11. Exploring Immersive Analytics for Built Environments

Tom Chandler[1], Thomas Morgan[1], and Torsten Wolfgang Kuhlen[3]

[1] Monash University
tom.chandler@monash.edu
tom.morgan@monash.edu
[2] RWTH Aachen University
kuhlen@vr.rwth-aachen.de

Abstract. This chapter overviews the application of immersive analytics to simulations of built environments through three distinct case studies. The first case study examines an immersive analytics approach based upon the concept of "Virtual Production Intelligence" for virtual prototyping tools throughout the planning phase of complete production sites. The second study addresses the 3D simulation of an extensive urban area (191 square kilometres) and the attendant immersive analytic considerations in an interactive model of a sustainable city. The third study reviews how immersive analytic overlays have been applied for virtual heritage in the reconstruction and crowd simulation of the medieval Cambodian temple complex of Angkor Wat.

Keywords: immersive analytics, crowd simulation, interactive simulations, immersive 3D visualisation, augmented reality, tangible interfaces, CAVE

11.1. Introduction and Overview

In simulations of the built environment, the 3D format is much more than a graphical gimmick. It is crucial to discern essential information such as line-of-sight issues, material flows and intersections of objects, all which are best remedied by a full 3D, dynamic reconstruction. Such spaces, whether mediated through a laptop screen or in a specially constructed CAVE environment, are spatially immersive because they depict environments–a factory, a cityscape, a reconstructed cultural landscape–in a realistic way. When navigating through these environments, the user has the benefit of familiar and proportionate visual cues such as walls, streets and computer-generated characters that provide the virtual space with perspective and relative position. While analytic features can overlay and augment these ready-made immersive environments, the organization of 3D models into sets or layers can also function as visual analytic features in their own right. For example, 3D model sets of architectural typologies, vegetation assemblies or walking figures can be interchanged in real time, revealing contrasts, patterns and interdependencies in the data that weren't previously apparent.

K. Marriott et al. (Eds.): Immersive Analytics, LNCS 11190, pp. 331–357, 2018.
https://doi.org/10.1007/978-3-030-01388-2_11

The utility of 3D models have been familiar to architects, engineers and urban planners for some time, but new Immersive Analytic interfaces, augmentations and overlays promise to significantly extend the interpretation and exploration of virtual environments. This chapter explores the application of Immersive Analytic features in simulations of built environments in three case studies that are situated variously in the present, the future and the past. The first case study examines an immersive analytics approach based upon the concept of 'Virtual Production Intelligence' for virtual prototyping tools throughout the planning phase of complete production sites. The second study models the addresses the 3D simulation of an extensive urban area (191 square kilometres) and the attendant immersive analytic considerations in an interactive model of a sustainable city of the near future. Notably, both of these case studies propose models of present built environments as a foundation for the prototyping of future scenarios. The third study takes a landscape of the present day–the world heritage listed Cambodian temple complex of Angkor Wat–almost a millennium back into the past, and reviews how immersive analytic overlays have been applied for the reconstruction of the complex and the simulation of crowds over a 24-hour cycle in medieval times.

11.2. Case Study 1: Planning and Visualisation in Production Engineering

In high-wage countries, the complexity of production processes has increased dramatically. It has been state-of-the-art for quite some years now to make use of virtual prototyping tools throughout the planning phase of complete production sites. In such virtual production systems, a tremendous number of process parameters have to be determined with the goal of making high-quality products. While these process parameters are highly interdependent, the knowledge about how to optimize production processes is dispersed among various experts from different fields working on diverse aspects of designing and optimizing procedures, which makes production planning a highly collaborative and interdisciplinary task. Last but not least, the data generated in virtual production environments is highly heterogeneous.

To meet the challenges above, a first attempt has been made just recently to follow an immersive analytics approach that is based on the concept of "Virtual Production Intelligence" (VPI), which in turn builds on the concept of business intelligence in terms of data aggregation, data condensation, and data exploitation. The VPI strives at an integrative, collaborative decision support for complex production planning processes in different design domains. Factory and Machine are two design domains, where virtual reality (VR), scientific visualization (SciVis) and information visualization (InfoVis) are used in combination to support planners in evaluating and reviewing layouts before implementing them in real factories and machines.

The work presented here is a Research Area in the Cluster of Excellence "Integrative Production Technology for High-Wage Countries". The Cluster is a major long-term project started in 2006, in which over 25 research departments of RWTH Aachen University, Germany, are collaborating with industrial companies across disciplines on new technologies in multiple fields of production, like virtualization, individualized production, hybrid production systems, and self-optimizing production. Within the VPI Research Area, the following RWTH institutes work closely together: The Institute of Information Management in Mechanical Engineering, the Visual Computing Institute, the Department of Factory Planning from the Laboratory for Machine Tools and Production Engineering, and the Nonlinear Dynamics of Laser Manufacturing Processes Instruction and Research Department.

The following paragraphs do not present original research by themselves, but instead are mostly an excerpt from various, previously published outcomes of the project [19], focussing particularly on the immersive analytics aspects of the VPI.

In the field of production engineering, immersive analytics should reduce and accelerate planning efforts and increase planning efficiency by providing an integrative analysis in immersive virtual environments. VPI mainly follows this approach and integrates the planning and optimization of the two design domains Factory and Machine. The heart of the VPI platform is an analytics engine providing, among other things, domain-specific data exploration, advanced data mining techniques, modules for correlation and sensitivity analyses, and a consolidation of planning data in the involved design domains based on domain ontologies [5,19]. The VPI analytics engine is operated via a Web-based platform but can, on top of that, be coupled to an immersive virtual environment, making the VPI an immersive analytics tool. The VPI offers a seamless explorative space by linking and integrating geometric, scientific as well as information visualization.

Design domain factory In the context of factory layout planning, VR-based support tools have already been realized by others, demonstrating that a visualization and interaction in immersive environments provides a significant added-value as compared to non-immersive, monitor-based systems (see e. g. [6, 9]). One of the strengths of walkthroughs in immersive VR is that planners can experience a factory in its original scale. This can, e. g., be used to manually check workplace visibility in a rather natural fashion and thereby ensure important lines of sight between related workplaces. While the existing tools allow for virtual walkthroughs of factory models and the modification of the factory layout from within the virtual environment by placing or rearranging shop-floor equipment, the VPI tool goes beyond this by leveraging visual analytics and providing a much higher level of integration of heterogeneous planning data into the model. With the VPI, geometrical data of the factory layout is combined with visualizations that offer access to additional relevant planning data, with the goal to further increase the usefulness of VR-based factory design. By this, the VPI wants to make a fundamental contribution to the realization of the vision of the digital factory. The overall VPI solution hypothesis is that a combination of semantical enrichment of planning data by a comprehensive information model on the one hand, and smart human-machine interaction via 3D, immersive user interfaces on the other hand, is advantageous for factory planning and the optimization process. Solutions of simulations of factory or production processes are linked together in such a way that interdependencies in the planning tasks can be identified and analyzed flexibly.

Within the VPI project, we developed a prototype of an immersive analytics factory planning platform called the *Factory Layout Planning Assistant (flapAssist)*. A key goal of flapAssist is to not only create an appealing rendering of factory layouts but to elaborate methods and tools that support the factory planners in their daily work. flapAssist offers traditional functionality, e. g., virtual walkthroughs (see Figure 1) and positioning of machines, but combines them with visualization concepts that have been newly developed to better support the planning and optimization process. Above all, material flows provided by the VPI analytics engine are a crucial type of planning data. In flapAssist, material flows are presented as a 3D overlay embedded into the immersive representation of factory models. As shown in Figure 2, inter-machine material flows are visualized via color-coded arcs. Alternatively, a card-style visualization can be used to visualize the accumulated material flow in which specific machines participate. Since material flow matrices can become rather complex, leading to a cluttered, confusing visualization, planners have the possibility to interactively filter out the flow information of those machines that are not in the focus of interest for the current task. Beyond the visualization of material flow data, flapAssist offers a support tool to aid the planner to optimize the position of individual machines with respect to the material flow matrix. Since material flow costs are not the only criterion affecting the factory layout, a fully automated approach is not effective here. Instead, the tool simply suggests where to move a user selected machine, taking into account the entire material flow for that machine. Via online communication with the VPI analytics engine, planners get immediate feedback on their changes and can thus quickly iterate between different optimization scenarios (see Figure 3).

Design domain machine To optimize configurations of complex machines, an analysis of high-dimensional parameter spaces is often necessary. As an example, in laser cutting machine tools typical parameters are laser power, laser pulse duration, beam radius, beam thickness level, Rayleigh length, and beam focal position. To identify optimal machine configurations for a specific production task, it is not feasible to simulate a machine for all possible parameter combinations in a straight-forward way. Instead, the process running on such a machine is simulated for selected combinations of parameters only, while other configurations are interpolated from these selected combinations via metamodels. For the design domain Machine, therefore, we identified metamodeling techniques as being crucial for the VPI approach. For an optimization of machine parameters based on metamodels, it is important to analyse correlations between parameters. Thus, the concept of linked multiple views seems appropriate to explore the data. We, therefore, developed a new visual analytics application called *Metamodel Slicer (memoSlice)*, which mainly builds on an interactive exploration of metamodels via linked multiple views. While a detailed introduction to metamodels in production

Fig. 1: Immersive virtual reality support for factory layout planning.

engineering would go far beyond the scope of this section (for details see e. g., [1]), the visualization concept of memoSlice shall be briefly described here (see figure 4 and [16]).

As required by domain experts, an adequate workflow should allow for an integrated overview as well as in-depth analysis of metamodels. Following this workflow requisite, memoSlice offers three main views of the data. Via the scatterplot matrix view, planners can gain a fast outline of the data. It contains one scatterplot for every possible pair of parameter-criterion combination and thus gives an overview of the distribution of values.

For a detailed examination, the hyperslice view [36] shows a matrix of 2D plots which display all possible axis-aligned slices through one focal point of the data space. In the matrix, the graphs shown on the diagonal represent axis-aligned 1D slices through the common point, allowing for a quick navigation through the parameter space.

In contrast to the HyperSlice view, the 3D view visualizes a single three-dimensional slice instead of several two-dimensional ones. The 3D view contains three crossed 2D-slices whose intersection point matches the projection of the focal point. Additionally, a direct volume rendering of a 3D-slice is superimposed. Thus, the 3D view does not give an overview over all dimensions but instead shows all information along three visible parameters, thereby providing more context than 2D slices.

Integrative analysis of the design domains Factory and Machine While a planning tool for the design domain Factory, like flapAssist, quite obviously can profit from an immersive analytics approach, since a lot of geometrical data is involved, an application like memoSlice for planning and optimization of the design domain Machine follows a more “conventional” visual analytics approach, where the

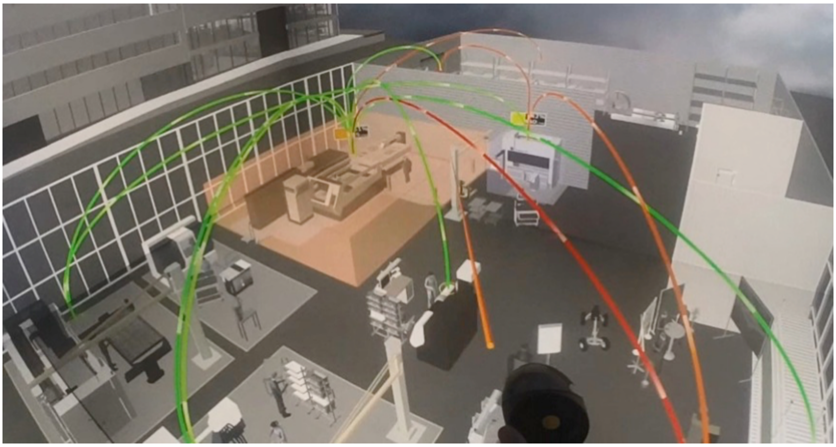

Fig. 2: Visualization of material flow.

aspect of immersion, if at all, plays a minor role. In fact, in the concrete example of memoSlice, the only 3D aspect that might profit from a stereoscopic, egocentric perspective, is related to the direct volume rendering in the 3D view, whereas the other views are considered as classical InfoVis visualizations, where a desktop environment might be more appropriate.

Having the ultimate goal in mind, i. e., to provide an integrative and explorative analysis of both design domains, flapAssist and memoSlice should no longer be stand-alone applications though. In a fully integrative planning and optimization process it should, for instance, be possible to instantaneously experience how machine configuration variations affect factory-related parameters, like material flows. Therefore, we have integrated memoSlice into flapAssist, where it serves as a widget for machine configuration (see Figures 5 and 6).

To enable interactive exploration of such a highly integrative scenario, a high degree of responsiveness is an absolute prerequisite (also see Chapter 5). Thus, the Virtual Production Intelligence relies on task-based parallelization techniques with a specific user-centred prioritization scheme and streaming updates to guarantee fast response times, without impairing the user's workflow. Even as stand-alone applications, the computational loads in flapAssist and memoSlice are rather high and require high-end computing hardware and parallelism. In memoSlice, for instance, the update of hyperslice views during user interaction requires the computation of radial basis functions for every pixel. Therefore, hierarchical parallelization schemes are applied here that are specialised for this use case.

Data annotation in immersive analytics tools An important issue of existing VR-based factory planning tools and, most probably, any VR-based analysis tool in science and technology, is the lack of efficient workflows for creation and access to annotations. Therefore, we developed an annotation framework, which can be

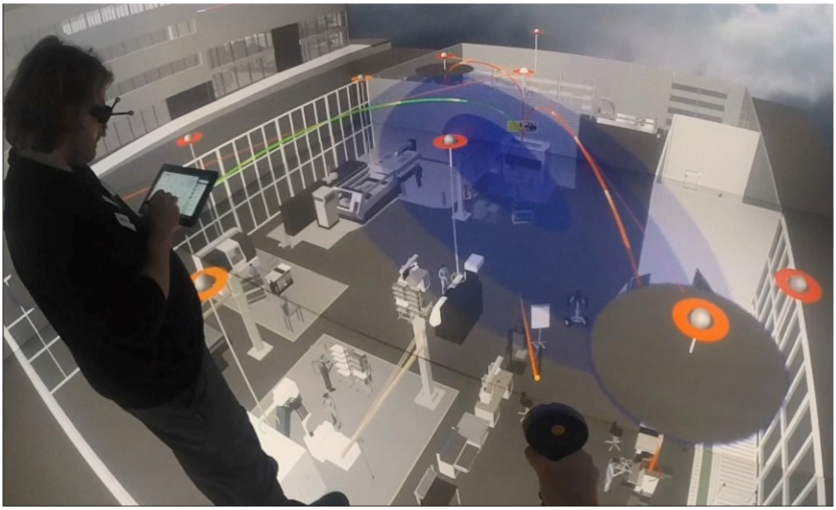

Fig. 3: Interactive optimization of a digital factory's shop-floor layout.

used with flapAssist and memoSlice but, in principle, with any other immersive analytics application, too [28].

As previously described in [19], a variety of annotations can be created via our framework, like labels or viewpoint annotations, and linked to different types of data, like text, voice comments, images, or sketches. To capture all these different data types, the annotation system provides a selection of interaction metaphors. Since standard keyboards typically do not work for entering text from within an immersive virtual environment, in our annotation framework text entries can be accomplished via smartphone input or, alternatively, via a speech-based system called *Swifter*, that we specifically designed for simplicity while maintaining good performance (see Figure 7) [28].

An annotation server based on Web services stores the annotation data in a standardized way so that it should be easy to integrate it not only into the VPI platform, but also into any other infrastructure. By this, annotations become accessible to all involved parties, thereby supporting truly collaborative factory planning.

Conclusions and lessons learned from this use case With the Virtual Production Intelligence platform, a first prototypical immersive analytics demonstrator for a use case in production engineering is under development, focusing on an integrative approach for planning and optimization tasks in the design domains Factory and Machine. A formal evaluation in terms of comprehensive user studies is still to be performed to evaluate whether an immersive analytics approach is, in fact, superior to a non-immersive Visual analytics approach for this use case. Expert studies for the evaluation of single system components have been already conducted with promising results. By means of the VPI approach, combining

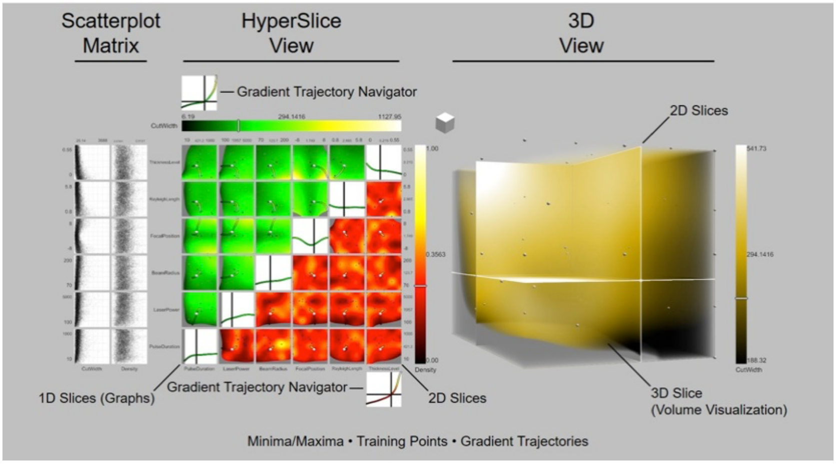

Fig. 4: The linked multiple view concept of memoSlice.

life-sized virtual walkthroughs of entire factories with further meta information of factory layouts and single machines, it becomes possible to compare design alternatives in an intuitive, cost-effective and timely fashion.

While real-time, low latency visualization and interaction are an issue in any immersive VR application in general, our case study clearly demonstrates that performance is an even bigger challenge in immersive analytics scenarios. Due to a high integration level of heterogeneous, partially large, data and computationally expensive analysis algorithms, efficient data management, as well as advanced, hierarchical parallelization strategies with user-centric priority schemes, need to be applied to make the analysis an explorative, human-in-the-loop experience. Furthermore, we consider a powerful annotation framework as crucial for any immersive analytics tool, making it possible to extract information and insight from immersive analysis sessions. Annotation data should be stored in a standardized way, so that it can be used in a variety of setups and further processed by collaborators in subsequent immersive or non-immersive analysis sessions.

11.3. Case Study 2: Immersive Analytics in the Interactive Visualisation of Sustainable Cities

Achieving liveable, sustainable and resilient cities of the future requires an in-depth understanding of the complexities, uncertainties and priorities that characterise urban communities, environments and infrastructures. In this context, infrastructure planners, designers and managers are seeking innovative processes that communicate across communities and lead to robust planning outcomes.

This chapter outlines immersive analytics considerations in the interactive exploration of virtual city environments under different future scenarios. The project's case study is the Elster catchment, an extensive area in the south-eastern

Fig. 5: Factory and machine planning in an immersive environment, integrating geometric and abstract information in a unified analytics tool.

suburbs of Melbourne, Australia. An immersive 3D model of the suburb, alongside support for interactive climate scenarios, was assembled at Monash University to engage various stakeholders and communities as the region moves into a more climate-uncertain future. Such stakeholders include the local government authorities that intersect with the Elster Catchment, the relevant water boards and government bodies, and local residents. As such the model presents an ideal case study for the communication of complex data to a broad range of users of varying technical and disciplinary expertise. It as an example of immersive story-telling (see Chapter 6)

This case study aims to address a set of core concerns that cut to the heart of the idea of the 'model' in the emerging field of immersive analytics. The first core concern is the idea of the model as simulation. The case study embeds practices and approaches that anticipate *immersive analytics* presence in the design, communication, and encountering of urban space.

A 'fluid' model and an engine for cities The project commenced with a comprehensive water and rain-loading model for the entire catchment of the Elster Creek in south-eastern metropolitan Melbourne. The watershed rises in the former sandy heathlands of the south-east, falling down toward Port Phillip Bay. Its lower reaches, which drain what was once the Elwood swamp, are extensively channelized and subject to regular flooding. The creek wraps up a microcosm of issues facing Australian urban ecologies; both vulnerable and dangerous; a

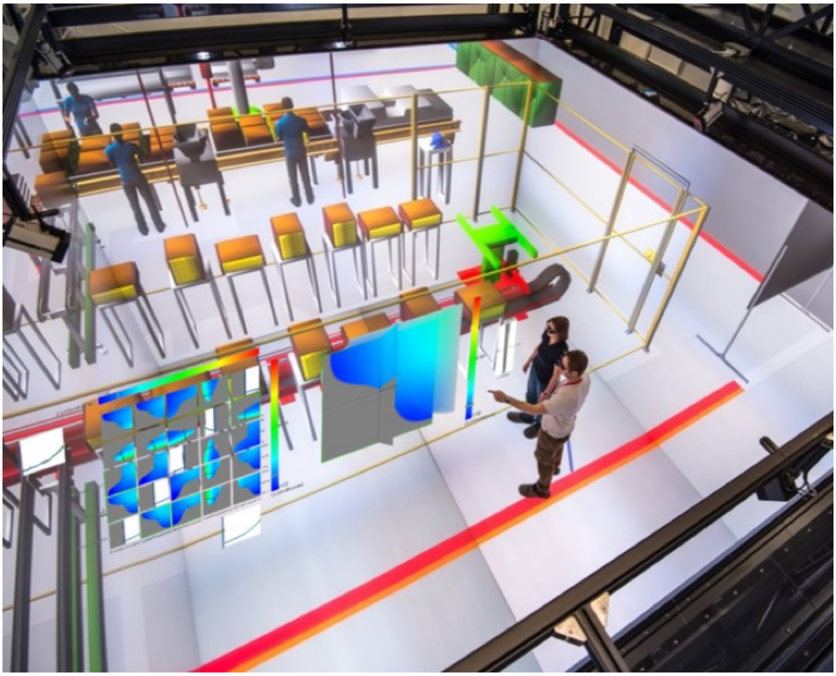

Fig. 6: Integration of factory and machine planning, bird's eye view into the RWTH Aachen aixCAVE.

contested space between the natural and the man-made whose exigencies are only exacerbated by climate change.

The rain-loading model considered the watershed as a cohesive system in which planning upstream could deliver substantial benefits downstream. The model offered a way of recording the impact of a series of small and medium scale interventions, such as water run-off retention, streetscaping, and domestic rainwater tanks, during a set of forecast rain and tidal surge events.

The model, while comprehensive, suffered from many of the issues that affect complex data visualisation in that it requires existing spatial/analytical skills and a passing understanding of the methodology of the study. Other ideas engaged through the model–like changes in the morphology of the urban space, shifts in density and housing approaches–are also less evident in the static outputs of the model-as-maps. The core issue, however, was an absent sense of tangibility and immediacy. That is, the scope and scale of the flooding impacts were not immediately legible, and the ways in which differing scenarios might impact the same space of the city was not readily apparent. Finally, the model embedded issues with what Ian Bogost terms simulation fever–the issue around completeness and accuracy in the simulation of events.

Bogost places emphasis on the narrative capacity of simulations, suggesting that; "A simulation is a representative of a source system via a less complex

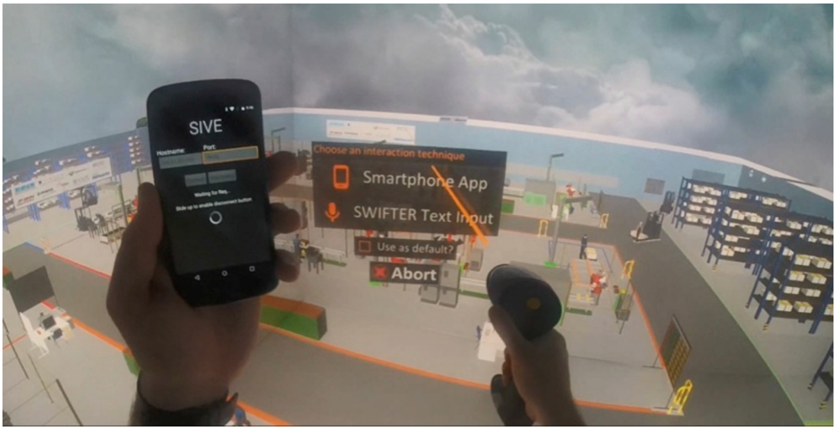

Fig. 7: Creating annotations for smartphone-based communication with the VPI analytics engine.

system that informs the user's understanding of the source system in a subjective way." [3, p. 98] More importantly, these simulations provide 'fungible' insights, narrative structures into reality, rather than a one to one mapping between the model and reality. Bogost goes on to suggest that; "What simulation games create are biased, nonobjective modes of expression that cannot escape the grasp of subjectivity and ideology." [3, p. 99] While Bogost uses this in the sense of interactive simulations as games, we want to argue that this holds true for all interactive and immersive spaces.

As the model was to be used to buttress and support a series of community outreach workshops that were aimed at clarifying and communicating the outcomes, the cooperative research centre for water (CRC for Water) commissioned an interactive virtual model of the suburb of Elwood that could illustrate a set of outcomes from the earlier water modelling.

The last two decades have seen a number of key texts on the idea of computable and simulatable cities. At one level this is an offshoot of the concepts and ideas bandied around by the cybernetic model of the city as a digital/human organism in the 1960s. We might find this in the ambivalent attitude of figures such as DeLanda [10, pp. 95-96], and the strange reductive processes of actor-network-theory (ANT) in the context of the city. A second thread (no less reductive) is the drive to image and model large cities, sites that appear to resist imagining *in toto*. Work by Pascal Müller *et al.* [23, 26], at the ETH, then Procedural, and finally at ESRI, sets the stage for the comprehensive modelling and imaging of cities.

So there is a question around the nature of this material. Approaches to 3D city models tend to proceed from the structure of the transportation network–either generated or inherited from spatial data. These are used to create developable blocks which can be further subdivided to serve as the starting point for generating

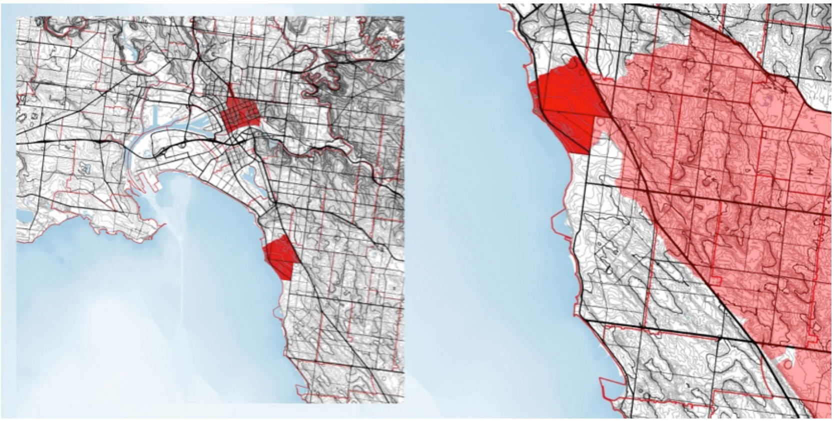

Fig. 8: The suburb of Elwood sits 8 kilometres south of Melbourne's CBD (left). The site forms the creek-mouth of the Elster catchment (right) and is subject to regular pluvial and fluvial flooding.

geometry based on land-use, lot size and orientation. It is, however, possible to neatly sidestep this process; the best asset for the initial generation, in this scenario, is the existing footprint of the dwelling–which embeds a kind of latent history of the site and the tectonics. 'Swamped,' an architectural design studio that ran parallel to the CRC investigation, furnished a set of building footprints as shapefiles with associated metadata. This was material that would have proven useful to the architectural students in the studio and indexed the apparent age of the property, the number of stories and the presence of at-grade car-parking.

These simple data points, combined with the footprint, provide seeds for the construction of dwelling geometry across the Elster catchment. The particular grammars act as analytical placeholders, distilling (visual) meaning out of otherwise undifferentiated footprints and data-tables. Combined with a road network, Cityengine is able to provide an index of street 'frontages' and the grammars are able to adjust to reflect this. The grammars are often able to identify the presence of extensions or renovations as the building footprint has additional vertices or extends beyond the expected scale of the historical type.

As a general procedure, the model involved a high-level separation between formal styles–historical styles associated with the age of the building. Elwood is a heterogeneous mix of styles–ranging from late-Victorian terrace houses, through federation and bungalow styles in the regions reclaimed from the swamp, to post-war brown-brick apartments, and contemporary developments along commercial arteries and the foreshore.

The 3D spatial model allowed for the superimposition of analytical material across the site–and allowed us to use it as a design seed–embedding further narrative intelligence into the model.

Areas subject to inundation or heat-stress varied depending on the parameters of the underlying simulation. The analytical component was represented as a set

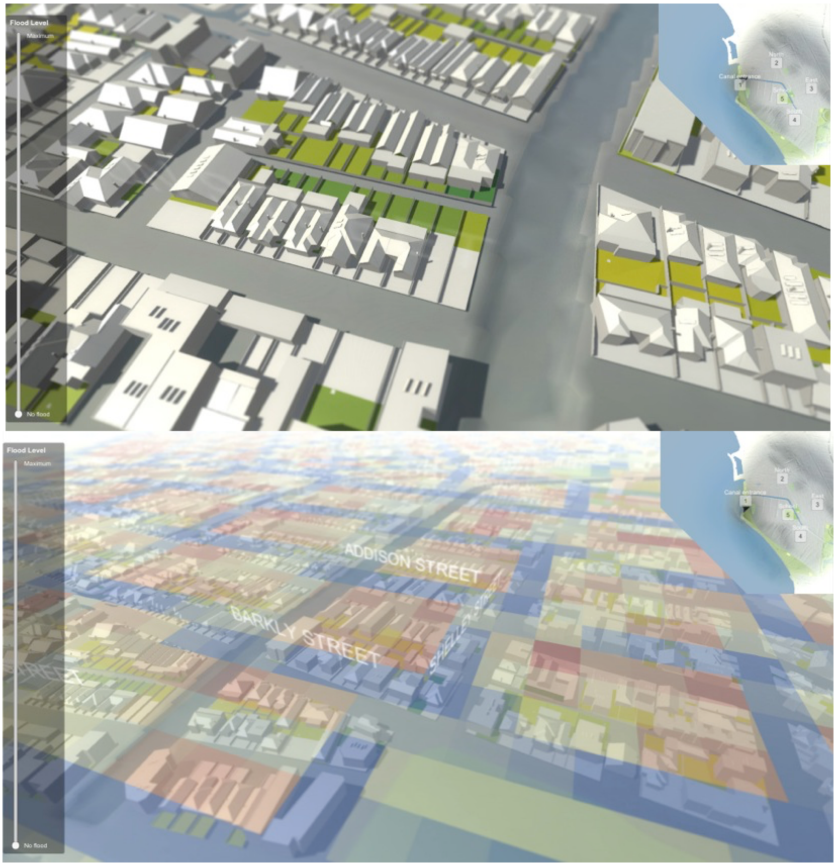

Fig. 9: Heat Island modelling was used to seed ‘stressed’ environments in the model–this affected the type, extent, and condition of vegetation. The nature of this information shifted from explicit (temperature) to implicit (vegetation type) as the user navigated down to street level.

of colour shifts or modifications to underlying models; areas subject to repeat flooding were paler, or the generative grammars of buildings were adapted to puncture and reveal underlying structure as if they had fallen into disuse and disrepair. Areas subject to increased heat-stress had a corresponding die-back in projected vegetation, and a ‘browning out’ in their resulting colouration. The simulation tracked changes in ground-cover that could influence future micro-climates. At the large scale, these were represented using the grid-cells of the simulation–giving a comprehensive overview of temperatures and stresses. But at the finer building scale, these were represented directly–imaging stressed and thriving gardens, and showing the extent of future green-space for each building lot.

Other dynamic details were used to code elements of the overarching narrative–incremental features such as water-tanks or solar panels could be instanced across the model, depending on the ethical valence of the scenario. Water sensitive streetscape features–such as swales and planting–could be instanced in appropriate sites based on stresses identified in the underlying model. Similarly, the stresses placed on building stock by rising waters and heat-stress can be directly, if hyperbolically, imaged using the narrative register of dilapidation and decay (see Figure 10).

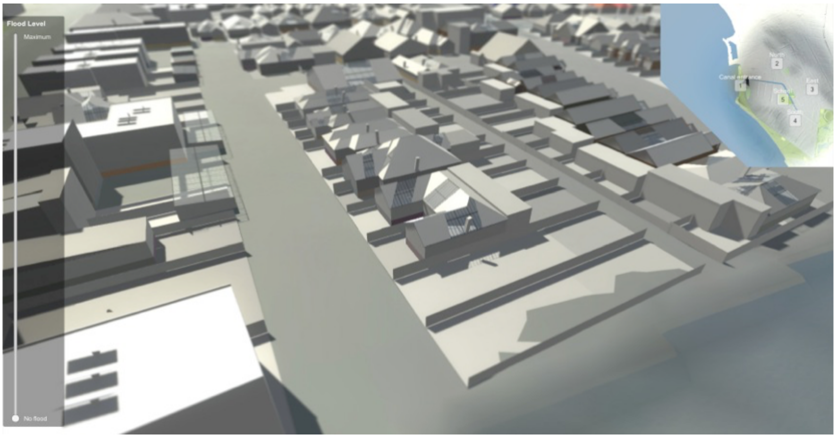

Fig. 10: Impacts of climatic shifts were conveyed through the 'image' of the built fabric of the suburb.

Critical and narrative elements are also communicated through procedural means–the surrounding city is dynamic, and hints at larger economic and socio-political conditions that might be attached to the underlying data model–the condition and coherence of the city skyline changing in response to the valences and registers of the scenario–shifting from ruin to shining, coruscating object as the overall story shifts from climate disaster to anti-fragile responses.

Finally, the model is able to communicate using other channels–i. e., spatial awareness, things that are coded as stories and performances. This scaffolds off a capacity for coherent immersive environments to function as narrative or story spaces, and in doing so communicate complex insights. While the immersive analytics field is young, similar affective powers are evidenced in the landscapes of games and simulations. In this mode, the purpose of the model or simulation is more than just a site for representation or computation. As critic Steve Fuller outlines, a possibility is that; "The point of *virtual reality* is to realise the latent potential of the actual world, typically by getting us to see or do things that we probably would not under normal circumstances but could under the right circumstances." [15, p. 105]

Fig. 11: The model offers the capacity to image extreme scenarios, and scaffold them off recognisable visual miscellanea to further entrench the enormity of future changes.

For a site in the built environment, which engages with multiple stakeholders and necessitates creative re-imaginings of the world, this is the direct benefit of an immersive analytical model. Where immersive analytics differs from virtual reality in that it not only allows for the consideration of other possibilities but also offers concrete and direct information and data that allow stakeholders to act on these possibilities.

Conclusions and lessons learnt from the use case Immersive Analytic models of the built environment have the capacity to act as 'transformative narratives' – tools to prompt speculation and corresponding act on such speculation. Renata Tyszczuk frames transformative narratives as:

"... those stories that have an open-endedness that makes space for the unforeseen–a future we cannot ever really understand. Such stories have agency: they can provoke us to think about how we might live with the prospect of uncertain futures, how we can prepare for situations we cannot anticipate, to think through our responsibilities to others and help develop our adaptive capacities." [33, p. 133]

The Elwood visualisation was commissioned to serve a straightforward purpose–the communication of sea-level and storm surge scenarios in an uncertain future. However, the combination of the robust underlying data model and recognisable built spaces allowed for powerful synergies. The immersive analytical model allowed the discrete time-steps of a severe weather event to be experienced as an *event*–as a space and a territory that unfolded around the viewer/participant, with storm clouds rolling in across the bay and floodwaters engulfing the old Edwardian houses along Barkly Street.

From an architectural and urban design perspective, the immersive analytic model differs from normative design models. It is not a design outcome but a continuum of potential responses–a set of Tyszcuck's 'transformative narratives.' Critically, they allow what are often disparate disciplines (design, engineering, planning) to communicate in a common site and model and to share this with a broader community. The model ties together various data sets that would not commonly be intersected or superimposed–and in their correlation creates new adjacencies and awarenesses. The completed asset immerses the wide array of actors and stakeholder in a rich and contingent set of futures, and ultimately, compel us all to engage and to *act with new knowledge*.

Citizen actors/agents engaged with the model with a sense of desperation or alarm. The mixture of contingent scenarios (in which some houses were rendered uninhabitable) reasserted the contingent nature of all actions in this space. The discussion turned toward the myriad of choices available to planners and designers in this space–to ways of bypassing path determinism. As we move forward, a manner and method to capture this insight would be desirable–increasingly integrating this into the immersive, analytical environment–closing the cycle of the decision and discussion loop.

Fig. 12: The interactive visualisation forms the crux of further discussions with community, stakeholders, and designers.

11.4. Case Study 3: Crowd Simulations in a Virtual Model of Medieval Angkor

The third case study in this chapter overviews a research project at the Faculty of IT at Monash University that attempts to visualize how the Angkor Wat complex might have operated almost a millennium ago. Broadly, the aims of this project are twofold. The first aim is to leverage from data obtained in recent archaeological surveys to craft a comprehensive virtual reconstruction of Angkor Wat as the centrepiece in the medieval Cambodian metropolis of Angkor in the 12th century; this is the ‘built environment’ component of this study. The second aim is the simulation and tracking of thousands of animated ‘agents’ as they enter, exit and circulate around the complex–the immersive analytic component of this study. As an academic exploration that uses simulation to test how assumptions can be made more precise, the primary users of this system are archaeologists and historians. That said, provided it can be mediated in an educational context, an interactive simulation of daily life medieval Angkor Wat would make an ideal exercise in the Australian high school history curriculum, where the civilization of Angkor was recently nominated as a key depth study.

Constructed in the 12th century by the Khmer king Suryavarman II, the temple of Angkor Wat in Cambodia is a world famous heritage site and the largest religious monument on earth. The temple’s well-preserved stone architecture and decorative reliefs have been the subject of extensive scholarship [20,24,29,30], but until recently, the wooden settlement that once lay within the temple’s enclosure walls remained unknown. In 2013, LIDAR archaeological surveys confirmed a grid pattern of roads and household ponds, suggesting a regular layout of dispersed but substantial wooden dwellings [13].

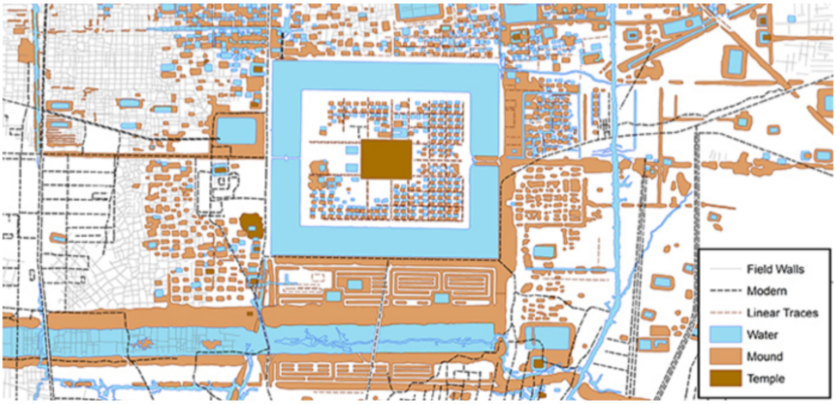

Fig. 13: A preliminary archaeological map of Angkor Wat (at centre, surrounded by a square moat) and its environs based on an analysis of Lidar imagery. *Image courtesy of Damian Evans / Khmer Archaeology Lidar Consortium*

Our simulation of the Angkor Wat enclosure in historical times is guided by two primary constraints in space and time. The temporal constraint was a single day set at some indeterminate date in the 12th century. In contrast to broad scale ‘virtual archaeology’ studies that plot change over centuries, such as the Rome Reborn Project [14], this simulation focuses specifically on just 24 hours: a day in the medieval life of Angkor Wat. The spatial constraints are delineated in the 1500 by 1300 metre space of the Angkor Wat enclosure bounded by its 200-metre wide moat. However, these spatial limits are only meaningful in describing the Angkor Wat complex as a two dimensional map and not as a simulated built environment. If we were only concerned with simulating the hypothetical flow of crowds around Angkor Wat, a map-based traffic simulation could have sufficed. A flat map (as in Figure 13) with multiple layers of cartographical data would allow us to plot the movement of human figures as points moving along the 2D vector lines around polygon boundaries. However, while a purely two-dimensional simulation would undoubtedly have analytic qualities, it would hardly be immersive

In a similar manner to the accompanying case studies in this chapter, the data that facilitates spatial immersion in the reconstruction of Angkor Wat in medieval times is not only intrinsically 3D in format, but comes with its own particular complexities.

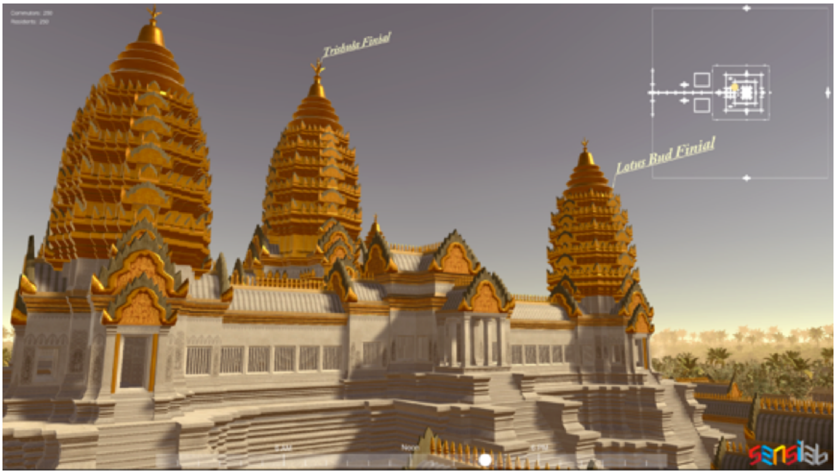

Fig. 14: Floating annotations describe various art historical elements in the temple's architecture.

To begin with, there is the 3D modelling of the ornate and architectural complexity of the stone edifice of the temple, its gates and walls. Though worn and at risk of collapsing in places today, the stone architecture of Angkor Wat is very much still in place. As such, its virtual reconstruction doesn't require the careful weighing of varying degrees of uncertainties as those deduced from faint imprints, such as the gladiatorial school or auxiliary fortress of Carnuntum [25, 34]. The reconstruction of the urban settlement surrounding the temple revealed by Lidar mapping, however, did need to account for a number of uncertainties. Angkor was essentially a green and wooden city, thick with trees shading the timber and thatch dwellings of its inhabitants (see Figure 15). Our simulation had to account for the likely species of trees and plants for Angkor in the 12th century drawn from epigraphic [18], archaeological [27] and historical botanic studies [2]. The patterning of vegetation was randomised and clustered in open spaces between the approximately 283 mounds and 250–300 ponds within the Angkor Wat enclosure [31, p. 1444]. Moreover, because the excavation data supported a scenario where "temple personnel of modest material wealth occupied relatively insubstantial, perishable structures located on mounds in the immediate vicinity of the temple" [31, p. 1452], our models of thatched, roofed houses were drawn from studies in traditional Khmer architecture [11, 32]. Given these uncertainties, one of the primary advantages of a 24-hour cycle was the opportunity to alter, adjust and re-order architectural and vegetation assemblies at initialization, each time the simulation started up.

Fig. 15: The cultural landscape surrounding the temples, showing wooden dwellings on raised stilts, vegetation assemblies derived from archaeological pollen cores, and smoke from cooking fires.

The 3D animated models of human figures presented their own complexities, for they had to be planned as mobile collections of artefacts, insignia and motifs. The adornment of the Angkorian populace, so well attested in the bas reliefs [22, 30], opens up the visualisation of the fabrics that cloth them [17] of the jewellery that adorns them, and even–speculatively for medieval Angkor–the magical talisman tattoos on their skin. And then there are the things that they carry–the balanced baskets and shouldered loads–and, if we are to visually model social hierarchies and distinctions, the colours and sumptuary motifs on the parasol that retainers hold aloft behind over their masters (see Figure 16).

This virtual space, rich in the hues of greenery, patterned cloth and gilded spires, constituted the immersive foundation for augmented textural, temporal and cartographic information. Floating descriptive 'signifiers' in white text would appear above walking figures and architectural elements when the viewer moved close to them (see Figures 14 and 16). A plan map of the temple enclosures in the top right of the screen specified the position of the viewer's camera as a yellow circle (see Figures 14, 15 and 16) and the time of day was recorded in a slider at the base of the screen. Each agent in the visualisation was also tracked by summary statistics; a coloured circle indicated their category (explained below) and floating lines of text listed their destination and the time each agent had been active within the visualisation space. Though these overlaid features were arguably interactive and dynamic in the immersive space of the visualisation, the purpose they served was more descriptive than analytical.

If immersive analytics can be broadly defined as the exploration of how new interaction and display technologies can be used to support analytical reasoning and decision making [7] then in this simulation of Angkor Wat the key analytic concern was the visualization of the cumulative movement of agent groups. Early attempts in envisaging agent motion trails and patterns of aggregate activity in the visualization of the Angkor Wat enclosure consistently ran up against limits in processing power and memory, not to mention the limits of the Unity engine in tracking vast numbers of animated meshes. Consequently, we resolved to split the simulation and visualisation into two separate but linked applications; the visualisation would render the details of the animated scene, and the simulation would deal with the tracking the agent paths under differing scenarios. And there were quite a lot of agent paths to deal with.

Current archaeological estimates suggest that at its peak, the Angkor Wat complex may have housed up to 4,500 residents within its walls. If preliminary inferences can be made from the Ta Prohm inscriptions, the temple was serviced by a workforce of 25,000 and a support population of 125,000 people [12, p. 1410]. Pending further excavations and historical research, not a great deal is certain about the daily activities of the historical 'agents' we are attempting to simulate. Admittedly some speculation is involved in this venture and we must extrapolate the numbers, categories and activities of our agents from the Ta Prohm and Preah Khan temple inscriptions [8, 21]. Our intention was to construct a hypothetical space that presents not one but several possibilities, and can accommodate for varying shades in between. Figure 16 shows our initial attempts at visualising these estimates as crowds of walking figures or 'agents' moving in, out and around the Angkor Wat complex over a 24-hour period. Luckily, only a fraction of the estimated workforce of 25,000 needed to be present at any one time in the simulation. Working on the assumption that the populace attached to the temple worked on an alternating fortnightly roster, we could reduce the workforce the total numbers of agents on any given day by 50%. Given that many of the agents essentially disappeared when they entered the temple or exited from the edges of the simulation, a further reduction could be made by only tracking agents who were both moving around and in view. What the agents were doing at any given time depended upon how we categorized them.

We divided these crowds of animated 12th century Cambodians into four broad categories: *visitors, residents, commuting workers* and *suppliers.* All agents were guided by broadly similar rules, though each agent category was given a

Fig. 16: Each agent is tracked by summary statistics. A coloured circle indicates their category (blue for commuting workers, red for residents, green for suppliers, yellow for visiting elites) and a floating 'signifier'- revealed when the camera approaches close by–shows their destination and the time they have been spent within the confines of the visualisation space.

different agenda for navigating the space of the enclosure. The *resident* category of agents circulated between the residential blocks and the temple but did not venture beyond the temple walls. The *commuting worker* category, by far the most numerous, entered the complex from any gate and make their way to a random point in the enclosure, where they would remain for a number of hours before leaving the complex by the same way they came. The *supplier* category of agents behaved much like the commuting workers except that they did not venture inside the complex but instead approached the gates, stopped briefly, and then returned the way they came. The *visitor* category–comprised of high-status individuals and their retainers–entered the temple through the central western gates, make their way along the interior stone causeway, after a random number of hours, returned the way they arrived.

Our simulated 24 hours is provisionally set in the dry season and we designed two 'pulses' of activity during daylight hours; the first was from dawn until 11am and the second from around 2pm to dusk. During the hottest time of the day, between 11am and 2pm, activity would be more subdued. During these times, the agents were programmed to avoid travelling and to seek shade, preferably near the wooden dwellings, where they would sit in groups to pass the time in company. Between 2pm and 6pm (dusk) commuting workers would again start to arrive or leave the complex.

When testing scenarios about how the workforce of Angkor Wat would have entered, exited and moved around the temple, a dynamic overlay that superimposed each agent's motion trail over time and space as moved around

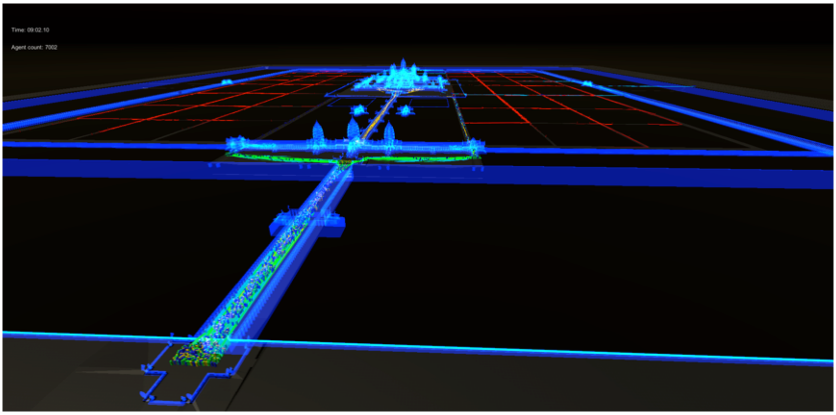

Fig. 17: A test visualisation of the tracking of agent paths during the early hours of the day. The mass of green and blue on the western causeway the red colours indicating residents and the green indicating suppliers. The trails mark the passage of the agent over time. *Image created by Kingsley Stephens*

the complex (see Figures 17 and 18) was a valuable device to visualise the consequences of agent decisions en masse. By turning off the textures and lighting and adopting an 'x-ray view' in the model, the residual colour coded trails the agents left behind as they ventured toward their destination meant that key thoroughfares and efficient routes became overlaid and thickened with colours the more that agents opted for them. Each agent was programmed to follow the shortest path (using an A* algorithm) to travel from point A to B, but with the proviso that certain paths were likely off-limits to certain agent categories and so too were certain gateways. Because the agent paths were coloured according to their category, the emergent effect was a 'heatmap' that visualised the overall distribution, destination and circulation of thousands of agents over the course of the day. Because slight changes to the rules governing the agent categories could result in vivid colour and phase changes in the map, the effect of the parts–thousands of individual characters making simple decisions–could be analysed as a whole.

As well as tinkering with agents numbers in general hypothetical tests, such as shifting the drop-offs in morning and afternoon activities and adding or subtracting numbers of agents in each category, we could also move toward testing our model against more particular scenarios. For example, restricting the central western gates to the elites (the visitor category) would see the animated populace branch out from a multi-coloured stream into a trident divided by social hierarchy, where the commuting worker and supplier categories separated from the visitor category and made their way into the complex through alternate gates. Another test concerned hypothetical structures. The recent discovery of multiple breaches in the north and south walls of the complex suggested that these too were once used as thoroughfares in medieval times [4]. While there is no evidence of bridges across the moat on either of these sides of the complex today, the existence of bamboo bridges of light construction in medieval times is a plausible hypothesis to test. The ability to toggle a hypothetical bridge across the north and south sides of the moat 'on and off' meant that agents would cross the bridges to achieve their destination if the bridges were present, with the result that the agent paths weave in and out to link the interior of the complex to the outer sides of the moat like the struts of a spiders web. If the bridges were absent, agents would have to take an alternate path to their destination by a considerably longer route.

Conclusions and lessons learnt from the use case Given the inherent challenges of visualising the past with incomplete data, our simulation was bound to be more explorative than predictive. Does a simulation proposing new access points into Angkor Wat support the hypothesis that wooden bridges existed there in medieval times? In lieu of hard evidence such as the remains or imprints of bridge posts under the mud of the moat, the answer remains equivocal. However, anything that can force assumptions to be more precise remains valuable in the investigation of the civilisation of Angkor.

Certainly, more could have been made of interactive aspects of the simulation especially in enabling the user to alter the weightings of agent numbers and goals in real time. As it was the simulation's variables and the layout of the 3D structures could only be altered offline. It followed that aspects of interactivity and immersion while spatially 'deep' were restricted in function. In navigating the simulation, the user's interaction was limited to a free roving camera, however, whether the user moved through the sky or along the roads, the texturing and animation of the thousands of 3D models populating the space ensured a visual consistency that meant the user was immersed in a cohesive virtual world. The addition of sound, belatedly mentioned here, underscored the spatial immersion with a sonic dimension. Samples drawn from field recordings were assigned specific locations within the enclosure and coded so that they faded in and out of earshot as the user moved between their source points. Thus, the sound of a conch shell horn emanating from the inner sanctums of the temple is loud when the user is near it, but muted when they are further away. And if the viewer floated the camera close to groupings of animated people they would discern fragments of conversations in Khmer, but if they moved the camera up above the tree line these sounds gradually diminish, and instead they would hear the wind ruffling the leaves of the sugar palm trees.

Though superimposed diagrams, floating 'signifiers' and motion trails that faded over time augmented the simulation with some analytic features, the key analytic feature of the immersive space was the library of 3D models themselves. It follows that in an immersive, contextual historical space where every evidence-based reconstruction relates to those around it, there is something to be said for the immersive analysis of the historical 'likelihood' of the 3D reconstructions themselves. For example, when the simulation was presented on a 3840 x 1080 multi-projector interleaved display at Monash SensiLab, visiting historians and archaeologists, while spatially and sonically immersed, were able to analyse whether particular 3D reconstructions in the Angkor Wat enclosure were plausible or implausible for medieval Angkor given the other models–including walking figures–around them.

Finally, there were the unexpected, serendipitous visualisations that emerged from the application of immersive analytics at Angkor Wat, one such example being the overnight visualisation in Figure 18.

Fig. 18: A test visualisation of the tracking of agent paths through the enclosure overnight. *Image created by Mike Yeates*

We have no firm evidence for activities at Angkor Wat through the night. At any rate, when night fell upon our virtual model our agents had almost all departed or retired indoors. However, our simulation accommodated for a small number of residential agents walking back and forth from the temple throughout the evening. The overlaying of agent paths in Figure 18 showed the trails of these wanderings along dark roads not unlike the slow exposure photography of cars through an urban centre at night. Though we have not yet modelled and animated the small ceramic oil lamps these nocturnal agents would have carried, the visualisation was nevertheless a surprising one. Here was a world famous monument that had been digitally reconstructed countless times, but always in the hard ambient occlusion of a simulated sun. This overnight visualisation suggested an intriguing view of the paths taken over night at a preindustrial centre lit only by the moon.

11.5. Conclusion

While the word model conjures up different meanings according to the discipline using it, a broad description is that a model is a mental representation of an external phenomenon [35]. A model might be simple, complex, metaphorical or physical. It can be conveyed through prose, diagrams or cultural practice. Perhaps the most easily recognised definition of a model is a scaled representation of a real object, as in a 'model' aeroplane, or a scale architectural model constructed carefully from cut segments of wood and plastic. While the built environments depicted in the illustrated figures of the above case studies are all digital, they share an easy visual interpretability because they seek to mimic visual and spatial reality directly. The added third dimension that all these case studies share means that they have transformed factory floor plans, city street atlases and archaeological survey maps beyond cartography into cohesive and self-contained virtual worlds. The fact that the virtual environments in the case studies discussed above can accommodate realistic 3D human figures further underscores the fact that they depict a *model* of reality that is recognisable, and recognisably 'immersive' in a way that a 'business model' or 'theoretical model' isn't.

This familiarity offers certain advantages as well as challenges. Because the environments incorporate familiar spatial cues navigation is intuitive whether the viewer walks through the environment in first person mode or floats above the scene. However, the fact that we are mostly navigating within an immersive environment presents a challenge to embed annotations and other documentation and to allow these to be viewed consistently within the environment without losing immersion, occlusion issues, and unnecessarily complicating the user experience. The issues of motion sickness and latency are more critical in simulations of the built environment than in other subject fields where graphics are usually centred and not presented as scale models of real spaces.

Despite these caveats, the advantages discussed in the above case studies are manifold. By enabling life-sized walkthroughs in immersive VR planners can not only experience a factory in its original scale but also compare, with the aid of analytic meta information, design alternatives in an efficient and intuitive manner. For the visualisation of sustainable cities, the combination of the underlying data model and recognisable built spaces allowed for powerful synergies; namely the discrete time-steps of a severe weather event to be experienced as an *event*–as a space and a territory that unfolds around the viewer. And, finally, there are the unexpected insights that emerged from the application of the immersive analytic tracking of animated characters–agents–following simple rules as they navigate reconstructed paths and avenues in a re-built environment of long ago.

References

1. Al Khawli, T., Eppelt, U., Schulz, W.: Sensitivity analysis of laser cutting based on metamodeling approach. In: Handbook of research on computational simulation and modeling in engineering, pp. 618–639. IGI Global (2016)
2. Blench, R.: A history of fruits on the southeast asian mainland. Occasional Paper 4, 115–37 (2008)
3. Bogost, I.: Unit operations: An approach to videogame criticism. MIT Press (2008)
4. Brotherson, D.: The fortification of Angkor Wat. Antiquity 89(348), 1456–1472 (2015)

5. Büscher, C., Meisen, T., Schilberg, D., Jeschke, S.: VPI-FP: an integrative information system for factory planning. International Journal of Production Research (2015) doi: doi:10.1080/00207543.2015.1057298
6. Caputo, F., Di Gironimo, G., Marzano, A.: A structured approach to simulate manufacturing systems in virtual environments. In: Congreso International de Ingegneria Grafica. pp. 73–80 (2006)
7. Chandler, T., Cordeil, M., Czauderna, T., Dwyer, T., Glowacki, J., Goncu, C., Klapperstueck, M., Klein, K., Marriott, K., Schreiber, F., Wilson, E.: Immersive Analytics. In: Big Data Visual Analytics (BDVA) 2015. pp. 73–80. IEEE eXpress Conference Publishing(2015) doi: 10.1109/BDVA.2015.7314296
8. Cœdès, G., Cordier, P.: La stèle de ta-prohm. Bulletin de l'École française d'Extrême-Orient 6(1/2), 44–85 (1906)
9. Corves, B., Loftus, M.: Designing an immersive virtual reality interface for layout planning. Journal of Material Processing Technology 107(1–3), 425–430 (2000)
10. DeLanda, M.: A new philosophy of society: Assemblage theory and social complexity. A&C Black (2006)
11. Dumarçay, J., Royère, P.: Cambodian architecture: Eighth to thirteenth centuries, vol. 12. Brill (2001)
12. Evans, D., Fletcher, R.: The landscape of Angkor Wat redefined. Antiquity 89(348), 1402–1419 (2015)
13. Evans, D.H., Fletcher, R.J., Pottier, C., Chevance, J.B., Soutif, D., Tan, B.S., Im, S., Ea, D., Tin, T., Kim, S.: Uncovering archaeological landscapes at Angkor using lidar. Proceedings of the National Academy of Sciences 110(31), 12595–12600 (2013)
14. Frischer, B.: The Rome reborn project. How technology is helping us to study history. OpEd. University of Virginia (November 10, 2008) (2008)
15. Fuller, S.: Humanity 2.0: What it means to be human past, present and future. Springer (2011)
16. Gebhardt, S., Hentschel, B., Kuhlen, T.W., Al Khawli, T., Schulz, W.: Hyperslice visualization of metamodels for manufacturing processes. In: IEEE Vistualization Poster Proceedings (2013)
17. Green, G.: Traditional textiles of Cambodia : cultural threads and material heritage. Published and distributed in Thailand by River Books, Bangkok, Thailand (2003)
18. Jacob, J.: The ecology of Angkor. Evidence from the Khmer inscriptions. In: Stott, P.A. (ed.) Nature and Man in South East Asia, pp. 109–127. School of Oriental and African Studies, London (1978)
19. Jeschke, S., Kampker, A., Kuhlen, T.W., Schuh, G., Schulz, W., Al Khawli, T., Büscher, C., Eppelt, U., Gebhardt, S., Kreisköther, K., Pick, S., Reinhard, R., Tercan, H., Utsch, J., Voet, H.: Virtual production intelligence (VPI). In: Integrative Production Technology, pp. 177–251. Springer International Publishing (2017)
20. Mannikka, E.: Angkor Wat: Time, space, and kingship. University of Hawaii Press, Honolulu (1996)
21. Maxwell, S.: The stele inscription of preah khan, angkor: text with translation and commentary. UDAYA, Journal of Khmer Studies 8, 1–114 (2008)
22. Maxwell, T.S.: Of gods, kings, and men: the reliefs of Angkor Wat. Silkworm Books (2006)
23. Müller, P., Wonka, P., Haegler, S., Ulmer, A., Van Gool, L.: Procedural modeling of buildings. ACM Transactions On Graphics (ToG) 25(3), 614–623 (2006)
24. Nafilyan, G.: Angkor Vat, description graphique du temple. Ècole française d'Extrême-Orient, A. Maisonneuve, Paris, (1969)

25. Neubauer, W., Seren, S., Hinterleitner, A., Doneus, M., Löcker, K., Trinks, I., Nau, E., Pregesbauer, M., Kucera, M., Verhoeven, G.: The discovery of a gladiatorial school at Carnuntum. In: 10th International Conference on Archaeological Prospection. pp. 423–426. Austrian Academy of Sciences (2013)
26. Parish, Y.I., Müller, P.: Procedural modeling of cities. In: Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques. pp. 301–308. ACM (2001)
27. Penny, D., Pottier, C., Fletcher, R., Barbetti, M., Fink, D., Hua, Q.: Vegetation and land-use at Angkor, Cambodia: A dated pollen sequence from the Bakong temple moat. Antiquity 80(309), 599–614 (2006)
28. Pick, S., Kuhlen, T.W.: A framework for developing flexible virtual-reality-centered annotation systems. In: Proceedings of the IEEE Conference on Virtual Reality, SEARIS workshop (2015)
29. Roveda, V.: Khmer mythology secrets of Angkor Wat. Weatherhill Inc., Trumbull, Connecticut (1997)
30. Roveda, V., Poncar, J.: Sacred Angkor: The carved reliefs of Angkor Wat. Published and distributed by River Books, Bangkok (2002)
31. Stark, M.T., Evans, D., Rachna, C., Piphal, H., Carter, A.: Residential patterning at Angkor Wat. Antiquity 89(348), 1439–1455 (2015)
32. Tainturier, F.: Wooden architecture of Cambodia: A disappearing heritage. Center for Khmer Studies, Publishing Dept, Phnom Penh, Cambodia (2006)
33. Tyszczuk, R., Smith, J., Clark, N., Butcher, M.: ATLAS: Geography, architecture and change in an interdependent world. Black Dog Publishing (2012)
34. Valdelomar, J.T., Brandtner, J., Kucera, M., Wallner, M., Sandici, V., Neubauer, W.: 4D investigation of digital heritage an interactive application for the auxiliary fortress of Carnuntum. In: Digital Heritage, 2015. vol. 2, pp. 81–84. IEEE (2015)
35. Webb, M.E.: Computer-based modelling in school history. In: Martin, A.S.L., D, Y. (eds.) Information Technology and the Teaching of History International Perspectives, pp. pp 211–218. Harwood, Amsterdam (1997)
36. van Wijk, J., van Liere, R.: Hyperslice: visualization of scalar functions of many variables. In: Proceedings of the 4th Conference on Visualization. pp. 119–125 (1993)