Beyond Simulation: Designing for Uncertainty and Robust Solutions

Sean Hanna
Bartlett School of Grad. Studies
University College London
s.hanna@ucl.ac.uk

Lars Hesselgren
PLP Architecture
LHesselgren@plparchitecture.com

Victor Gonzalez, Ignacio Vargas
Next Limit Technologies
victor@nextlimit.com

Keywords: Design, decision support, interaction techniques, simulation, visualization.

Abstract

Simulation is an increasingly essential tool in the design of our environment, but any model is only as good as the initial assumptions on which it is built. This paper aims to outline some of the limits and potential dangers of reliance on simulation, and suggests how to make our models, and our buildings, more robust with respect to the uncertainty we face in design. It argues that the single analyses provided by most simulations display too precise and too narrow a result to be maximally useful in design, and instead a broader description is required, as might be provided by many differing simulations. Increased computing power now allows this in many areas. Suggestions are made for the further development of simulation tools for design, in that these increased resources should be dedicated not simply to the accuracy of single solutions, but to a bigger picture that takes account of a design's robustness to change, multiple phenomena that cannot be predicted, and the wider range of possible solutions. Methods for doing so, including statistical methods, adaptive modelling, machine learning and pattern recognition algorithms for identifying persistent structures in models, will be identified. We propose a number of avenues for future research and how these fit into design process, particularly in the case of the design of very large buildings.

1. INTRODUCTION

The development of contemporary technologies of simulation has yielded many techniques for deriving results of high quality and precision. As these technologies are predominantly computational, Moore's [1965] law has ensured a steadily increasing speed and precision at an exponential rate that should continue to improve these results into the future, allowing the same tools to apply in greater detail to ever larger projects. But where do we go from here? What are the future directions for research and development in simulation?

Design projects undertaken by architects and engineers are of a scale that is unprecedented in history. Not only do highly programmed buildings such as the international airports of Beijing and Dubai encompass a square kilometre or more of floor area, but the planning of entire new cities, in all their functional complexity, is becoming commonplace, especially in Asia. There are few models for such projects, and little that intuition and experience can hope to contribute, so the need for modelling and simulation in virtually every aspect of the design and planning process has become ever more clear. Moreover, the collaborative teams required to realise these projects are of similarly unprecedented scale, and require effective communication. Their members may be distributed geographically, as well as temporally, and may even change over the duration of the task, making detailed virtual models ever more relevant as a requirement of collaboration.

These factors indicate a greater need for simulation, but also make that simulation far more difficult.
The size and complexity of projects ensures this, as does the obviously disastrous cost of mistakes on such a scale. This position paper outlines current limits or difficulties in the state of the art, then suggests possible solutions and where research efforts should be made.

2. DIFFICULTIES

In practice, there are a number of important limits to what simulation is capable of and how it can be used. Examples are given in this section of a number of current difficulties: resource dependent limits, unknowable design parameters, 'wickedness' of design problems, the process of design in practice, and miscommunication inherent in the use of models.

2.1. Difficult to simulate

For many tasks, the complexity of the situation alone makes simulation exceedingly difficult, either because of the time or resolution required to generate a usable solution.

The Pinnacle, at nearly 300m tall, is designed by Kohn Pedersen Fox to be London's tallest building. In addition to the structural and wind load problems typical of its height, its double-skin of partially overlapping glazing panels introduces additional complexity as it forms a scaled 'snake-skin' of a singly and doubly curved façade. In this case, the effect of the new building on air flow in the area was of concern, particularly the possible impact on pedestrian areas at ground level, which might be adversely effected by winds redirected and amplified by the building's extreme size. In this, a series of computational fluid dynamics simulations were instrumental in guiding the design. Performed in X-Flow by Next Limit, these contributed to a ‘skirt’ near ground level in which the vertical surfaces of the glazed tower flare outward to form a canopy and redirect air flow at the ground (Figure 1).

But the building skin has other requirements at a finer level of detail that present more difficulties. The double-skin design is intended to perform passive cooling and ventilation via the cavity between the two layers (Figure 2). Outside air is allowed to enter each glazing unit through an opening at the bottom, rises as it draws heat from the building, and is drawn out laterally through the vertical slot between overlapping panels. This flow also is relatively straightforward to simulate, however the building consists of 8,500 units, each with a unique shape angle and position with respect to prevailing winds, and as the local flow of air is altered by any change made to overlapping, neighbouring panels, the evaluation of how any particular design behaves requires the modelling and simulation of air flow with respect to 12,000 independent angles of glazing. Optimisation of the position and angle of each panel required many iterations of this, and thus a cost in computation time of approximately two weeks, each time a significant change was made in the building shape.

Such a task is not uncommon, and a time frame of two weeks is acceptable for occasional testing, but hardly “real-time”. The case is particularly noticeable as the overall shape of the building was modelled parametrically using Bentley’s Generative Components, and so could be easily modified in many other respects. The optimisation process thus sits somewhat outside the normal process of negotiating the interdependent systems that make up a complex building. Phenomena of much greater precision were ruled out entirely. What about the noise due to the acoustic effect of air flow on each panel? What about rain? These could not have been simulated accurately enough to be of any real use. Although possible in principle, they are at present beyond feasibility due to another level of magnitude in their complexity.

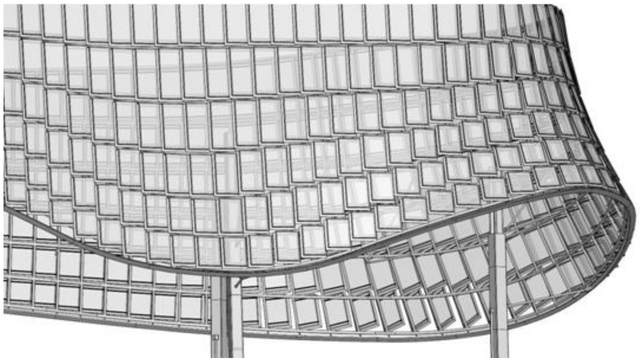

Figure 1. Detail of partially overlapping gazing in the Pinnacle ‘skirt’. Image: Kohn Pedersen Fox.

Figure 2. Exterior air enters the cavity through the opening at the bottom (left). Air rises as it heats up and is drawn to the left, exiting through the vertical slots between overlapping planes (right). Image: Kohn Pedersen Fox.

Figure 3. A single simulation gives precise values for wind velocity, but some regions can be particularly sensitive to initial conditions. Image: Next Limit Technologies.

2.2. Impossible to simulate

Compounding the difficulties above, for some design problems we do not even know the values of all the variables involved. In the case of the Pinnacle it is certain that the state of surrounding buildings will change in future, and possibly affecting the air flow drastically enough to make analyses of the current environment obsolete. For many complex problems, the precise states of relevant variables cannot be measured, or (as is the case with many kinds of human behaviour) there is insufficient knowledge on how to even model the system.

Even in the relatively stable state of an unchanged urban environment, most phenomena to be simulated are continuous and can take on any of an infinite range of real values. In such cases, the probability of simulating the exact values for wind speed, direction or other factors approaches zero (Figure 3). This can often pass without causing problems, but if conditions lie within an instability regime, in which a minor change in the wind causes big differences in performance, then the simulation becomes useless.

In practice, one makes a series of best guesses, and then plans for multiple scenarios. A number of other towers are currently planned for the City of London, and one can use the current state of their designs in a model for the vicinity of the building. But these are only coarse guesses, subject to change, and if the result of the simulation is highly sensitive to initial conditions they may not always suffice. Complex phenomena are dependent on many factors, and for many design problems it is impossible to collect all the relevant data at the outset. Unfortunately, this is often just the type of problem designers face.

2.3. ‘Wicked’ problems

Such difficulties in simulation are made explicit in Rittel and Webber’s [1984] definition of the “wicked problem”, which by its nature resists any kind of clear definition. Unfortunately, design in disciplines such as architecture and planning is described as dealing almost exclusively with such problems. The brief is relatively ill-defined relative to the real range of problem considerations, the perception of the problem itself may change radically as design progresses, and the solution is typically arrived at by a unique process that cannot be predicted in advance.

Rittel and Webber list ten points that describe the nature of this wickedness. For such problems, the problem domain itself cannot be defined. It is perhaps misleading even to consider the design task as dealing with a *problem* at all, in the traditional sense, as this has no “definitive formulation” [ibid., point 1] in terms of boundaries or objectives. Moreover, wicked problems “do not have an enumerable (or an exhaustively describable) set of potential solutions, nor is there a well-described set of permissible operations that may be incorporated into the plan” [ibid., point 6], so no finite set of rules can be considered that might guarantee they are solved. Even with the fastest computer available, the notion of a solution space can simply expand endlessly.

Even given a wealth of potential solutions, the act of testing them is problematic in itself. Solutions to wicked problems have no “immediate [or] ultimate test” [ibid., point 4]. Aside from the inability to define an objective, the unpredictability of the system in question and emergent nature of its behaviour mean that any proposal may generate repercussions, or “waves of consequences” into a future beyond the point at which the test is made. Moreover, “every solution to a wicked problem is a “one-shot” operation; because there is no opportunity to learn by trial-and-error, every attempt counts significantly” [ibid., point 5]. One therefore cannot experiment with the various possible options, trying out, for example, various versions of a motorway or an urban development, because the cost of these is so high, and each is “essentially unique” [ibid. point 7].

The architect may frequently be of the attitude: “once we define it as problem and solution any competent engineer can deal with it”. This is not a comment on the skill level of any member of a particular profession, but merely an observation on the way problems in disciplines are often framed—the classic formulation of engineering problem, or an optimization scenario, is clear. Optimisation might be resource expensive, but it consists simply of defined solution spaces, constraints and performance measures. Unfortunately, these are the wicked part of wicked problems. The bigger and more complex the system, the more we are forced to rely on models to aid in understanding and designing for them, but our certainty about the results of these models also decreases as complexity grows.

2.4. Fit with the design process

The design process is characterised by rapid change, requiring frequent remodelling, re-simulation or re-optimisation. The working relationships of the team both in design and construction also change from project to project, as requirements themselves change, and the structure of communication must be reconstructed to some extent to reflect this.

The features of wicked problems noted above are in stark contrast to the comparatively constrained design domains of, for example, the automotive and aerospace industries—the industries in which optimisation and simulation are used most successfully. These are the industries primarily responsible for the tools (e.g. Abaqus, for finite element analysis and CATIA, for parametric CAD modelling, are both by Dassault, the latter created directly within its aviation division) and currently remain their greatest influence and market. Within these industries, much more is clearly defined and constrained in advance about the design objective, methods of manufacturing and channels of communication. The processes of design and fabrication are usually known, and therefore highly streamlined. A greater proportion of variables are known throughout the process, thereby justifying a greater investment of time and resources to set up a model that is known to be useful at the outset of a new production line. A spoiler on a car, for example, presents a complex aerodynamic situation, but many cars have them and so they are well understood. Architecture has few such spoilers. Where a good deal of systematic, refined and explicit knowledge can be reliably used and traded in these specific domains, the practice is necessarily messy with respect to the built environment.

Attempts have been made to systematise design more generally. In Simon's [1996] proposal for a "science of the artificial", he attempts to rectify the seemingly intuitive and "cookbooky" nature of how it has traditionally been taught and practiced. But such attempts have been opposed, for example by Schön [1983], who criticises this picture of engineering in which problems are well defined and ends are agreed a priori as "Technical Rationality", an essentially Positivist view that is somewhat limited. Schön argues that real design occurs only by an extended practice of re-evaluation and reflection, in which the working definitions of the problems are refined in parallel with their solutions. This cycle of reinterpretation is often observed as an essential feature of design [Snodgrass and Coyne 1999; Lawson 2006] and creativity [Czikszentmihalyi 1988] in practice, and it would seem necessary to support it with different types of tools.

2.5. A false sense of accuracy

The precision of engineering simulations belies the fact that they are ultimately built on statistical measurements of a significant variation of real cases, and this basis may be unknown to the end-user. The simulated behaviour of a single beam of given dimensions is known only because of real tests on many samples of a similar material, with the potential range of structural properties taken into account as a factor of ignorance. But while the models themselves are founded on statistical approximation, this unfortunately doesn't show in the result of simulation. Precise values are given for minimum material tolerances, and these are used in a structural model, along with similarly precise geometry for idealised members, connection details, etc. The result, naturally, is just as precise. For a trained engineer familiar with the factor of ignorance implicit in the original assumptions, the level of real accuracy may be estimated with a little thought, but this can easily be forgotten in practice.

Lawson [2006] gives examples of two types of dangers in how this affects design. The first is that the precision of calculation itself conveys an image of knowledge that does not actually exist. The ease with which computers perform hidden calculation to many decimal places obscures the scientific notion of significant digits, and the polished graphical display can give the uncanny impression of authority. This is typical with students, who "sometimes submit thermal analyses of their buildings […] calculated down to the last watt. Ask them how many kilowatts are lost when a door is left open for a few minutes and they are incapable of answering." [Lawson 2006 p. 70] The unknown quantity of heat loss due to use far outweighs any minor benefits of a few watts here and there, but in this case the designer's ignorance of the major factor in the margin of error is too easily concealed by the clarity of the simulated output. The precision of this output is too often taken to indicate accuracy.

Even if factors contributing to the performance of the design are well known, the second danger is that easily accessible statistics about one factor may influence the designer to emphasise it over more important ones. The kind of precise measurement in Boje's [1971] account of the seconds lost in every opening and closing of an office door, for example, exemplifies a kind of "numerical measuring disease" [Lawson 2006] that might sway a designer toward open plans as they ignore the far more important social and interpersonal factors dependent on spatial separation. This influence may be the result of any of a number of properties inherent to our perception of statistics: they make the factor in question more frequently visible, more explicit and more easily explained to others.

3. POSSIBLE SOLUTIONS & RESEARCH AIMS

In the example of the Pinnacle given above, 8,500 glazing units with 12,000 angles of placement are all unique and require different solutions for optimisation. As a set, however, there are many common features: overall structural hierarchy, materials, but also how they respond to varying wind conditions. There may well be regularities to be found in these that can help in determining the results of the simulation or design optimisation, in the same manner as the statistical regularities assumed in everything from structural capacity of a beam to the variation of annual climate. These regularities are not likely to be statistical, however. In a statistical approach, the variables are defined a priori. Here, and for interesting design tasks, they are unknown at the outset. This section describes directions for research in several areas that may help to deal with the limits and problems of the previous section. Several possible solutions are described: the mapping of broader state spaces, a change in design goals, and 'smarter' modelling techniques.

3.1. Multiple runs of the simulation

The single run of a simulation results in a prediction of behaviour for a precise set of conditions—a specific temperature, humidity, wind direction and velocity, and sunlight, for example, constitute the weather. Designers, however, are usually interested less in this than in climate. Planning for a building that is intended to last for many years, they need a range of varying conditions that must be accommodated, not an instantaneous snapshot of a specific one. In the case of wind, a 'wind rose' captures a range of possible input parameters of wind velocity and direction, as they might be distributed probabilistically for a given location. Designing for this range then means running a series of multiple simulations under differing conditions within this given range, and often weighting any conflicting recommendations according to their likelihood or importance.

Multiple runs are often done in an ad-hoc fashion due to time constraints, or more methodically in the case of sensitivity analysis to test perturbations around a single solution under investigation. With enough such runs of a simulation, however, one might begin to build a more systematic overall picture of the effect of a particular design parameter on the behaviour of the system as a whole: The width of building element A has a non-linear but reliable effect on the wind velocity in zone B, at least when the wind direction is between 150 and 230 degrees; If the building remains constant, the wind direction has a quantifiable non-linear relationship with the load on element C, at least below a threshold velocity D. These relationships, as complex as they may be, constitute the state space of the system—an abstract, high-dimensional space in which each point represents a different version of the design, its environment or boundary conditions. Even if the relevant relationships between variables are difficult to know in advance, they can emerge when sufficiently frequent samples are taken. By mapping this in detail, one can discern a great deal more about the kinds of effects that ranges of design choices will have, and the ranges of conditions within which one may operate. Moreover, in the language of complexity science, this state space will likely contain certain regions of divergence and instability, and others that form basins of attraction. Identifying these, and their limits, would allow one to design for stability over time by mapping a (intuitive or systematic) description of the stability and instability of a given configuration.

Many approaches to multi-objective optimisation, including Pareto optimisation [Deb 2001], take a variation on this approach. In these, many solutions are evaluated to determine a range of possible optima that trade one parameter off against another. The final decision as to which solution is used may be deferred to a later time.

It appears that the knowledge gained by many runs of a simulation can have a direct effect on the designer. While the simulation of how people move through spaces is far more complex than the physical behaviour of inanimate systems, Space Syntax methods of analysis [Spiliopoulou and Penn 1999; Hillier and Shu 2001] have proven reliable in doing so. Part of the reason is the acknowledgement that the prediction is ultimately founded on the cumulative results of a vast number of people—in the simulation of visual agents [Turner 2006], a single agent moving through a building will trace a path that appears unlike that of a normal person, however the total effect of a large number of agents in a virtual model will correlate highly with the movement of real people in the actual space. Designers working with such agents in real time have been observed to change their interaction with the developing plan from one of first person manipulation of elements, to one of engagement with or accommodation of the agents themselves. It appears likely that instead of imagining walking a single path through a building as an aid to design, the view of many simultaneous simulations allows the designer to think more abstractly in terms of the overall behaviour relevant to the building.

To fully exploit this exploration of design spaces by multiple simulation, a fuller understanding is required of complex systems in general, and any specific design domain in particular. From its inception over half a century ago, complexity science has explicitly acknowledged the difference between systems that can be reliably predicted statistically, and those complex systems which cannot [Weaver 1948]. Given that the identification of a regularity in a previously unconsidered set of variables may allow the latter to become predictable, this distinction may not be absolute. The relevant questions for any given domain are just what sort of regularities is it possible to find? Structural systems are generally more stable then fluid dynamics, for example. An understanding of what causes phase changes, and what tools, variables and resolutions are most appropriate to model them, is still relatively unexplored in domains relevant to design.

3.2. Change in goals: robustness

If one is to design a built environment that is robust and sustainable as conditions change, the attempt to predict, or futurology, is less tenable than providing an adaptable infrastructure. Designing for a sustainable future is largely about identifying persistent structures across scales—everything from road networks to floor to ceiling heights—that have been viable and robust in the past, and ensuring they continue. The result should maintain adaptability even when more precise predictions inevitably turn out to be wrong.

A change in how we conceptualise our goals for design to explicitly acknowledge robustness in spite of variation may be required. In specifying an invariant objective, optimisation normally targets single optima which may be unstable to perturbation when apparent project goals change rapidly during design or real-world conditions turn out to be somewhat different from those predicted. Somewhat less optimal plateaus of stability are preferable. Technical research required here overlaps with that of the other suggestions in this section: increased computational power allows exploration of search spaces and multi-objective optimisation (§3.1) and structured approximations (e.g. low resolution models) may be derived by running a truncated optimisation process during early stages of design, to be completed in detail later (§§3.3 & 3.4).

3.3. Increased speed and smarter models

Performing optimisation (as in §2.1) or reliably mapping a state space (§3.1) require numbers of simulations of progressively higher orders of magnitude, and making these multiple runs feasible requires faster simulations. These are guaranteed by current trends in the increasing availability of computing power: its cost will continue to decrease (Moore's law [Moore 1965]), and the adoption of grid and cloud (internet based) computing will make better use of it by sharing otherwise dormant resources. But these are only incremental improvements, and a step change from single simulations to an overview of a complex state space requires a vastly larger number of simulations. Moreover, it is likely that constantly growing projects, increased pressure on project timelines and new demands for detail will negate much of this benefit. This is even without the possible counter effects of increased demands from software known as "Wirth's law" [Wirth 1995]. In addition to better hardware, the step change may be affected by the development of smarter models for use in simulation.

In the simplest case these might be based on statistical approximations—low resolution working models, for example. Mesh sizing for finite element analyses takes such an approach in attempting to use the largest element dimensions possible to reliably capture relevant details, thereby increasing resolution in some zones and decreasing it elsewhere [Langham and Grant 1999]. In mapping a large state space of multiple simulations, the basins of attraction or regions of greatest sensitivity to initial conditions may be the same across a broad range of resolutions, and a far lower resolution may be used in some areas.

In more complex cases, patterns or otherwise hidden correlations in data might be found via more advanced statistical techniques or machine learning algorithms. Hanna [2007] uses such an approach in the optimisation of cellular structures consisting of many thousands to millions of unique cells. The modular nature of the individual units provides enough regularity to allow a function to be derived that can replace the optimisation and simulation entirely. A support vector machine is used to map local stress to an optimal cellular structure by training on data taken from between 100 and 600 previously optimised samples. The result can actually improve performance (verified in simulation) over traditionally optimised versions, and increase speed in the order of tens of thousands of times faster. The patterns derived from large data sets may extend to much less clearly defined properties as well. Similar machine learning methods have been used to extract and manipulate arbitrary patterns from spatial arrangements of desks in the workplace [Hanna 2007a] buildings [Laskari et al. 2008] or entire cities. In the latter case [Hanna 2009] the geographical location of a city has been shown to correlate with a number of properties of its form that are non-discursive in the sense that they are not easily spotted or described explicitly by a human observer, but the computer can derive them from plan data and thereby classify cities as to their location with a significant degree of accuracy.

Research into machine learning in general is necessary here, in addition to more domain specific investigation. In many cases, reliable predictions are possible because of the underlying stability of a different part of the system. The consistent patterns predicted by Space Syntax of social interaction [Spiliopoulou and Penn 1999] and crime [Hillier and Shu 2001] in addition to human movement, for example, are due to the relative stability of building layout or street networks over time. Identifying how various subsystems are related for any particular domain will help to identify strategies for design.

3.4. Adaptable, flexible methods of modelling

A design process in which change is frequent (§2.4) and problems are ill-defined (§2.3) would benefit from modelling and simulation methods that can easily and quickly adapt to new problem definitions.

To increase the speed of simulation when ideal levels of resolution are not known in advance, modelling techniques that readily allow changes in resolution may be used. Particle systems may be preferred to meshes of fixed topology, for example, because although the latter allow variation in detail at crucial edges, they are not easy to change over time. In a simulation of air flow over an automobile, Next Limit's X-Flow dynamically updates the number of particles in the system depending on the volume of turbulent air (Figure 4). The process uses a maximum of 50 million particles, but begins with only 5 million for a substantial saving of computation time. This flexibility is potentially more valuable if varying the resolution of particles also becomes possible, with particle (and therefore computation) density increasing over time in turbulent zones where detail is greater and decreasing where it is not needed. This is now being considered for development in the future. The principle can be extended in n-dimensions to the resolution of sampling of state spaces as described above (§3.1).

To fit with the design process and ill-defined problems, there are several ways in which it is possible to use and re-use intermediate results throughout the design process. A typical example of how this is frequently done already is the overall surface shape of a large roof, which might be optimised at a low resolution, while finer details such as the space frame modules that form the actual structure, the dimensions of the structural grid and even the local interruption of the structure by cuts for services might change frequently thereafter. Because their effects on the partial solution are local or non-existent, there is no need to revisit the initial optimisation. The development of a repertoire of such partial solutions is frequently employed in practice.

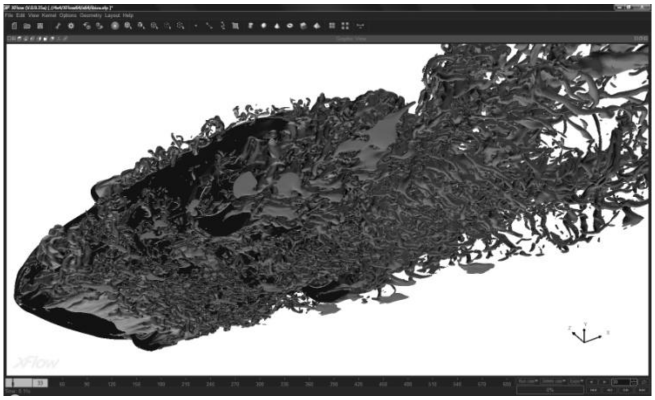

Figure 4. X-Flow dynamically updates the number of particles in the system depending on the volume of turbulent air. Image: Next Limit Technologies.

Figure 5. Multilight renders each light source separately so that they can be adjusted independently after rendering is complete. Different versions of the same scene can be produced without re-rendering. Image: Next Limit Technologies.

Simulation tools can accommodate these. In Maxwell render, Next Limit have developed the 'Multilight', in which each light source is rendered and separately as a partial solution, then to be mixed afterward in the final image. The user is then able to adjust the exact mix of light sources after the fact, in real time, to produce different versions of the scene (Figure 5). This results in increased ease of use for the user, as it does not require decisions about final light to be made early on, but allows a reflective process by which the effect of the light can be seen directly and immediately as the decisions are made later. It also allows far better communication with a client or among design teams. While renderings are often seen as somewhat final, they are crucial to the collective creative process, and such tools encourage engagement.

Basic research is still required. Multilight is possible because light is easily separable, but many design variables are not. The roof example above is strictly hierarchical in that decisions of detail design may be dependent on overall shape, but not vice-versa. Finding the points at which partial solutions to more complex models may be stable enough for re-use is another task for the investigation of state spaces and machine learning research mentioned above (§§3.1 & 3.3). Technical development will then be required in the development of tools that allow the use of partial solutions. This use of simulation and models at intermediate stages, without clearly defined start conditions or a fixed end solution, is intimately related to the way designers work, and ultimately, some re-education of designers themselves may also be required.

4. CONCLUSION

As designers take on larger and more complex tasks, this paper has suggested that the single analyses currently provided by most simulations display too precise and too narrow a result to be maximally useful in design, and instead a broader description of what simulation can do, and how it can be used, is required. It has attempted to outline ways to make our simulations, and our buildings, more robust with respect to the uncertainty we face in design: a better exploration of the range of solutions, changes in how we perceive optimisation goals, and the use of statistical and machine learning algorithms to do so.

In planning efforts for future research, there is certainly the tendency to refine the simulation methods we have to ever finer degrees of precision. This is helpful, but only in context of understanding the real needs of the design tasks. None of the methods in question are in any danger of becoming "technologies looking for an application", but this paper has aimed to present the ways in which they need to be used, in the hope that newly developing technologies will improve designers' understanding of uncertainty and robustness and better equip them to deal with the complex, ill-defined problems they face.

References

Boje A (1971) *Open-plan Offices*. London, Business Books.

Czikszentmihalyi, M (1988) Society, culture, and person: a systems view of creativity, *in* Sternberg, RJ (ed), *The nature of creativity: Contemporary psychological perspectives*, Cambridge University Press, Cambridge, pp. 325-339.

Deb K (2001) *Multi-objective optimization using evolutionary algorithms*, John Wiley and Sons.

Hanna S. (2007) Inductive machine learning of optimal modular structures: Estimating solutions using support vector machines, *Artificial Intelligence for Engineering Design, Analysis and Manufacturing*, 21: 1-16.

Hanna S (2007a) Automated Representation of Style by Feature Space Archetypes: Distinguishing Spatial Styles from Generative Rules, *International Journal of Architectural Computing.*

Hanna S (2009) Spectral Comparison of Large Urban Graphs, In Koch D, Markus L and Steen J, *SSS7: Proceedings of the 7th International Space Syntax Symposium*, Stockholm.

Hillier B and Shu S (2001) Crime and urban layout: The need for evidence, in *Ballintyne S, Pease, K and McLaren V: Secure Foundations: Key Issues in Crime Prevention and Community Safety*. IPPR, London.

Langham AE and Grant PW (1999) Evolving rules for a self-organizing finite element mesh generation algorithm. In *Proc. 1999 Congress on Evolutionary Computation* , pp. 161-168, Washington D.C., USA, July 1999. IEEE Computer Science.

Laskari S, Hanna S and Derix C (2008) Urban identity through quantifiable spatial attributes: Coherence and dispersion of local identity through the automated comparative analysis of building block plans, in Gero JS and Goel A (eds) *Design Computing and Cognition '08,* Springer.

Lawson B (2006) *How Designers Think: The design process demystified, Fourth Edition*. Elsevier, Oxford.

Moore GE (1965) Cramming more components onto integrated circuits, *Electronics* 38(8).

Rittel HWJ and Webber MM (1984) Planning problems are wicked problems. In Cross N (ed.) *Developments in Design Methodology*, John Wiley and Sons.

Schön D A (1983) *The Reflective Practitioner: How Professionals Think in Action*. Basic Books, New York.

Simon HA (1996) *The Sciences of the Artificial.* 3rd ed. MIT Press, Cambridge, MA.

Snodgrass AB and Coyne RD (1997) Is designing hermeneutical? *Architectural Theory Review*, 2(1): 65–97.

Spiliopoulou G and Penn A (1999) Organisations as Multi-Layered Networks: face to face, e-mail and telephone interaction in the workplace, *2nd Intl. Space Syntax Symposium*: 1-24.

Turner A (2006), First steps in evolution of animats for pedestrian simulation. Presented at *Design Computing Cognition '06, Workshop on Models of Machine Learning in Creative Design*, Eindhoven.

Weaver W (1948) Science and Complexity, *American Scientist*, 36: 536.

Wirth N (1995) A Plea for Lean Software, *Computer*, 28(2): 64-68.