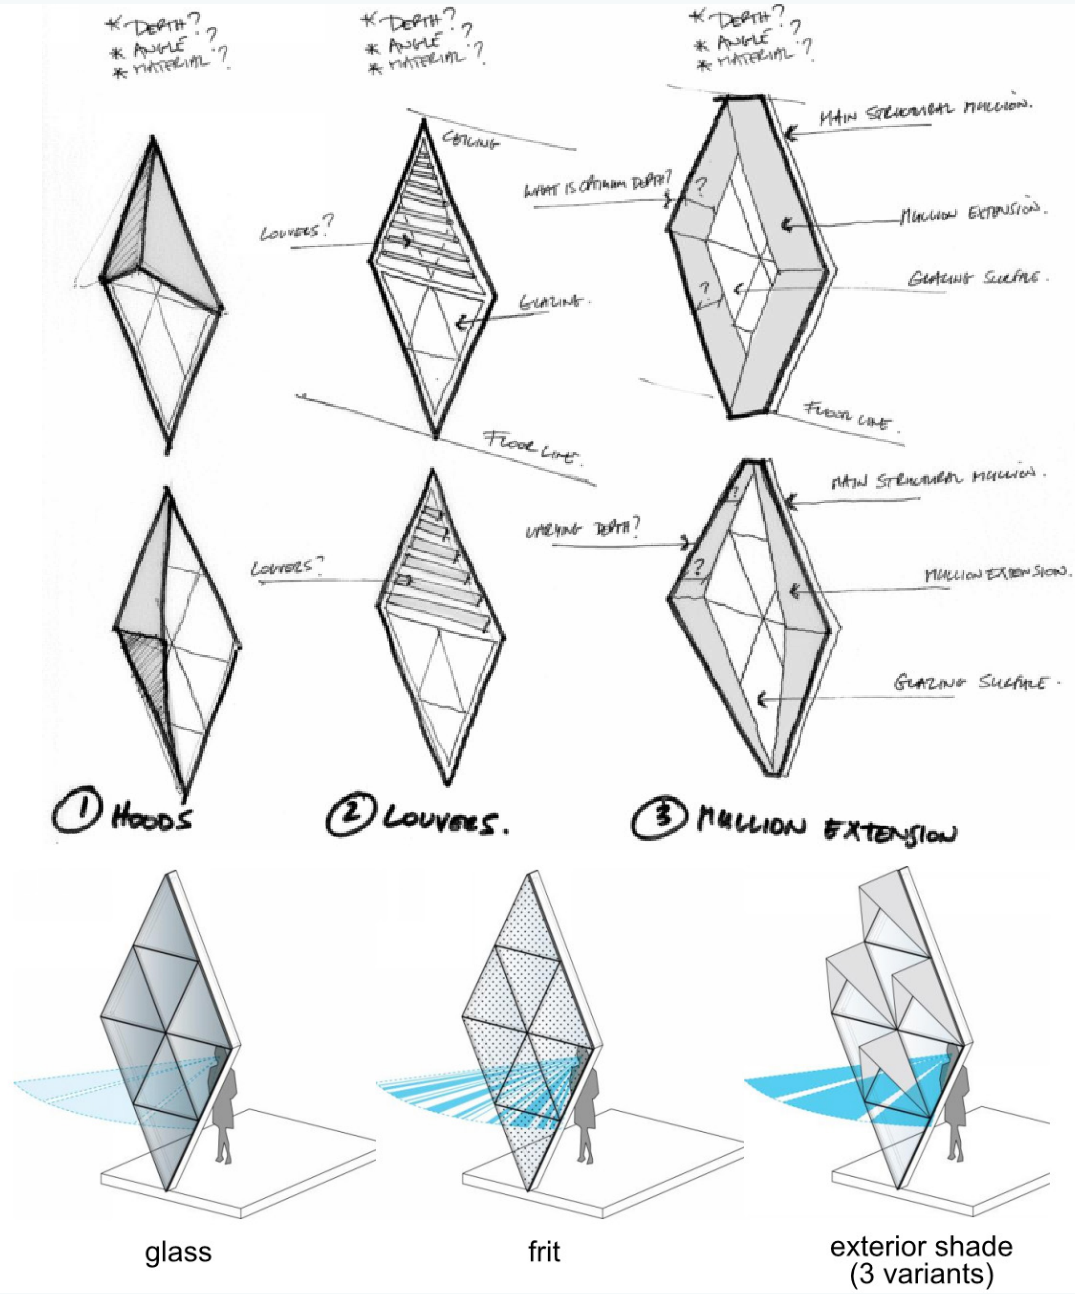

Figure 6. Design options for facade shading, SAHMRI (Source: Atelier Ten)

Energy

With the largest solar gains experienced on the northeast facade, and the least on the southwest facade, one benchmark that was targeted early on was to reduce the northeast facade load by 50% to more closely match the southwest facade. Testing showed that with a 50% frit on the northeast, only a 32% reduction could be realised; however, all nine variations of exterior shading devices were capable of meeting this. Here the value of the simulation was in freeing the designers to select an external shading strategy optimised for cost and buildability rather than purely for energy savings.

The analysis used an hourly annual energy simulation tool, eQUEST, and simulated a pair of typical spaces: a perimeter office and a perimeter laboratory. These spaces were modelled with simplified versions of the shading options and the performance was simulated across a range of orientations (figures 7 and 8). Figure 7 shows how the load varies over the time of day and year (the closer to red, the higher the solar load). This method of abstracting the design problem allowed a rapid response to the emerging design ideas. While it didn't capture the energy consequences of these design options for every possible perimeter space, it captured the majority of the conditions in the building enough to make well-informed decisions about optimising the building facade. Figure 8 shows an example where a 900mm top hood reduced solar load compared to no shading by 52%.

While reducing solar load through passive design principles should always be the first step in facade design, it can be advantageous to understand the contributors to annual energy consumption to ensure the design team focus their efforts on the highest building energy consumers. In SAHMRI's case, due to the function of the offices and laboratory spaces, high lighting and plug loads from night operations were expected. This meant that reducing solar gain had a comparatively small effect, with an 8% reduction in energy intensity in offices and only 5% in labs (Figure 9). As energy intensive spaces, the labs used up to four times the energy per unit area when compared to the office spaces, mostly as a result of equipment. The team knew this would be the case but wanted the broader project team to be aware of the extent of energy savings in advance, avoiding the risk of poor decision making based on lack of performance context information.

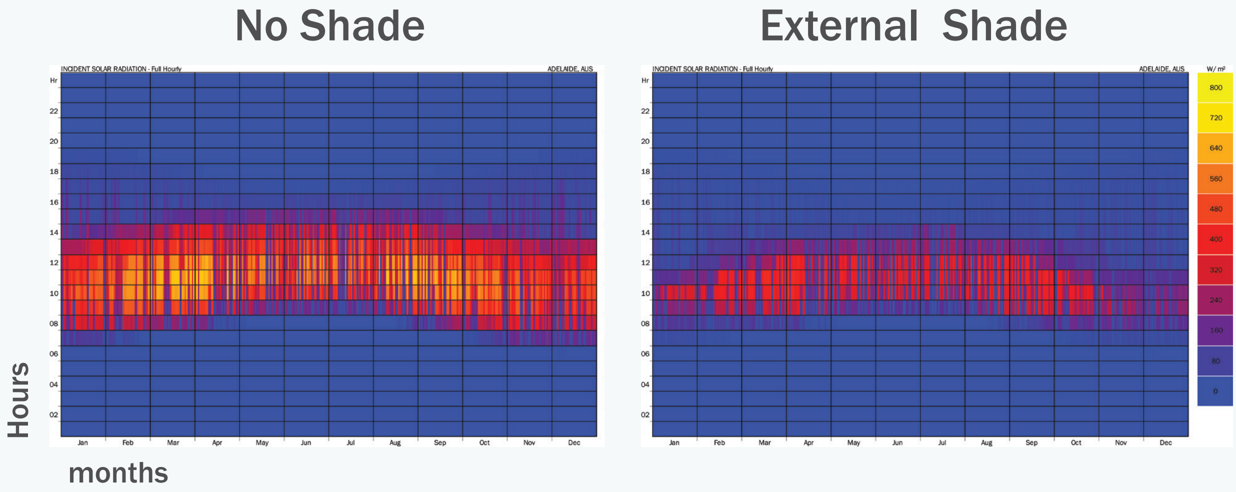

Figure 7. Reduction in external incident solar radiation calculated using DAYSIM (red indicates high solar load) (Source: Atelier Ten)

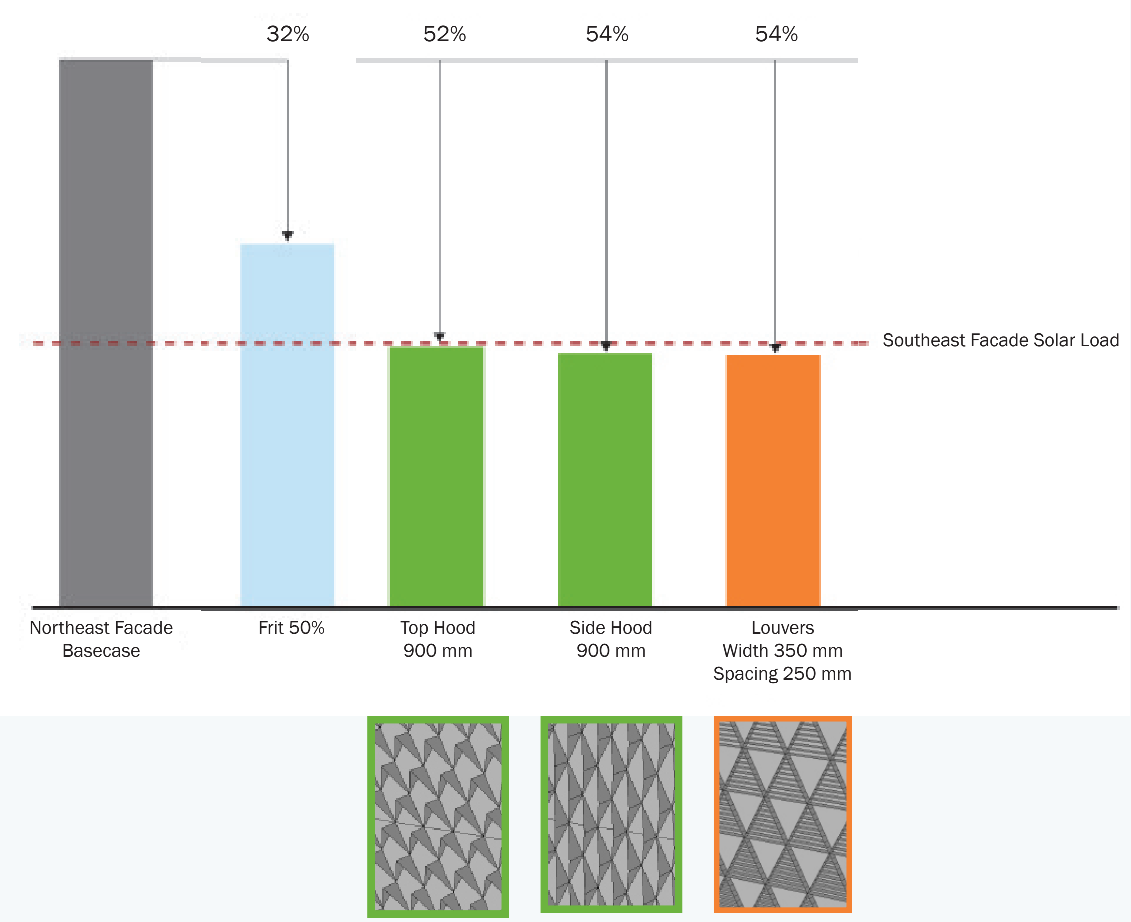

Figure 8. Transmitted solar energy calculated using eQUEST (Source: Atelier Ten)

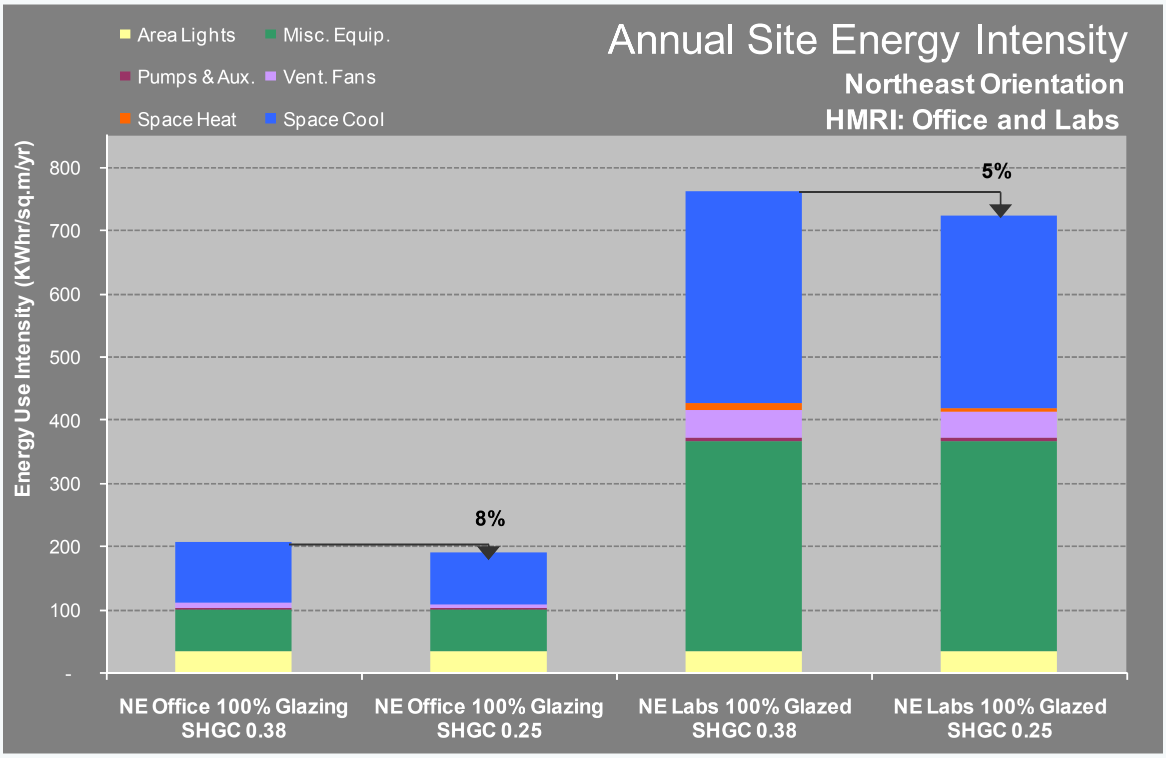

Figure 9. Annual energy intensity reductions across different glazing solar heat gain coefficients (SHGC) calculated using eQUEST (Source: Atelier Ten)

Visual comfort

With the facade significantly affecting thermal comfort, daylight and glare – some of the critical conditions for a healthy workplace – Woods Bagot and Atelier Ten focussed the performance discussion around holistic performance and making an attractive workplace. This factor (among others), now commonly referred to as wellness, was significant to the building occupants, who were competing with other leading university research institutions globally to recruit talented researchers to their program.

With daylight and glare being next on the list of variables to optimise, a 3D model of a typical ten metre open-plan workplace bay was developed, representative of the majority of the building area along the open eastern facade (the west facade is lined with solid-walled service rooms, shielding the building occupants from strong late afternoon, low-angle sun). This model was developed with the intent of optimising the top hood exterior-shading approach, preferred by the architects for different orientations. This exercise took some variants of the extended mullion principle and applied them to the hood dimensions. All daylight analysis used this representative single bay. A daylight model was created as per the process shown in Figure 10 below.

There is a fine line between balancing availability of natural light and minimising the risk of glare from the sun, as often improving one will have a detrimental effect on the other. Traditionally, the daylight factor has been used, however values of 2% and above can create visually uncomfortable high-glare spaces that are too bright and have significant contrasts across the working plane. The 2% daylight factor is used in the Green Star – Design & As Built rating tool, version 1.1. While metrics for measuring the appropriate amount of light required to perform a task are fairly straightforward (but advancing steadily), glare is infinitely trickier as it depends on a person's location as well as their view angle.

A combination of illuminance (daylight) and luminance (glare) simulations were used for SAHMRI's internal spaces. An average annual illuminance of 300 lux was the target for the interior light level, as well as developing a more uniform level of lighting from the facade to the interior. In terms of glare, the facade options were tested using daylight glare probability (DGP), at the time a new visual comfort metric. Values in this metric lower than a threshold of 0.35 equate to imperceptible glare. Values below 0.40 and 0.45 are perceptible and disturbing respectively, while values above 0.45 are considered intolerable. While it might be unrealistic to remove all instances of glare without significantly impacting the function of a window itself, the aim is to minimise the frequency of glare, so as not to cause discomfort more often than not. Since the SAHMRI project, current tools are able to simulate annual average conditions, rather than point-in-time conditions. This allows designers to achieve a balance throughout the year, providing appropriate natural light levels for most of the time that spaces are occupied.

Figure 11 shows the difference in annual average illuminance for a fully unshaded northeast facade compared to a south-facing facade. With the hood design falling somewhere in between these two, the aim was to get as close to the south-facing facade performance as possible by applying different geometrical hood designs to the northeast facade. This was achieved by lowering the tip of the hood by 500mm, with each triangular glazed element having equal sides of roughly one metre long. This solution comes close to approximating the same levels of light experienced on a south-facing facade which receives no direct sunlight.

The option of perforating the top hood and applying a secondary shading device was also investigated; however, this was shown to have minor benefits when compared to the solid hood. Comparison of glare risk was done quantitatively for particular days and times, as well as qualitatively throughout the year, to develop a true understanding of when glare was likely to occur and its potential severity. As can be seen in figure 12, glare due to direct sunlight will always occur during early morning, regardless of the presence of exterior shading; however, this can be significantly reduced in the hours that follow with the hooded designs incorporated on the SAHMRI facade.

In atrium areas towards the centre of the building the hoods are shallow or eliminated to allow more daylight and direct sun into the transitory space that also serves to transfer natural light deeper into the building (Figure 13). Conversely, deeper hoods are used in areas where daylight and glare need more control, providing more uniform conditions between the facade and the interior. This was rationalised from a cost and buildability perspective to 89 different hood shapes spread across nearly 5000 triangular glass panels.

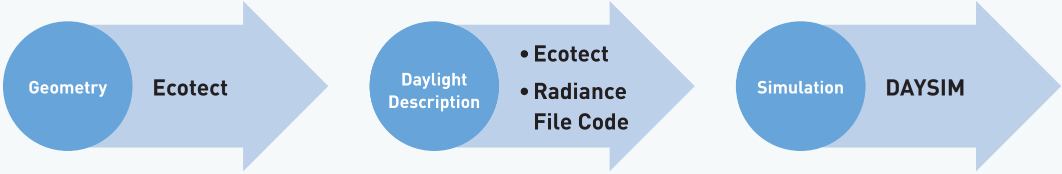

Figure 10. Daylight model creation process (Image: Authors)

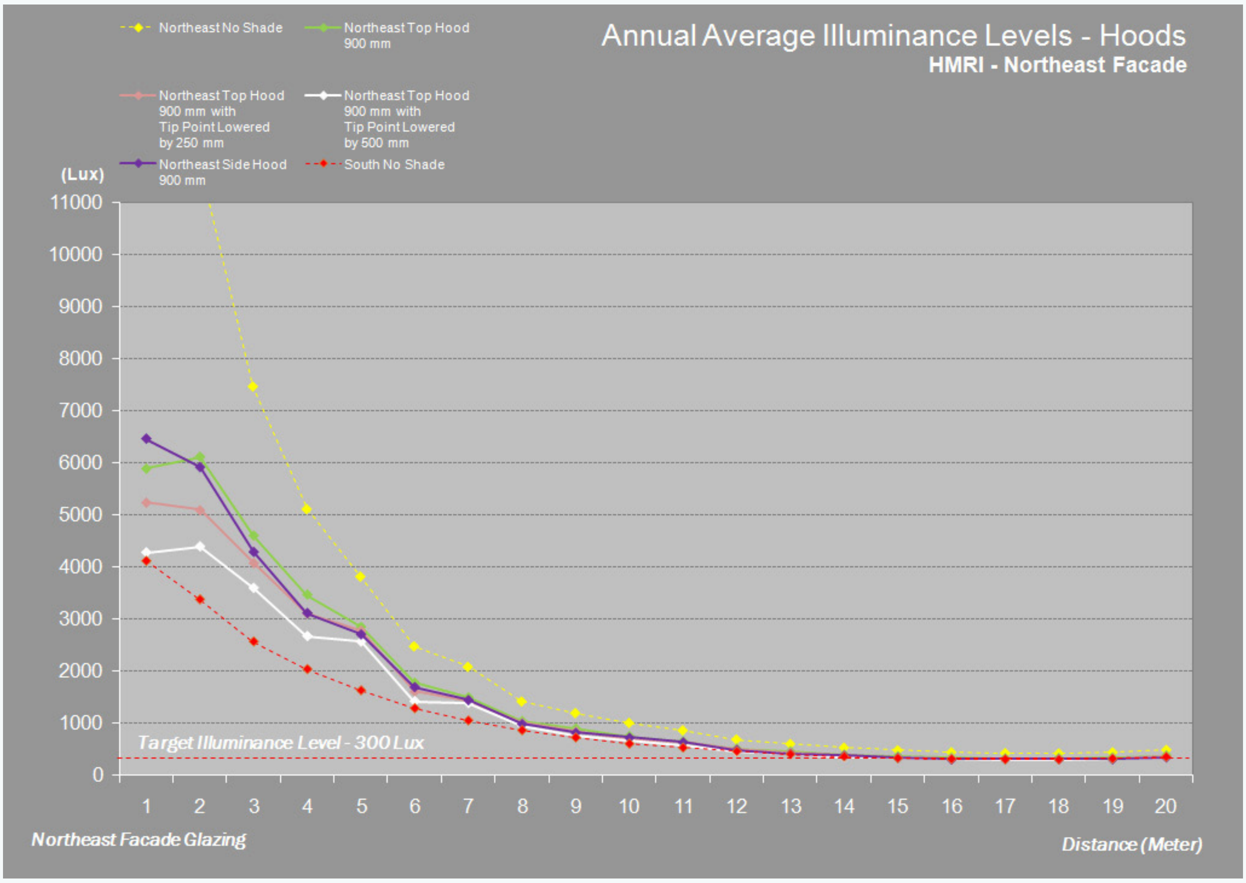

Figure 11. Annual average illuminance levels for different shading hood configurations, calculated using DAYSIM (Source: Atelier Ten)

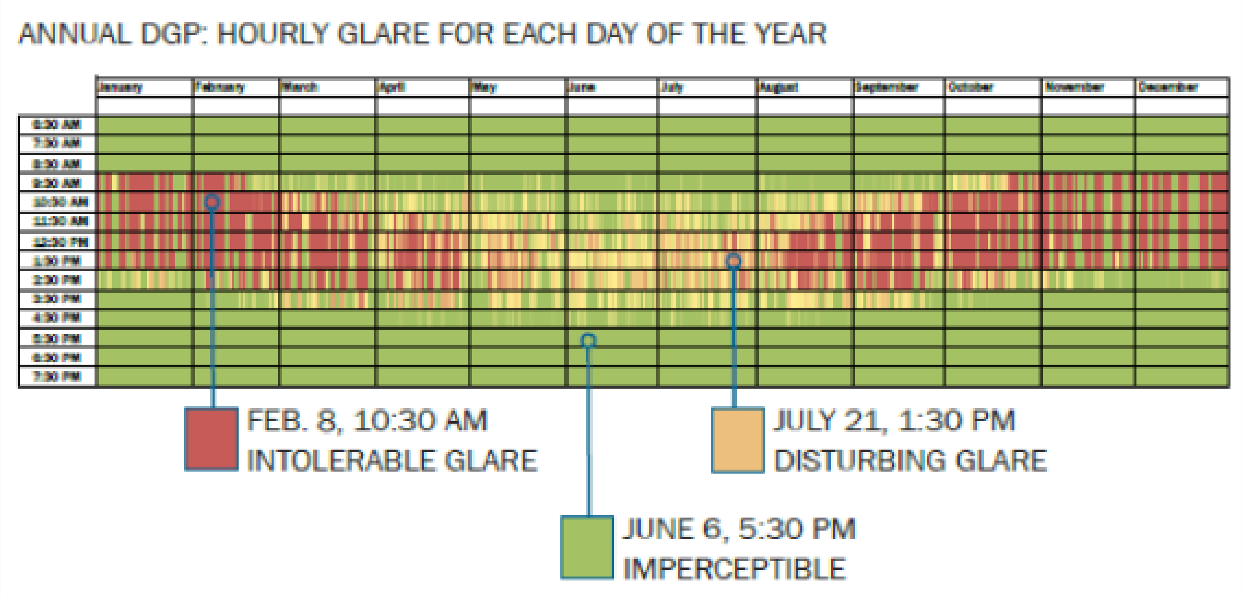
ANNUAL DGP: HOURLY GLARE FOR EACH DAY OF THE YEAR

🟥 FEB. 8, 10:30 AM
INTOLERABLE GLARE

🟧 JULY 21, 1:30 PM
DISTURBING GLARE

🟩 JUNE 6, 5:30 PM
IMPERCEPTIBLE

Figure 12. Annual glare risk assessment, calculated using Radiance (Source: Atelier Ten)

Combining metrics to achieve a successful outcome

Given the number of options to be considered and the sometimes conflicting performance of energy vs daylight vs glare, the balancing of these metrics to reach an optimised solution takes time and can often be intimidating. Add to this the performance on different orientations and the variables multiply even further. This did not stop the SAHMRI team achieving their goal of a truly environmentally responsive facade. Intuitive parametric modelling was used to determine the ideal size, angle and location of the hoods for each space type and orientation. Tools such as Rhino, DIVA and Grasshopper can now be used to efficiently optimise designs across a number of performance variables, which makes the approach used for SAHMRI even more impressive. Although this approach appears to involve a number of different software packages, most have the ability to 'speak' to each other, reducing the need for multiple models to be created.

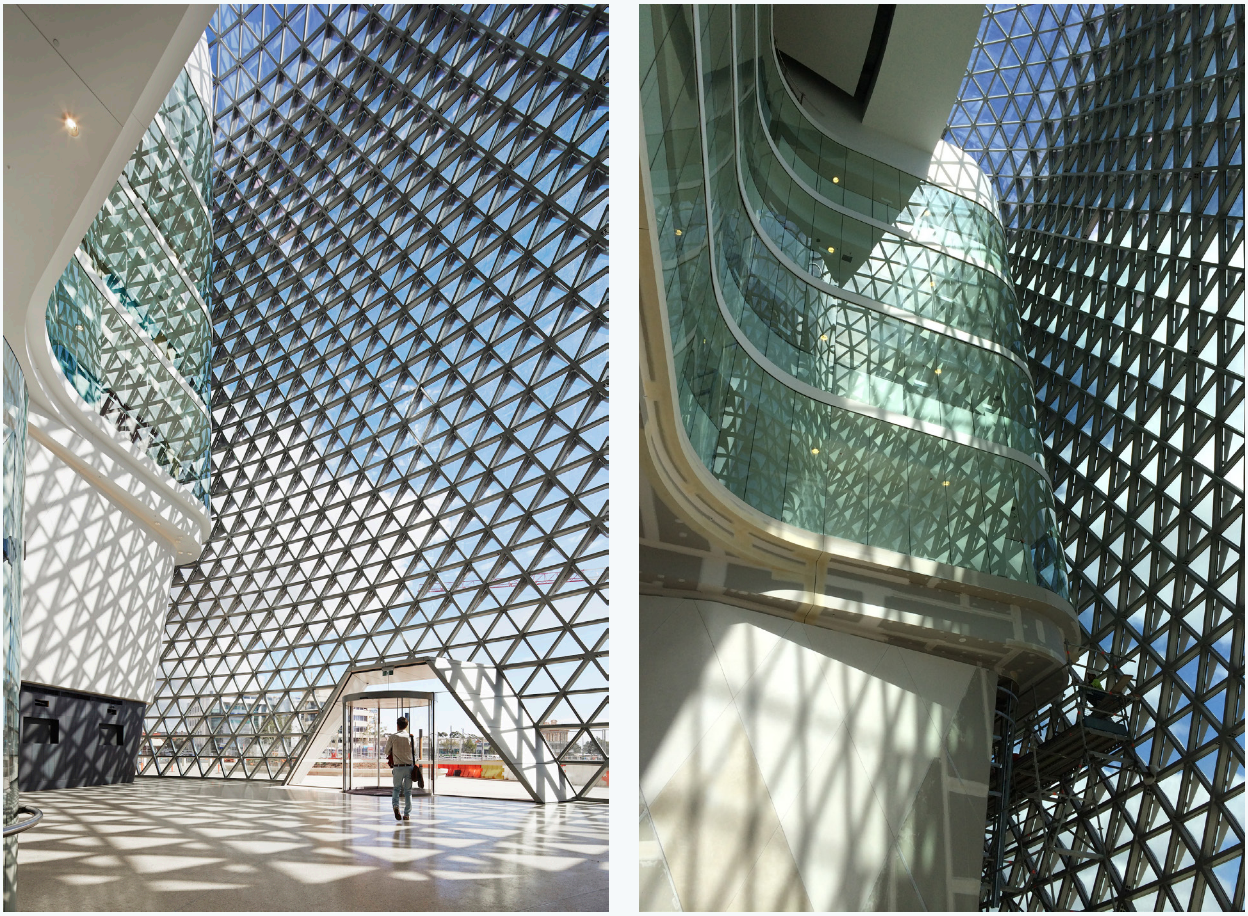

Figure 13. SAHMRI atrium space (Photos: Peter Clarke)

Some words of caution

Quality assurance

The phrase 'garbage in, garbage out' is as relevant to building simulation as to any other form of computing. Every simulation modeller has been caught out by poor assumptions, data entry errors, and other mistakes both innocent and careless. The temptation is also ever present to accept results at face value with the result that these mistakes reticulate through to the final results without question.

While a good deal of quality assurance in simulation is just simple process – checking inputs, documenting assumptions and the like – the other critical component of quality assurance is to experiment with the model across a range of scenarios. It is common to find that when a model is subjected to new conditions it shows up errors in the base case model that need rectification. This means it is particularly important not to treat simulation as a single point process, as is often undertaken for compliance. The more design and operational variants a model is subjected to, the more possible it is to build confidence in the model. This process works well with the use of simulation through the whole design and construction process, as documented above, because the model is updated and repurposed multiple times. Pure compliance simulations have a far higher probability of unseen errors.

Absolute versus relative performance

The traditional role of simulation – outside compliance – has been to inform design options. For this process (and indeed in many compliance processes) the simulation is used as a comparison of the relative performance of options rather than a prediction of absolute performance. Both of these roles have validity but care is required in translating a simulation model between relative and absolute; in some cases this may not even be possible.

An absolute prediction requires calibration to be representative of a building in operation. It follows that a model built for compliance or design option comparison almost certainly does not have the features necessary to predict absolute performance. However, this problem is not always understood by simulators or their clients, leading to unwelcome surprises in post construction when, for instance, a building does not achieve absolute energy figures predicted by a simulation produced for NCC JV3 or Green Star. A building operational performance simulation needs to accurately capture the design, and more importantly, construction of a building, as well as subtleties such as occupancy and operation. When combined with actual weather data, the simulated performance should be within 5% of actual performance, providing an excellent tool to understand what areas can be adjusted to improve performance in real life.

Timing

Time needed to create a simulation model varies widely depending on what is being modelled. Some software tools, especially those designed for use during concept design, can provide peak load and energy consumption results instantaneously once a 3D model has been created or imported; whereas a fully detailed thermal simulation of a detailed design can take much longer. For detailed models, the initial production of the building geometry, services and operation is time consuming and can take several weeks (thus best suited to larger commercial projects that will benefit from detailed modelling). Once the base model is produced, adjusted design and operational scenarios are typically quite easy and quick to produce provided they do not involve major changes to overall building design.

In the early days of simulation, the process of running individual simulations could take more than a day. However with the advent of faster computers and cloud computing, simulation times are rapidly declining and most simulations take only minutes to run; overnight would be considered a long run. This is important as it means that simulators have far more opportunity to refine the simulation to ensure results are robust.

It is worth touching on the potential opportunities for sharing models between architects and simulators, which, over recent years, has been claimed to be a highly effective way of collaborating. Architectural models are often developed with extensive detailing that causes robustness issues within thermal modelling or CFD software packages, and often, the most efficient way to generate the required outputs is for the engineers to create a simplified model. That said, there is without doubt opportunity to refine this process in the future by engaging the environmental design team early in the design stages to truly maximise collaboration and model sharing.

Conclusions

Computer simulation gives design teams the opportunity to test building performance during the design process through to post construction. It is possible to add high level cost implications to these approaches, so that the building has the ability to operate, not only at an environmental advantage, but sometimes financially advantageous as well. Correctly used, the simulation can return many times its cost to the project by enabling smarter decisions that replace industry rule-of-thumb and intuition with detailed, project specific performance and cost information. However, often simulation is used purely for compliance with NCC JV3, Green Star or NABERS; this undervalues the simulation and wastes much of its value potential to the project.

This article has set out an explanation of how simulation can be used at each stage from concept design through to post-construction. Common pitfalls of the simulation process around timing and quality assurance have also been discussed.

Overall, it is concluded that design teams should use simulation as an integral part of the design process to enable better design decisions to be made throughout the entire construction process.

About the Authors

Nicki Parker

Nicki has recently taken up the position as NDY's Sustainability Manager for Brisbane after an 11 year career in building physics and sustainability. Nicki is passionate about using computational analysis techniques to demonstrate the performance of a particular design, and working with design teams to optimise appropriate solutions. In particular, she enjoys working with clients and colleagues alike to expand their knowledge and appreciation of such numerical modelling techniques, while working towards high sustainability aspirations. Nicki is the Vice-President of the Australasian affiliate of the International Building Performance Simulation Association (IBPSA).

Quentin Jackson

Quentin is an Environmental Modelling and Sustainable Design specialist and leads Aurecon's Sustainability team in Queensland. He has a passion for working with diverse teams to deliver projects that achieve practical sustainable outcomes without costing the earth (both environmentally and financially)! He has extensive experience in green ratings and has helped develop many of the tools currently in use in both New Zealand and Australia. He is also currently a Green Star Accredited Professional, Assessor, Expert Reference Group and State Industry Group member for the GBCA. Quentin is an Affiliate Director of the IBPSA.

Dr Paul Bannister

Dr Paul Bannister is Director of Innovation and Sustainability at Energy Action. He is an internationally recognised authority on energy efficiency in commercial buildings and is the primary technical author of the NABERS Energy and Water ratings systems. He is also a Director-at-Large of the IBPSA and President of IBPSA's Australasian affiliate. He has over 100 publications in the field of energy efficiency.

Paul Stoller

Managing director of Atelier Ten's Australia offices, Paul Stoller is recognized for environmental planning and design consulting work on large-scale campus, community and urban building projects. Paul's work in Australia includes major museums in Sydney and Perth, an award-winning research laboratory and a highly sustainable innovation district in Adelaide, a botanic conservatory in Canberra, and the original environmental concepts for Federation Square in Melbourne. Paul also teaches in the School of Architecture at University of Technology Sydney.

Australian
Institute of
Architects

DISCLAIMER