Causality in Hospital Simulation Based on Utilization Chains

Gabriel Wurzer and Wolfgang E. Lorenz

Vienna University of Technology
Treitlstrasse 3 (Architectural Sciences / 259.1)
Vienna, Austria, 1040
{firstname.lastname}@tuwien.ac.at

Keywords: Agent-Based Simulation, Analysis, Debugging.

Abstract

The operation of complex buildings (e.g., airports, hospitals, industrial facilities, penitentiaries) is commonly simulated forward in time: agents arrive and perform their prescribed tasks, utilizing resources and space as required. When trying to understand the model's state at a certain point in time, say, "why is this resource over-utilized?", one must either guess or run the simulation again to determine what the cause is. Our contribution lies in the introduction of causal chains into the workflow of an agent-based simulation, so that an end user (in our case, process planner and hospital architect) can get a further insight into the intermediate simulation result at a certain point in time, without having to re-run the simulation.

1. INTRODUCTION

In the past half year, we have been simulating a large central operations theater (eight operation rooms minimally, the total number dependent on simulation output). With this work, we regularly had to hunt down errors and inconsistencies. Even when having full access to the model's code (*white box approach*), this was already quite a challenge, as there are many causes for an observed behavior (activities of an agent, signals, queuing and prioritization, etc.). For someone without access to the code (i.e. *black-box approach*, typically end-users), finding the erroneous spot might be even more troublesome: one could either guess (basically using *trial and error*) or resort to more sophisticated protocol-based methods such as (Rogsch and Klingsch 2011). Clearly, it is hard to argue for the trustworthiness of a model under such circumstances. Our major goal is thus to give end-users a greater flexibility in exploring a model during its development phase.

2. END-USERS NEED MORE THAN RESULTS

Process planners and hospital architects (who are our end-users) are not confident in simulation results lest they can understand how they were computed. For example (see the left side of Figure 1), they will want to see not only a utilization graph but also *dependencies between different resources*. It is common to shift this responsibility to the visualization (see the right side of Figure 1), meaning that one has run a simulation multiple times in order to understand the relationships between the patients and resources fully.

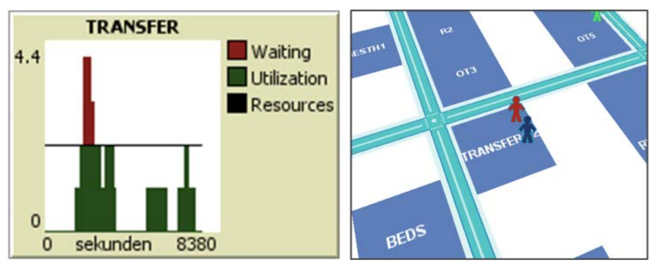

Figure 1. No causality shown in utilization graphs (left) and simulation view (right).

What if we could instead *stop the model* at a certain point in time and find out the *chain of events* that led to the current state, without having to restart? This is the core idea and main contribution of our paper, which we detail under Section 4. Summarizing briefly, we might then ask—for a specific space under scrutiny—the following questions:

- *Why* are resources utilized (causal chains, Section 4.1)?
- *Who* is using certain resources (role-centric, Section 4.2)?
- *When* are resources utilized (time-centric, Section 4.3)?

It is our hope that, using such functionality, users will be less tempted to call for "realistic visualizations" for proving the credibility of a simulation but can instead resort to proper argumentation.

3. BACKGROUND

3.1. Simulation Model

We use an agent-based model in which each agent follows a fixed sequence of functions to visit (that is, the *medical pathway*, which might consist of functions ARRIVAL > [HOLDING AREA] > PATIENT TRANSFER SYSTEM > OT > RECOVERY). This is similar to a schedule-calibrated occupant behavior simulation (Goldstein et al. 2010) or User Simulation of Space Utilisation (Tabak 2008). We use real (but anonymized) pathways exported by the Hospital Information System (see Wurzer et al. 2012; Glock et al. 2013 for details) or generate these by using "standard sequences" from which we eliminate steps based on probabilities, also determining step durations using min/max service times per function.

Functions form the dynamic aspect of our simulation: they are resources which can hold a finite number of agents at a time (*capacity*) and ultimately determine when an agent is able to proceed with its process. Thus, one might call our model a more discrete/client-server-based approach than agent-based approach, even though we also use agents for their ability to navigate in space (along the circulation, in the preliminary schema of the hospital, using a pedestrian movement model if required). Each function is situated in a space (e.g. OT1), which is the basic planning unit that our end users are after. There are two modes in which functions are acquired. First, an agent can utilize a function directly because it is in his process (*function as activity*). Second, a function is utilized *as queue* when another function is unavailable (e.g., preoperative holding area).

Histories are recorded for both agents and spaces: agents record at which time which function was sought or acquired, or a queue was entered, whereas spaces record which agents have visited or queued for a contained function. Using these basic recording mechanisms, we can later infer a causal chain.

3.2. Related Work

Techniques that can deliver the functionality we propose have already been available for some time in the context of debugging. An illustrative example of interrogative debugging techniques, (Ko and Myers 2004) have proposed a system that can interrogate a program over "why" or "why not" a certain program state was reached, by recording program states in between. In contrast to this approach, we need not "infer" causality (cf. Pearl 2009; Hluch 2014), as this is explicitly given in the medical pathway that each agent follows (see Section 3.1).

Visually, our approach uses Sankey diagrams (cf. Tufte 1983, p.176) to depict flow between spaces. A similar technique that depicts temporal events from medical records has been previously presented by Wongsuphasawat and Gotz (2012), albeit with no connection to the building layout on which our approach superimposes the visualization. Such an approach can be seen as Focus+Context technique (Card et al.1999) where a space under scrutiny is shown in full detail together with flows from other spaces leading there (giving context).

4. CONTRIBUTION

4.1. Why are Resources Utilized?

The answer to this question lies in all spaces that an agent has crossed before coming to a space under scrutiny. As we have recorded all utilizations for these spaces, we may thus speculate if these have led to a bottleneck (also refer to Figure 2). If agents are utilizing a space below its capacity, no queuing occurs (labeled green in the left part of Figure 2). If, on the other hand, agents queue for that space (labeled red, see again the left part of Figure 2), there is a potential for that space to have become a bottleneck.

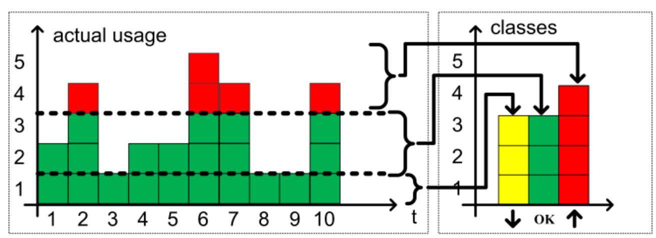

Figure 2. Utilization of a function. (left) Utilization with queuing depicted as red. (right) Classification into underutilized (yellow, <20% usage), well-utilized (20%-100%) and over-utilized (>100% =queuing).

The latter aspect deserves some attention, as there are two possible interpretations. Either we say that a space for which queuing occurs [even once!] can be the source of a bottleneck (*absolute utilization*). Or, we disregard minor queuing activities and focus on whether queueing occurs "most of the time" (*relative utilization*). In that context (refer to the right side of Figure 2), a space may be "under-utilized" (< 20% utilized), "well-utilized" (20-100%) or "over-utilized" (queueing occurs, thus >100%).

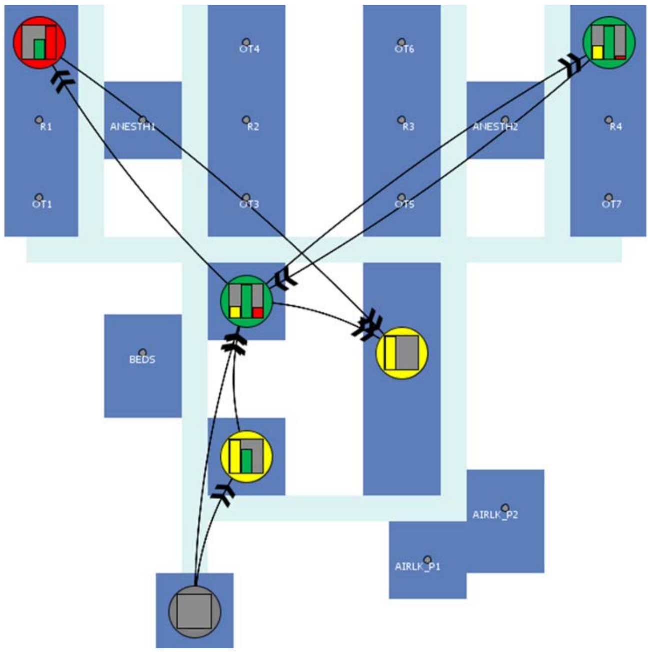

Figure 3. Causal chain depicting "why?" the space on the lower-right has been utilized (red = over-utilization, yellow = under-utilization, green = well-utilized). Actual screen shot from our simulation.

Figure 3 shows a utilization graph for the relative case. The overall utilization is shown by the color of each circle (yellow = under-utilized, green = well-utilized, red = over-utilized, inscribed into each space along the causal chain leading to the space under scrutiny). Under-, well- and over-utilizations are embedded into the same circle, allowing for an overview. Absolute utilization would be depicted in almost the same way, except that a space would turn red when a single over-utilization occurs (though this is only helpful in cases where resources are very scarce).

4.2. Who is Using Certain Resources?

To answer questions on utilization for a specific role (e.g. "are acute patients causing a bottleneck?"), a separation of causal chains *by role* is performed (refer to Figure 4). Colored edges now show patient type (red = acute, green = planned, red = day-clinic patients) and volume of patients crossing a space under scrutiny. As before, we show the relative utilization of the space as a circle (yellow = under-utilized, green = well-utilized, red = over-utilized), which we scale according to the patient volume. Large volumes of a specific patient type and visible utilization problems along the causal chain can help to pinpoint a bottleneck in the process for a certain group of building users.

4.3. When are Certain Resources Utilized?

This information is readily available in the utilization graphs (see the left of Figure 1); however, anchoring it to the chain of events may make patterns and dependencies in the temporal domain more obvious. We superimpose temporal utilizations over the causal chains (see Figure 5), by inscribing usage plots into circles showing relative utilization. A benefit of having such visualization is that planners can get a feeling for temporal dependencies without having to re-run the simulation.

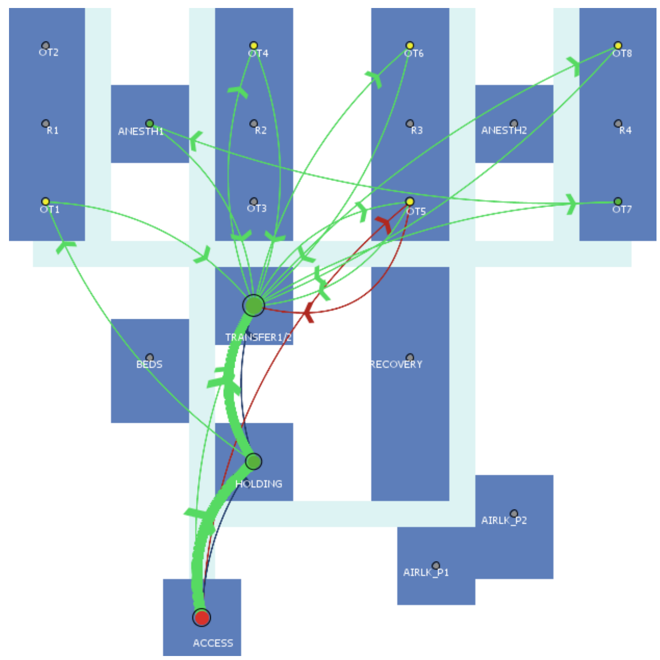

Figure 4. Utilization from a patient-centric perspective ("who?"). Patient flows depicted as colored edges (red = acute, green = planned, blue = day-clinic patients). Thickness of edges corresponds to patient volumes. Relative utilization of spaces shown as color-coded circles (red = over-utilization, yellow = under-utilization, green = well-utilized). Circle radius shows patient volume having crossed that space. Actual screen shot from our simulation.

5. DISCUSSION

Our approach is targeted at planners wishing to ask "why", "who" and "when" types of questions during the development of their model—either for analysis or for debugging purposes—leading to some confidence in the simulation results. The underlying model consists of agents visiting their respective medical pathways (or causal chains, as we would say). However, there is no need to have fixed sequences of visited functions as a basis for our approach -

it will just work fine with any other agent-based simulation for which agent histories and spatial utilization are recorded.

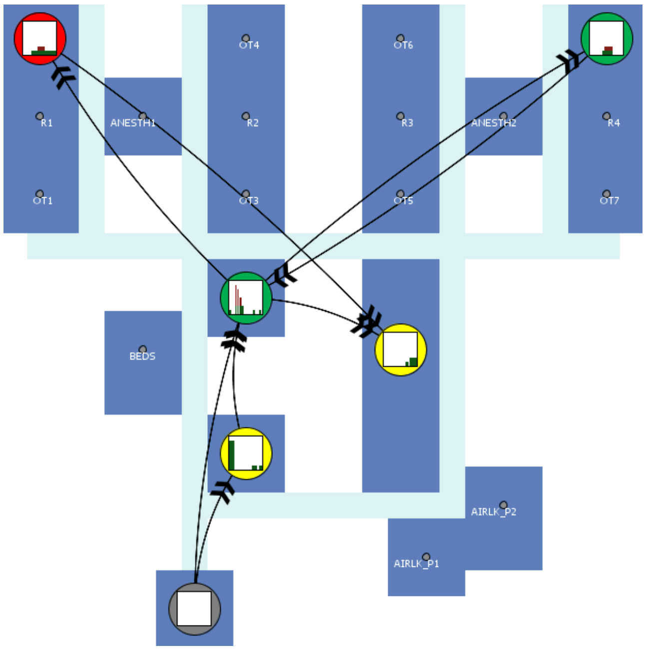

Figure 5. Superimposing utilization graphs over the utilization chain to show time-dependent causality. Actual screen shot of our simulation.

Another point worth discussing are the results of our underlying model (i.e., OT utilization, number of patients, room configuration and so on). We have intentionally left these out, since what we are interested in is *tool support* rather than a concrete case, even though our simulation was initially built with a 1100-bed clinic in Vienna in mind (but has since been extended so as to become more general-purpose OT planning software). Arguing further, we are only interested in preliminary phases of building design, where space programming takes place. In contrast, most papers in the same field target optimization of hospital units in the later phases of design, where concrete results can be given because the main decisions in space programming have already been made.

The chosen visualization types—"why", "who" and "when"—come from actually working with practitioners while developing our model. To be fair, we have yet no indication if our approach can be considered a "success" or "failure" in terms of aiding the design team for the general case. Proof or rejection of this statement would require us to test our model with one group using the visualization and one without—a task we have to leave for future work.

6. CONCLUSION

We have presented a visualization of utilization chains in agent-based models, which lets planners ask "why", "who" and "when" a certain space and its contained resources were utilized. Interrogating a model in this way can help users gain confidence and hunt down errors during the development of a simulation.

Acknowledgements

This work was sponsored in part by the ZIT MODYPLAN and the FFG STABLE HBO project. We wish to acknowledge fruitful discussions among the project members, especially with Anita Graser, Barbara Glock, Richard Wurglitsch, Niki Popper and Rudolf Linzatti.

References

CARD, S.K., MACKINLAY, J.D. AND SHNEIDERMAN, B. (Eds.) 1999. Readings in Information Visualization: Using Vision to Think. Morgan Kaufmann Publishers, San Francisco.

GLOCK, B., WURZER, G., BREITENECKER, F. AND POPPER, N. 2012. Reverse Engineering Hospital Processes Out of Visited Nodes. EUROSIM 2013, Cardiff. 312-317.

GOLDSTEIN, R., TESSIER, A. AND KHAN, A. 2010. Schedule-Calibrated Occupant Behavior Simulation. SimAUD 2010, Orlando. 79-86.

HLUCH, L. 2014. The Place of Causal Analysis in the Analysis of Simulation Data. Advances in Intelligent Systems and Computing. 245. 3.

KO, A.J. AND MYERS, B.A. 2004. Designing the Whyline: A Debugging Interface for Asking Questions about Program Failures. CHI 2004, Vienna. 6(1). 151-158.

PEARL, J. 2009. Causality: Models, Reasoning and Inference. Cambridge University Press, New York.

ROGSCH, C. AND KLINGSCH, W. 2011. To See Behind the Curtain - A Methodical Approach to Identify Calculation Methods of Closed-Source Evacuation Software Tools. PED 2010, Gaithersburg. 567-576.

TABAK, V. 2008. User Simulation of Space Utilisation. Ph.D. Thesis, Eindhoven University of Technology.

TUFTE, E.R. 1983. The Visual Display of Quantitative Information. Graphics Press, Cheshire.

WONGSUPHASAWAT, K. AND GOTZ, D. 2012. Exploring Flow, Factors, and Outcomes of Temporal Event Sequences with the Outflow Visualization. IEEE Transactions on Visualization and Computer Graphics. 18(12). 2659-2668.

WURZER, G., LORENZ, E.W. AND PFERZINGER, M. 2012. Pre-Tender Hospital Simulation using Naïve Diagrams as Models. I-WISH 2012, Vienna. 157-162.