Simulation-as-a-Service: Analyzing Crowd Movements in Virtual Environments
[Submitted title: Democratizing the Simulation of Human-Building Interactions]

Muhammad Usman, Brandon Haworth, Petros Faloutsos, Mubbasir Kapadia

2021

Democratizing the Simulation of Human-Building Interactions

Muhammad Usman*
York University
usman@cse.yorku.ca
www.eecs.yorku.ca/∼usman

Brandon Haworth
University of Victoria
bhaworth@uvic.ca
www.web.uvic.ca/∼bhaworth

Petros Faloutsos
York University
pfal@cse.yorku.ca
www.eecs.yorku.ca/∼pfal

Mubbasir Kapadia
Rutgers University
mubbasir.kapadia@rutgers.edu
www.cs.rutgers.edu/∼mk1353

Abstract

At present, environment designers mostly use their intuition and experience to predictively account for how environments might support dynamic activity. The majority of Computer-Aided Design (CAD) tools only provide a static representation of space which potentially ignores the impact that an environment layout produces on its occupants and their movements. To address this, computational techniques such as crowd simulation have been developed. With few exceptions, crowd simulation frameworks are often decoupled from environment modeling tools. They usually require specific hardware/software infrastructures and expertise to be used, hindering the designers' abilities to seamlessly simulate, analyze, and incorporate movement-centric dynamics into their design workflows. To bridge this disconnect, we devise a cross-browser service-based simulation analytics platform to analyze environment layouts with respect to occupancy and activity. Our platform allows users to access simulation services by uploading 3D environment models in numerous common formats, devise targeted simulation scenarios, run simulations, and instantly generate crowd-based analytics for their designs. We conducted a case study to showcase cross-domain applicability of our service-based platform, and a user study to evaluate the usability of this approach.

Keywords: simulation-as-a-service, crowd simulation, human-centric design, spatial analytics, building information modeling

1 Introduction

Analyzing how an environment layout impacts the movement and activities of its prospective inhabitants is a critical aspect of environment design process. While many established methods to evaluate environment performance, such as cost, structure, energy, and lighting mostly rely on static space representations, crowd simulation analytics account for the

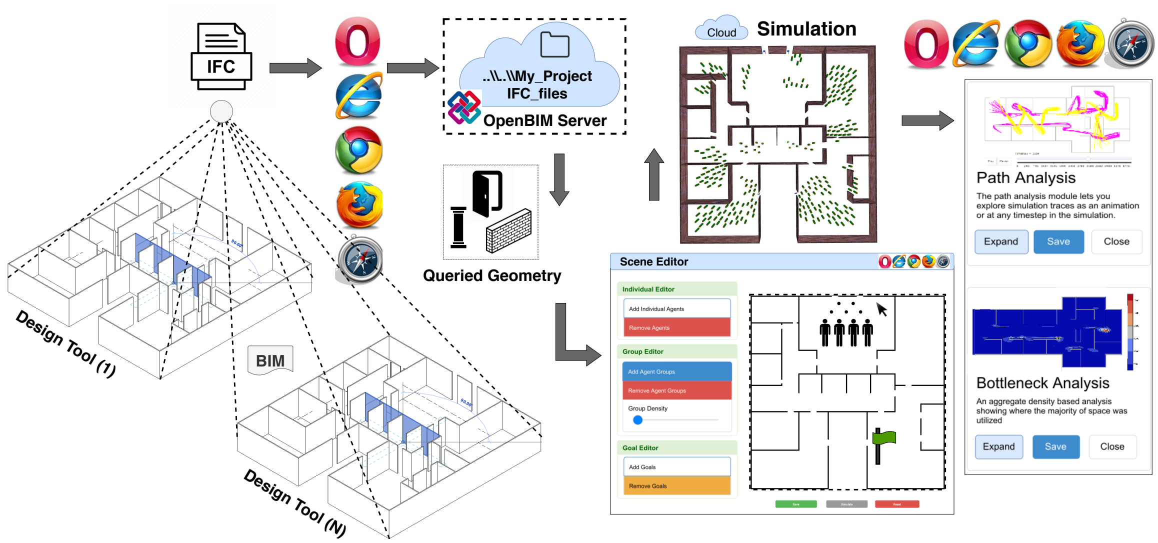

Figure 1: An overview of the proposed workflow. User uploads a 3D environment model (e.g. an IFC file) to the service, interactively sets up crowd configurations, runs the simulation and analyzes spatial visual and quantitative feedback for occupants–building interactions.

dynamic movement of people and their spatiotemporal impact on user experience, operational efficiency, and space utilization. Human-behavior simulation, however, presents high integration costs into environment design pipelines. Prior solutions demand deep expertise in a particular simulation platform and they require solving sophisticated interoperability challenges to import environment geometries, annotate spaces with semantics, define crowd behavioral parameters, generate simulation results, and visualize spatiotemporal data maps of space utilization.

To address these challenges, we adopt a Software-as-a-Service (SaaS) paradigm for software distribution and licensing using cloud computing [1]. This model has gained popularity in recent years and has several advantages both as a business model, but also for its users. It enables deep integration levels with other software in the work process to achieve targeted goals often in a cross-platform manner. In this way, specifically utilizing web-based or cloud services, allows tools to be used on-demand and across platforms without re-configuring core processes.

We develop and show how this approach may be used in a particularly challenging domain to support highly valuable processes. Often, firms which produce environment designs have a particular focus or set of focuses, which we will refer to as design domains. The design domains tend to be clearly defined by the environment uses, such as: outdoor urban settings, high-density housing, school/academic services, public services, commercial retail, industrial fabrication, warehousing, games, commercials, digital media, film, etc. Often the set of underlying tools used in the design process are the same or look the same. In early stages, procedural or prefabricated design processes are used to get numerous draft designs prepared quickly (e.g., using Grasshopper [2]). Within the design pipeline, advanced drafting tools are used to generate complete designs, retool, and reconfigure in an iterative process (e.g., using Autodesk Revit [3]). Finally, designs are converted to engineering blueprints and models using a variety of approaches (AutoCAD, 3D Printing, Model Services). In both, the beginning and middle stages of the pipeline, firms make use of their staple tools and configurations. These processes are often rigid and/or design domain specific. Adding tools to any one of these design domain processes can be difficult or prohibitive for a variety of reasons. A preliminary survey, including structured and unstructured questions, to experienced architects at three different firms with three different primary design domains revealed that tool adoption is very difficult and attrition is high for tools that do not seamlessly integrate into already existing pipelines, or are prohibitively expensive to do so (either monetarily, time expenditure, or acquiring expertise).

2 Literature Review

In recent years, crowd simulation has gained popularity in environment design processes. It enables designers to evaluate environment layouts from a human behavior perspective, and it supports advanced visualization of human-behavior metrics for enhanced communication and decision-making among design stakeholders. In this section, we summarize prior work on crowd simulation, crowd-aware environment design analysis, and recent developments of the software-as-a-service paradigm that can be used to provide a more seamless integration of crowd analytics into environment design pipelines.

2.1 Crowd Simulation

Crowd simulations provide time-based dynamics of potential human-building interactions in virtual environments. Research in this field has evolved quite rapidly over the last few decades [4]. Several simulation techniques have been proposed to imitate human-like movements [5, 6]. Particle-based simulation approaches revolve around computing local-level interactions for speed, motion and relative position of every individual to deliver large crowd behaviors [7]. Some approaches use social forces like repulsion and attraction to model individual occupant interactions [8, 9]. Hybrid rule-based approaches avoid future collisions in crowd [10]. An egocentric approach using affordances computes space-time plan for individual agents [11]. Reinforcement learning and deep learning have been used to model complex and more realistic crowd behaviors [12, 13].

2.2 Crowd-aware Environment Analysis and Design

In recent years, different approaches have focused on using crowd simulations to analyze environment layouts for making crowd-informed design decisions. Some efforts have been put in optimizing architectural elements to improve occupants' movement flow during evacuations [14]. Optimal egress routes are computed for built-environments using computational techniques which afford predictive egress planning [15]. Interactive methods are developed for rapidly optimizing small-scale virtual built-environments within user-selected parameters [16]. Interactive computational tools are developed to afford both static and dynamic crowd-aware analytics in built-environments [17]. In the area of parametric building design, a node-based framework was designed for the joint modelling of the built-environment, occupants' characteristics and the activities [18]. Following this work, another framework is presented to automatically explore, both individually and jointly, the building and crowd behavioral parameters, to design crowd-efficient building layouts [19]. These approaches, however, are often integrated into specific environment design workflow and require certain hardware/software infrastructures and expertise to be utilized by the general audience.

2.3 Software as a Service

Software-as-a-Service (SaaS) is a paradigm that is progressively gaining more traction in the industry because it separates the ownership, deployment, and maintenance of the software products from the end users (e.g. clients). This lets users utilize the software services on-demand by means of some client-side infrastructure (e.g. Application Program Interface (API), Web Interfaces, etc.) often via the internet [20]. A survey on modeling and simulation as a service discussed the advantages, limitations, and risks involved in using cloud-based simulation services–extracting the difference between Software and Simulation as a Service paradigms while noting the elasticity and ease of technical administration of the approach [21]. The work presented in [22] discusses cloud computing and virtualization platforms used for civilian and military modelling and simulation applications. Distributed architectures which use a model-driven engineering technique to extract geometric information of building model from CAD/BIM tools have been used as a remote service to run simulations and provides 3D visualization which can be visualized through an external third-party software tool (e.g. 3ds Max) [23]. In contrast, our platform is simulator agnostic in the sense that it uses a robust and modular underlying crowd simulation platform that specializes in continuous models. Allowing the user to choose what form they want their simulation to take. In another approach, a framework to model and simulate urban system simulations is presented on high performance cloud clusters [24]. A cloud-based framework is presented to remotely run simulations for studying the deployment of sensors in large facilities [25].

3 System Overview

In this section, we provide a detailed overview of our proposed cross-browser cloud-based simulation service to perform crowd-aware analytics for virtual and built environments. The system allows architects and designers (users) to upload 3D environment models, author simulation scenarios by setting crowd configurations, run simulations, and retrieve crowd-aware analytics for their designs. Figure 2 shows an overview of the interface of our simulation service. The basic interface of the proposed service allows users to visualize their environment designs in both 2D and 3D. The 2D visualization is an orthographic projection of the environment model (i.e., top–down view). The 3D visualization is a perspective projection of model from the top. Users can interact with their models by means of rotating and zooming around model's origin or using a fly through mode. A "Reset" functionality is also available to reset the camera to default view in 3D. The uploaded environments and their respective crowd-aware analytics generated by our service get saved to users' profile directories and can be accessed at a later time. Further details on individual functions is discussed below.

3.1 Environment Specification and Model Support

The system allows users to upload environments as Industry Foundation Classes (IFC), a BIM format. For an IFC, our system supports both IFC2x3 and $IFC4$ certifications (schema). The system does not limit users to use any particular environment design and modeling tool to generate their design models. Rather, IFCs can be sourced via any mainstream environment modeling platform. When an IFC file is uploaded, it is sent to an internal BIM model service and queried for geometric information of the model. The BIM service stores the model and sends back environment specifications (e.g. walls, doors, floors, etc.) to the system in an XML format.

3.2 Crowd Configuration and UI

The system allows users to define crowd configurations for design-specific simulation scenarios to run with their environment models. The "Scenario Editor" in Figure 2 summarizes a crowd configuration process. On the right is an environment layout of an exhibition space (e.g., art gallery). On the Left are the allowable actions a user can perform in the scenario editor. These include adding and removing individual occupants as well as occupancy groups, setting crowd-density levels (LoS) (i.e., number of occupants to spawn within an occupancy group), and adding and removing targets or goals for the occupants to walk to.

The Level of Service (LoS) provides a classification for pedestrian environments by assigning labels to different crowd densities [26]. Fruin provided these LoS classifications for different types of pedestrian environments, including stairs, queuing, and walkways. In this work, we selected the LoS for walkable areas (e.g., walkways). It has been used in traffic studies and crowd simulations to measure the quality of crowd movement (e.g., crowd flow) both for automotive and pedestrian applications. We summarized different LoS levels and their respective crowd density values in Table 1. An occupancy group is added by drawing a rectangle into the scene (Figure 2 – *Pink* region) and number of occupants to spawn within that group is calculated by multiplying area of that occupancy group region with selected crowd-density LoS level.

On double-clicking an individual occupant or an occupancy group shows a list of available targets in the current configuration. User can then select one or more targets from the list for the occupants (or occupancy groups). By default, the available targets are the exits (e.g., doors) in the environment. However, the platform does not refrain the users from removing/updating any of the existing targets or adding new ones. Any target or goal in an environment is a location in the space (e.g., point of interest) that the user wants the agents to visit during the simulation. For example, in a restaurant environment, a user-defined target could be a room representing a *Bar* for the agents to visit. Once a crowd configuration is created, user can then save it by selecting a "Save Config" action.

Like the environment models, the crowd configurations created by the users also get saved to users' profile directories and can be accessed at a later time. If the user uploads a new revision of an environment model for which the user already has created a crowd configuration, the user does not need to recreate the crowd scenario (unless the user wants

LoS Levels	Crowd Density	Interpretation
A	≤ 0.27	*Low Dense*
B	0.43 to 0.31	
C	0.72 to 0.43	*Medium Dense*
D	1.08 to 0.72	
E	2.17 to 1.08	*High Dense*
F	≥ 2.17	

Table 1: Level of Service (LoS) values and crowd density mapping from [26]. Density is measured in occupants per square meter.

to) and can reuse the existing crowd configuration to run simulations which the user has designed for an earlier version of that environment.

3.3 Crowd Simulation

We formally define a simulation scenario $S = \langle B, C \rangle$ which contains the specification of an environment layout B (e.g., geometric information like positions and attributes of walls, doors, floors, etc.) and the virtual crowd C. A virtual crowd $C = \langle C_I \cup C_G \rangle$ consists of individual occupants and occupancy groups, where $C_I = \langle O, A, P \rangle$ defines a collection of occupants O their desired activities A (e.g., egress evacuation, gathering in a meeting room, or other day-to-day scenarios) and optional parameters P of the underline crowd simulator being used, and $C_G = \langle O_G, A, P \rangle$ defines a collection of occupancy groups O_G their desired activities A and optional parameters P of the crowd simulator. For a similar set of activities, both individual occupants and occupants in occupancy groups, carry out the same behaviour during the simulation. If different activities are assigned to individual occupants, they will all carry out their own assigned behaviours during the simulation, which could be different from other individual occupants. The same, however, is not true for occupancy groups. All the occupants in an occupancy group carry out the same set of behaviours that are assigned to their occupancy group.

To run crowd simulations, our system uses an open platform, SteerSuite [27], which allows to run, analyze, and optimize crowd steering algorithms. Occupants are simulated using one of social forces (SF) [8], optimal reciprocal collision avoidance (ORCA) [28], predictive avoidance (PAM) [9], or footsteps planner (footstepsAI) [29, 30] models for steering behaviors (for demonstration we use SF in this paper). When a user selects a "Simulation" action, the system passes the current simulation scenario S with both environment specification B and crowd configuration C set by the user to SteerSuite (hosted on our cloud server) in an XML representation. The system can also generate a default evacuation scenario by analyzing the model to help users get started. Once the simulation is completed, occupant trajectories and other crowd-aware simulation statistics are send back to the user.

3.4 Simulation Feedback

After a simulation is completed, users may analyze their environment by selecting from an ever-expanding list of dynamic crowd analysis and visualization approaches. It allows users to visualize spatial quantitative and qualitative feedback from the simulation. Figure 3 shows occupants' trajectories (path analysis – Top) and density contours (bottleneck analysis – Bottom) respectively. The traces are shown in *Blue*, from starting position to the target, for all the occupants. To make the simulation experience intuitive for users, the system can also playback traces allowing the user to go back-and-forth in simulation timesteps with a help of a slider. The heatmap is a color-coded representation of an average occupant density per square meter, calculated for whole design space of the environment and for all the occupants, over the course of simulation. *Red* regions show areas of high density (e.g. potential bottlenecks), whereas *Blue* shows less dense areas.

The system also reports simulation statistics as quantitative numbers. These include minimum, maximum and average evacuation times and traveled distances over the course of simulation, as well as an average *Exit Flow* of occupants. The exit flow is calculated by dividing average evacuation time with total number of occupants completed the simulation.

4 Case Study

This section presents a series of case studies to demonstrate the effectiveness of the proposed simulation service platform in three different design domains, with models sources from three different pipelines, showing how the SaaS approach can be effectively and seamlessly used to inform decision making in the environment modeling pipeline.

4.1 Eatery Design

An eatery layout whether being designed for a restaurant, a food court or a cafeteria, has to comply with numerous applicable codes including accessibility, flow, and egress. For an egress, however, accounting for potential human–building interactions for future inhabitants is of vital importance.

In this use case, we demonstrate how the presented workflow can be used to analyze crowd dynamics of potential human–building interactions for two completely different simulation scenarios. A real-world restaurant style environment is created using Autodesk Revit. Figure 4 shows analytics for a restaurant environment for an emergent egress evacuation and a group dine-in scenario. For an egress evacuation (top row), using our service controls, we interactively added virtual customers in different spaces of restaurant with an objective (e.g. target) to move towards nearest exit. Crowd trajectories are shown in Blue, highlighting the paths virtual customers followed while moving towards exits. The color-coded heatmap highlights the bottlenecks in space which appeared during the evacuation, providing visual insights on potential human-safety hazards. For group dine-in scenario (bottom row), we added two different groups of virtual customers, entering the restaurant from different entrances, waiting in the lobby to be attended by a receptionist, moving to the bar, dining-in in the main dinning hall, going to the bathroom, visiting the manager, and heading back towards exits. Crowd trajectories are shown in different colors for each group to differentiate their activities and the paths they followed along them. The heatmap shows potential bottlenecks at the bar entrance and in the lobby. Average exit flow, traveled distances, and evacuation times are also shown in the figure.

4.2 Exhibition Design

A real-world exhibition style environment (e.g. an art gallery) is created using Rhinoceros. Figure 5 shows the analytics for an egress evacuation and a group-based exhibition exploration scenario. For egress evacuation (top row), we interactively added virtual visitors at different exhibit points in the art gallery with an objective to move towards nearest exit. Path analysis reveals that the obstacle in the middle hallway towards left-side, helped in forming multi-lane in the left-side of the gallery. The heatmap shows bottleneck in the middle hallway towards right-side of the gallery near the exit. These analyses showcased that a designer might want to consider adding an obstacle in the hallway towards right-side of the gallery as well, to help formation of lanes for egress, or make other design improvements accordingly. For group-based exhibition exploration (bottom row), we added two different groups to explore the gallery from one exhibit point to another, making stops, and then moving to the next.

4.3 Workplace Design

A workplace environment (e.g., an office) is created using SketchUp. Figure 6 shows analytics for an egress and a daily work-routine scenario of two different teams. In the egress scenario (top row), virtual employees are added in different spaces in office with an objective to move towards nearest exit. Path and bottleneck analyses are presented. The heatmap reveals multiple bottlenecks in the hallways near meeting rooms and cafeteria. For the daily routine scenario, we added two different teams in different spaces in the office and showed their work-routine activities including attending meetings, going to cafeteria, etc. Several bottlenecks in space are revealed in the heatmap.

Note. In the presented case studies, all the environments represent a single-storey layout. This is because, with the current work, we aimed to lay down the foundations (e.g., a preliminary working platform) towards democratizing the simulation of human–building interactions, to address the complexities and costs related to dynamic movement simulations to facilitate its adoption in the game, environment design and urban planning communities. The support of simulating multi-storey environments is one of the future directions we are currently working on.

Count	Mean	Median	Standard Deviation
6	72.5	72.5	4.85

Table 2: A summary of results from 10 participants for SUS, where the score range is from 0 to 100.

5 System Usability

A pilot user study is presented to evaluate the usability of our system. Ten (10) senior-level graduate students voluntarily participated in the experiment. All the participants reported prior experience with CAD tools to analyze building structures. Participants were tasked with using the proposed simulation service to author crowd configurations in a given residence environment (e.g., a house), and analyze the space for human occupancies. All participants used the system for a fixed amount of time (20 minutes). The participants were allowed to interact with the platform as they like while completing the crowd authoring task. For example, they were allowed to add agents, agent groups, and targets, or remove any previously added agents or goals. Afterwards, participants completed a System Usability Scale (SUS) [31] survey which is an established method in the literature to evaluate the usability of a system, and can be scaled to the range of 10 to 100, with a score higher than 68 to be considered above average and admissible [32]. SUS score is a compound measure of usability for a system which has been proved to be reliable. The summary of SUS scores are reported in Table 2. Our mean (72.5) and median (72.5) scores from SUS fall within the adjective range of “good” and “excellent” [33].

6 Conclusion

We present a cross-browser cloud-based crowd analytics framework for crowd-aware environment design. Eliminating all the hardware/software infrastructure dependencies, a single solution is proposed to bring the environment layouts from different environment modeling design tools (e.g. Autodesk Revit or Rhinoceros) into an interactive crowd authoring workspace, to set up design-specific crowd scenarios, remotely run simulations, and analyze feedback. A series of case studies are presented to showcase the effectiveness of this approach by analyzing environment designs for different domains with respect to human-occupancy. The usefulness of the system is evaluated with a system usability study (SUS) where participants rated their confidence as “above average”. We believe that the current system operates as a novel instrument for evaluating hypotheses in the SaaS and collaborative design using dynamic human analyses domains. Future research will involve adding support to allow the users to alter their designs within the web-browser workspace. As well, future work will include a comprehensive user study with novices and experts from the architecture and gaming community to evaluate this approach with more complex design and crowd behavioral tasks.

Acknowledgements

This research has been partially funded by grants from the ISSUM, OGS, and in part by NSF IIS-1703883 and NSF S&AS-1723869.

References

[1] Mark Turner, David Budgen, and Pearl Brereton. Turning software into a service. *Computer*, 36(10):38–44, 2003.

[2] Grasshopper. Retrieved: 2020-01-10. `https://www.rhino3d.com/6/new/grasshopper`.

[3] Inc. Autodesk. Revit. Retrieved: 2019-12-15. `https://www.autodesk.ca/en/products/revit/`.

[4] Daniel Thalmann and Soraia Raupp Musse. Crowd rendering. In *Crowd Simulation*, pages 195–227. Springer, 2013.

[5] Mubbasir Kapadia, Nuria Pelechano, Jan Allbeck, and Norm Badler. Virtual crowds: Steps toward behavioral realism. *Synthesis lectures on visual computing*, 7(4):1–270, 2015.

[6] Nuria Pelechano, Jan M Allbeck, Mubbasir Kapadia, and Norman I Badler. *Simulating heterogeneous crowds with interactive behaviors*. CRC Press, 2016.

[7] Craig W Reynolds. *Flocks, herds and schools: A distributed behavioral model*, volume 21. ACM, 1987.

[8] Dirk Helbing and Peter Molnar. Social force model for pedestrian dynamics. *Physical review E*, 51(5):4282, 1995.

[9] Ioannis Karamouzas, Peter Heil, Pascal Van Beek, and Mark H Overmars. A predictive collision avoidance model for pedestrian simulation. In *MIG*, pages 41–52. Springer, 2009.

[10] Shawn Singh, Mubbasir Kapadia, Billy Hewlett, Glenn Reinman, and Petros Faloutsos. A modular framework for adaptive agent-based steering. In *I3D*, pages PAGE–9. ACM, 2011.

[11] Mubbasir Kapadia, Shawn Singh, William Hewlett, and Petros Faloutsos. Egocentric affordance fields in pedestrian steering. In *I3D*, pages 215–223. ACM, 2009.

[12] Nicolas Heess, Srinivasan Sriram, Jay Lemmon, Josh Merel, Greg Wayne, Yuval Tassa, Tom Erez, Ziyu Wang, SM Eslami, Martin Riedmiller, et al. Emergence of locomotion behaviours in rich environments. *arXiv preprint arXiv:1707.02286*, 2017.

[13] Xue Bin Peng, Glen Berseth, KangKang Yin, and Michiel Van De Panne. Deeploco: Dynamic locomotion skills using hierarchical deep reinforcement learning. *TOG*, 36(4):41, 2017.

[14] Glen Berseth, Muhammad Usman, Brandon Haworth, Mubbasir Kapadia, and Petros Faloutsos. Environment optimization for crowd evacuation. *CAVW*, 26(3-4):377–386, 2015.

[15] Vincius J Cassol, Estêvão Smania Testa, Cláudio Rosito Jung, Muhammad Usman, Petros Faloutsos, Glen Berseth, Mubbasir Kapadia, Norman I Badler, and Soraia Raupp Musse. Evaluating and optimizing evacuation plans for crowd egress. *CGA*, 37(4):60–71, 2017.

[16] Brandon Haworth, Muhammad Usman, Glen Berseth, Mahyar Khayatkhoei, Mubbasir Kapadia, and Petros Faloutsos. Code: Crowd-optimized design of environments. *CAVW*, 28(6):e1749, 2017.

[17] Muhammad Usman, Davide Schaumann, Brandon Haworth, Glen Berseth, Mubbasir Kapadia, and Petros Faloutsos. Interactive spatial analytics for human-aware building design. In *MIG*, page 13. ACM, 2018.

[18] Muhammad Usman, Davide Schaumann, Brandon Haworth, Mubbasir Kapadia, and Petros Faloutsos. Joint parametric modeling of buildings and crowds for human-centric simulation and analysis. In *CAAD Futures*, pages 279–294. Springer, 2019.

[19] Muhammad Usman, Davide Schaumann, Brandon Haworth, Mubbasir Kapadia, and Petros Faloutsos. Joint exploration and analysis of high-dimensional design–occupancy templates. In *MIG*, page 35. ACM, 2019.

[20] Phillip A Laplante, Jia Zhang, and Jeffrey Voas. Distinguishing between software oriented architecture and software as a service: What's in a name? 2008.

[21] Erdal Cayirci. Modeling and simulation as a cloud service: a survey. In *Winter Simulation*, pages 389–400. IEEE Press, 2013.

[22] Erdal Cayirci and Chunming Rong. Intercloud for simulation federations. In *High Performance Computing & Simulation*, pages 397–404. IEEE, 2011.

[23] Sixuan Wang and Gabriel Wainer. A simulation as a service methodology with application for crowd modeling, simulation and visualization. *Simulation*, 91(1):71–95, 2015.

[24] Daniel Zehe, Wentong Cai, Alois Knoll, and Heiko Aydt. Tutorial on a modeling and simulation cloud service. In *Winter Simulation*, pages 103–114. IEEE, 2015.

[25] Rafael Pax, Jorge J Gomez-Sanz, Ismael Sagredo Olivenza, and Marlon Cárdenas Bonett. A cloud based simulation service for 3d crowd simulations. In *Distributed Simulation and R. Applications*, pages 1–5. IEEE, 2018.

[26] John J Fruin. *Designing for pedestrians: A level-of-service concept*. Number HS-011 999. 1971.

[27] Shawn Singh, Mubbasir Kapadia, Petros Faloutsos, and Glenn Reinman. An open framework for developing, evaluating, and sharing steering algorithms. In *MIG*, pages 158–169. Springer-Verlag, 2009.

[28] Jur Van Den Berg, Stephen J Guy, Ming Lin, and Dinesh Manocha. Reciprocal n-body collision avoidance. In *Robotics research*, pages 3–19. Springer, 2011.

[29] Shawn Singh, Mubbasir Kapadia, Glenn Reinman, and Petros Faloutsos. Footstep navigation for dynamic crowds. *Computer Animation and Virtual Worlds*, 22(2-3):151–158, 2011.

[30] Glen Berseth, Mubbasir Kapadia, and Petros Faloutsos. Robust space-time footsteps for agent-based steering. *Computer Animation and Virtual Worlds*, 2015.

[31] John Brooke. Sus: a retrospective. *Journal of usability studies*, 8(2):29–40, 2013.

[32] Jeff Sauro and James R Lewis. When designing usability questionnaires, does it hurt to be positive? In *SIGCHI*, pages 2215–2224. ACM, 2011.

[33] Aaron Bangor, Philip Kortum, and James Miller. Determining what individual sus scores mean: Adding an adjective rating scale. *Journal of usability studies*, 4(3):114–123, 2009.

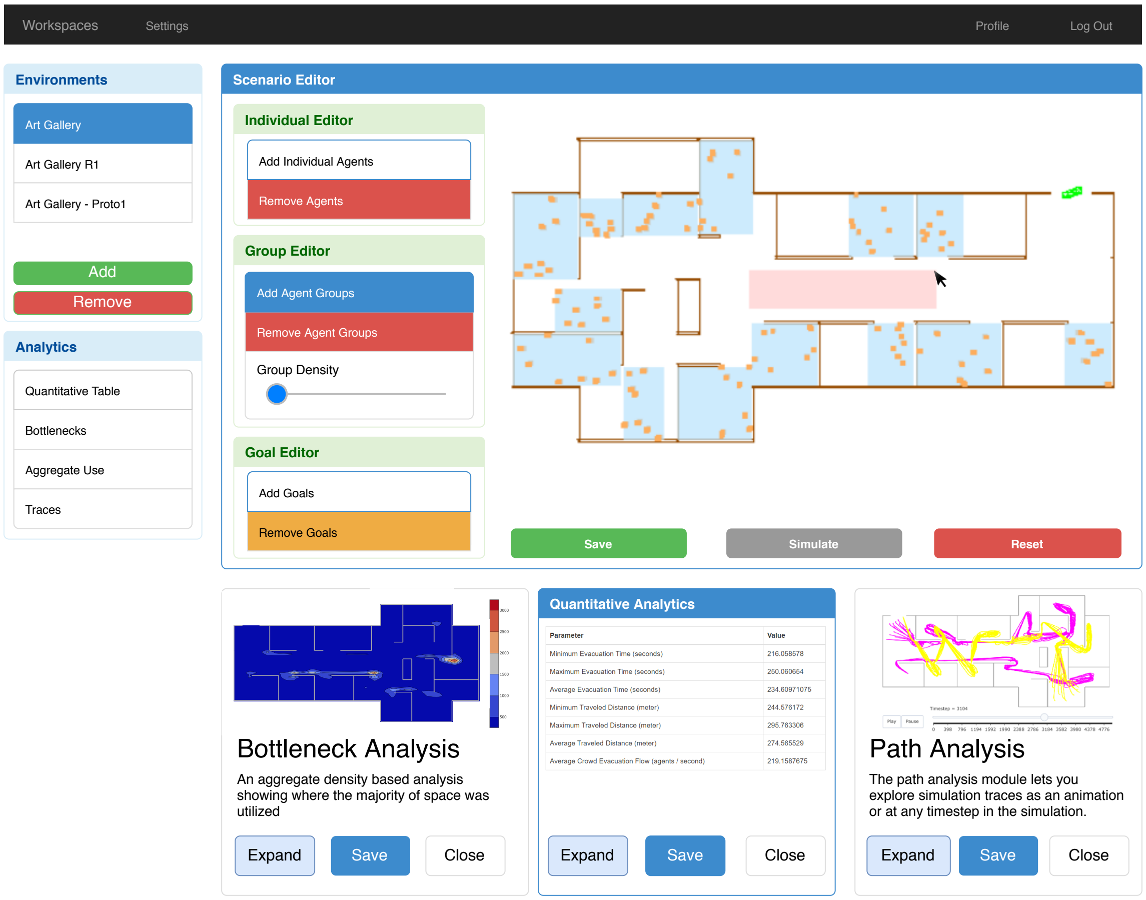

Figure 2: An overview of the workspace interface. It allows the user to layout, store, and retrieve workspaces for projects. Within the workspace, users can design scenarios, run simulations, and create analytics portfolios. In the scenario editor, the user is designing an egress scenario for an exhibition space. The goal for the agents is set at the entrance of the environment. As the simulation will start, the agents will start moving towards the goal (e.g., exit) to vacate the environment.

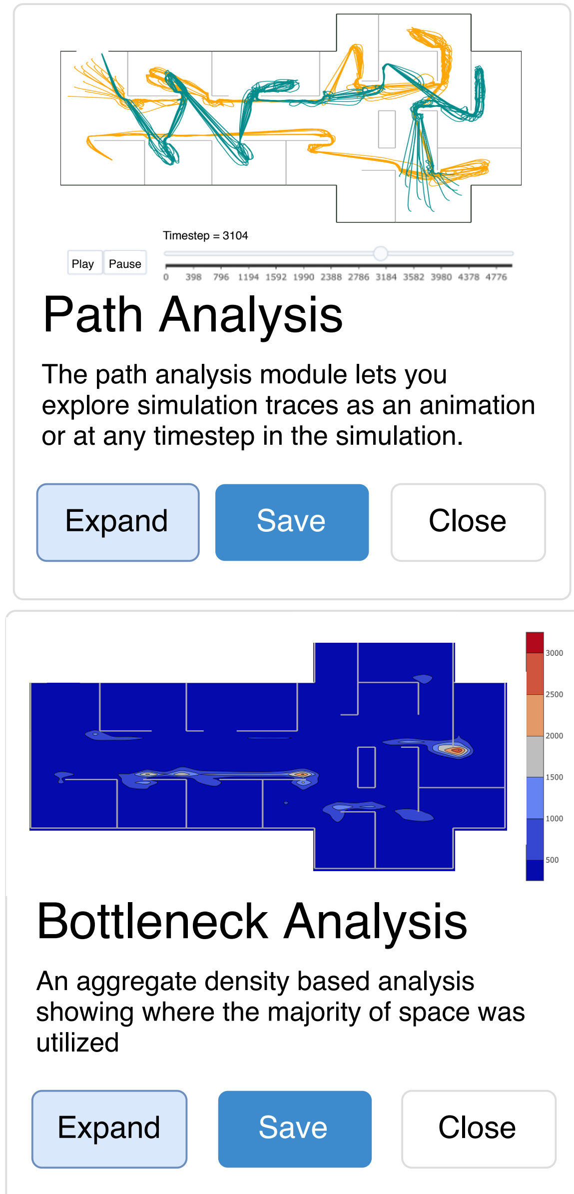

Figure 3: Qualitative tools afford quick exploration of simulation results and problems areas. Our bottleneck analysis thresholds aggregate occupancy maps to bring focus to various types of flow bottlenecks in designs.

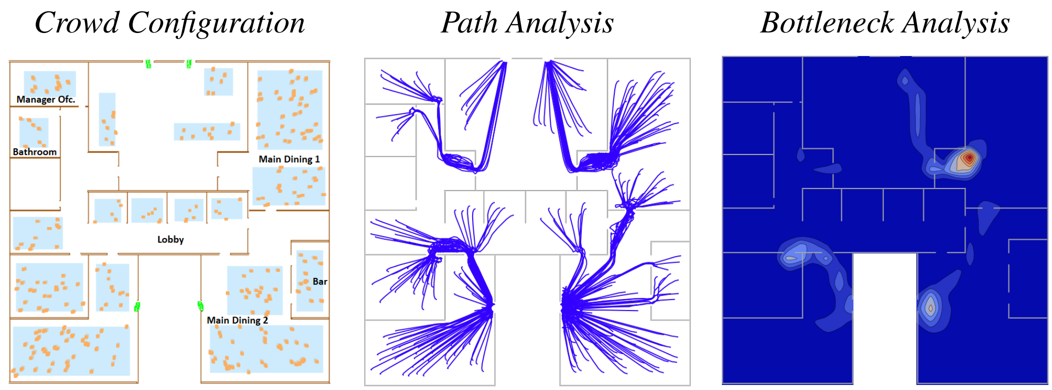

Evacuation: **Flow** = 19.31agents/s, **Distance** = 20.76m, **Time** = 19.75s

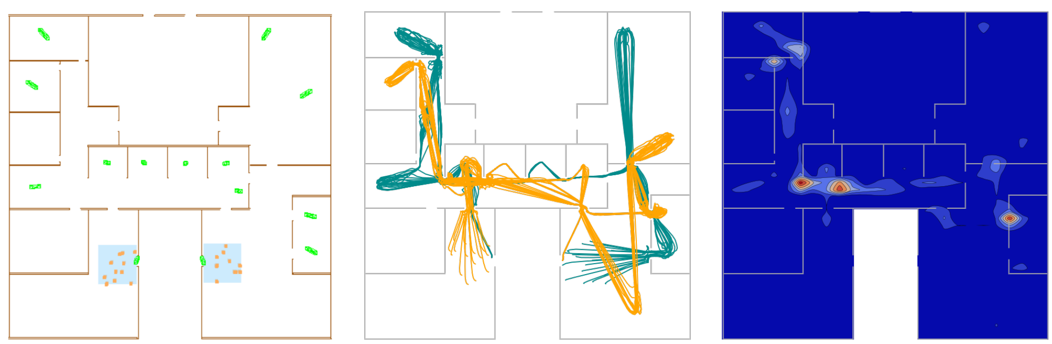

Group Dining: **Distance** = 198.7m — **Time** = 182.0s

Figure 4: Crowd analytics for an eatery environment (e.g. a restaurant). Two different scenarios are presented: Top – an egress evacuation where customers from different spaces in the restaurant moving towards nearest exit, and Bottom – a group dining where two different group people come to the restaurant, wait in the lobby, go to the bar, dine-in, go to the washroom and leave. In the left column, the crowd configuration is shown with spawn region in light blue and goals in green. Crowd trajectories are shown in *Blue* for egress and multi-colored for the group dining scenario. Crowd-density analysis is shown as color-coded heatmap (Red–Blue) where denser crowd areas (bottlenecks) are highlighted in dark red.

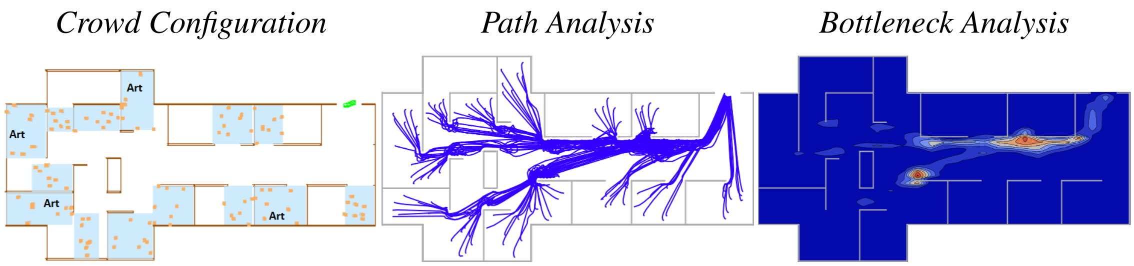

Evacuation: **Flow** = 27.45agents/s — **Distance** = 37.03m — **Time** = 36.12s

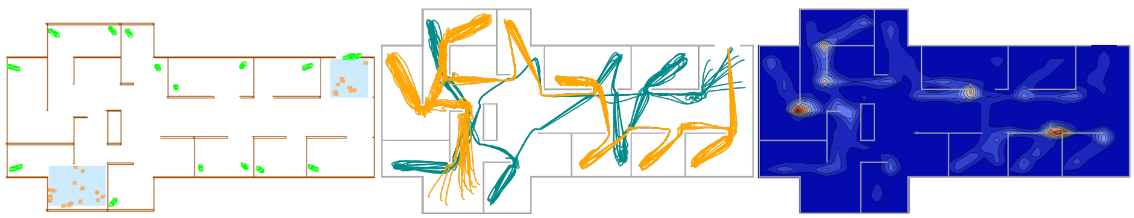

Group Exploration: **Distance** = 209.8m — **Time** = 208.6s

Figure 5: Crowd analytics for an exhibition environment (e.g. an art gallery). Two different scenarios are presented: Top – an egress evacuation where visitors from different spaces in the gallery moving towards nearest exit, and Bottom – a group exploration where two different group people exploring the gallery from one exhibit point to another. Crowd trajectories are shown in *Blue* for egress and multi-colored for group exploration scenario. Crowd-density analysis is shown as color-coded heatmap (Red–Blue) where denser crowd areas (bottlenecks) are highlighted in dark red. Crowd exit flow for egress evacuation, and average evacuation time (**T**) and traveled distance (**D**) for both evacuation and group dining scenarios are also reported. *Green* cylinders are the targets for agents, whereas *Orange* represent group agents.

Crowd Configuration *Path Analysis* *Bottleneck Analysis*

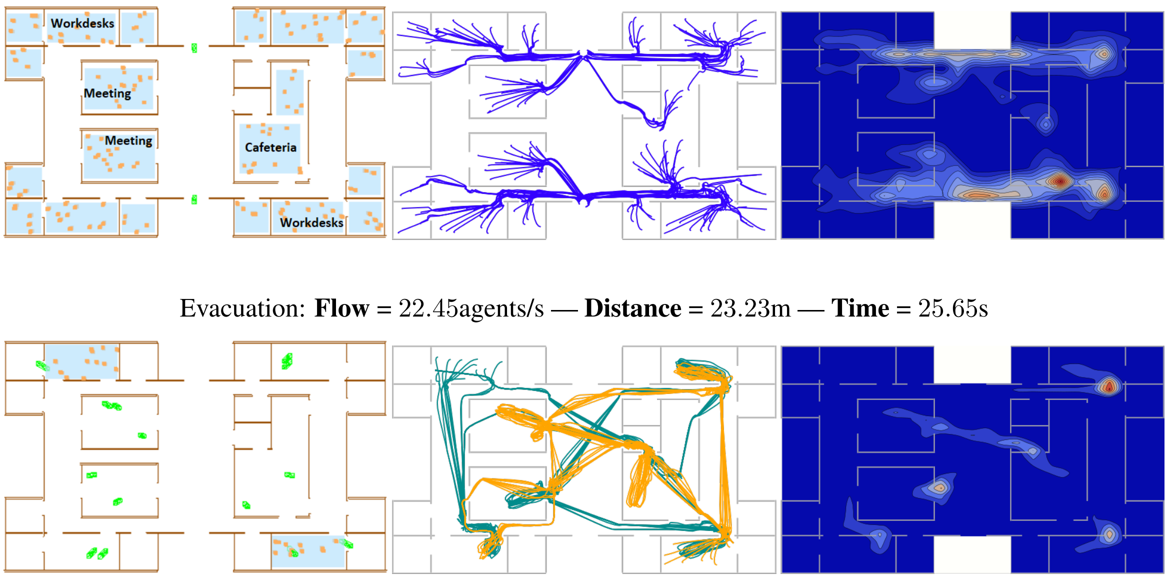

Evacuation: **Flow** = 22.45agents/s — **Distance** = 23.23m — **Time** = 25.65s

Work Routine of Teams: **Distance** = 235.8m — **Time** = 224.3s

Figure 6: Crowd analytics for a corporate work environment (e.g. an office). Two different scenarios are presented: Top – an egress evacuation where employers from different spaces in the office moving towards nearest exits, and Bottom – a daily work routine of two different teams attending meetings, going to cafeteria and visiting colleagues.