

Evaluating and Optimizing Level of Service for Crowd Evacuations

Brandon Haworth*
York University

Muhammad Usman†
York University

Glen Berseth‡
University of British Columbia
Petros Faloutsos¶
York University

Mubbasir Kapadia§
Rutgers University

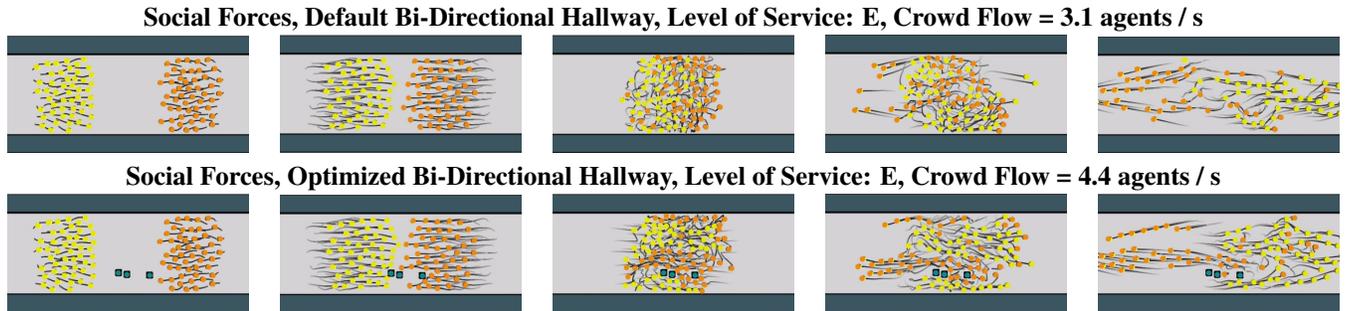


Figure 1: Simulation of social forces steering algorithm in a bi-directional hallway for different levels of service. Optimal pillar placements produce emergent lanes and increase the effective critical density, thereby increasing the effective level of service at higher crowd densities.

Abstract

Level of service (**LoS**) is a standard indicator, widely used in crowd management and urban design, for characterizing the service afforded by environments to crowds of specific densities. However, current **LoS** indicators are qualitative and rely on expert analysis. Computational approaches for crowd analysis and environment design require robust measures for characterizing the relationship between environments and crowd flow.

In this paper, the flow-density relationships of environments optimized for flow under various **LoS** conditions are explored with respect to three state-of-the-art steering algorithms. We optimize environment elements to maximize crowd flow under a range of density conditions corresponding to common **LoS** categories. We perform an analysis of crowd flow under **LoS** conditions corresponding to the **LoS** optimized environments. We then perform an analysis of the crowd flow for these **LoS** optimized environments across **LoS** conditions.

The steering algorithm, the number of optimized environment elements, the scenario configuration and the **LoS** conditions affect the optimal configuration of environment elements. We observe that the critical density of crowd simulators can increase, or shift **LoS**, due to the optimal placement of pillars. Depending on the steering model and environment benchmark, pillars are configured to produce lanes or form wall-like structures, in an effort to maximize crowd flow. These experiments serve as a precursor to environment optimization and crowd management motivating the need for further study using real and synthetic crowd datasets across a larger representation of environments.

CR Categories: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Animation; J.5 [Computer Applications]: Arts and Humanities—Architecture;

*e-mail:brandon@cse.yorku.ca

†e-mail:usman@cse.yorku.ca

‡e-mail:gberseth@cs.ubc.ca

§e-mail:mubbasir.kapadia@gmail.com

¶e-mail:pfal@cse.yorku.ca

Keywords: crowd simulation, crowd evacuation, steering algorithms, level of service, environment optimization

1 Introduction

The evaluation of pedestrian and agent-based crowd dynamics is of importance to several communities including architectural design, urban planning, computer graphics, and game design. Crowds of varying densities exhibit a number of behaviours in different contexts. In particular, bottleneck and evacuation scenarios, in which crowd members are compressed together through environmental and contextual features, produce interesting aggregate behaviours. For instance, panicked or evacuating crowds can produce turbulent behaviours such as trampling or reduced-to-blocked flow scenarios. Environment and level designers may attempt to characterize and study these scenarios through a variety of means.

Traffic and pedestrian dynamics communities often use a standard qualitative classification, Level of Service (**LoS**), to describe flow relationships with respect to volume, or density, in an environment. [Fruin 1971] provides the original descriptions of **LoS** definitions for pedestrian environments. This classification system provides an easy to use means of describing flow contexts based mainly on observations of typical crowd contexts. Using this system as a guide, a designer can attempt to model key aspects of the environment to meet a desired **LoS**. For example, a stairwell may be designed to accommodate **LoS C** capacity conditions on average. Although **LoS** can work well as a qualitative metric for the design and classification of environments, it lacks empirical evidence that generalizes across crowd types, contexts, and environment configurations.

Synthetic crowds can be simulated to extrapolate crowd flow for rigorous quantitative analysis, but these may be at odds with the

qualitative **LoS**. In order to successfully use computer-assisted design, it is important to have sound quantitative metrics that can be used as the basis for analysis and optimization. These should relate to the widely accepted **LoS** classification leading to effectively chosen design parameters.

In this paper, we perform environment optimization experiments across various **LoS** conditions using established agent-based crowd simulation techniques. We explore optimal environment configurations as they pertain to crowd densities and environmental contexts. These experiments reveal a sensitivity of flow-density relationships to local steering algorithms in optimal results. We show that **LoS** conditions in different optimal contexts produce both agreeable and disagreeable flow density relationships dependant on both the steering model and environment.

This paper makes the following contributions: An empirical study of flow-density relationships for steering simulators and benchmarks in the context of environment optimization. Our results provide insight into how **LoS** classification differs across steering simulators in different evacuation contexts, therefore it is unclear whether there is a standard **LoS** characteristic across synthetic crowd behaviours. This motivates the need for further experimentation and field observations.

2 Related Work

Crowds exhibit interesting macroscopic effects depending on context, particularly in high density situations. [Still 2007] reviews the complexity and difficulty involved in the simulation of crowds in these critical situations. However, the results of a large scale empirical study presented in [Karamouzas et al. 2014] show a simple power law describes real-world pedestrian interactions in a variety of conditions.

2.1 Crowd Simulation

Crowd simulation through local steering began with the seminal work [Reynolds 1987] describing the modelling of flocking and grouping behaviours. Since, numerous approaches have proven useful for solving the steering problem in a variety of contexts. These approaches include: force-field based [Karamouzas et al. 2009; Helbing et al. 2000] which model interactions as physical forces; geometrically-based [van den Berg et al. 2011; Guy et al. 2009] which decompose velocity-space and may include synthetic perceptions [Ondřej et al. 2010]; and multi-phase [Singh et al. 2011] which integrate multiple steering models to generate a final steering decision.

Approaches that take into account the nature of density and crowd size have been developed as well. The simulation of large scale dense crowds proposed in [Narain et al. 2009] is made possible by decoupling local collision avoidance through introducing unilateral incompressibility in the continuous domain. Generating realistic crowd simulations based on relating the common speed/density relationship of fundamental diagrams to dense simulated crowds was proposed in [Narang et al. 2015; Best et al. 2014].

2.2 Crowd Dynamics and Evaluation

Crowd dynamics and traffic flow has a long and rich history. In particular, pedestrian traffic flow exhibits unique dynamics at different densities which must be planned for, originally discussed in [Fruin 1971].

Evaluation of these dynamics and the correctness of their simulations is of particular importance and has spawned different ap-

proaches to crowd analysis. Fundamental diagrams were used as a measure of aggregate similarity in [Seyfried et al. 2010] which, by performing a large empirical study, highlights the importance and discrepancies arising from measurement. [Lerner et al. 2010; Lerner et al. 2009] proposed measuring similarity to real world data through a data-driven approach, highlighting the effects of density under different contexts. [Singh et al. 2009] proposed simulation of crowds over a series of challenging test cases using an extensible set of evaluation metrics for measuring the performance of steering algorithms. This work was extended in [Kapadia et al. 2011] to simulate crowds in a representative set of challenging test cases and compute the coverage of a steering algorithm and resultant quality of simulations.

2.3 Crowd and Environment Optimization

Recent work [Wolinski et al. 2014; Berseth et al. 2014b] has demonstrated the potential of automatically fitting the parameters of crowd simulation techniques to meet different performance criteria, such as maximizing crowd flow.

Environment optimization spans multiple domains wherein a specific goal is related to the best parameters of the environment by some criteria function. Most related to our work is the optimization of crowd affecting environment parameters. Optimization of features such as pillars has been studied in [Rodriguez et al. 2013; Jiang et al. 2014] for a single steering algorithms. In [Berseth et al. 2014a], static path analysis was proposed to procedurally generate difficulty optimal game level layouts. Recent literature shows that choice of steering algorithm affects the output of environment optimizations [Berseth et al. 2015].

Comparison to Prior Work. We complement prior work in crowd simulation, evaluation, and optimization by examining the sensitivity of agent-based steering algorithms and optimization results with respect to flow-density relationships in crowded environments.

3 Methodology

In this section we outline the concept of **LoS**, steering algorithms, and specific environments used in the analysis.

3.1 Level of Service of Pedestrian Crowds

Levels of Service (**LoS**), as defined in [Fruin 1971], provides a classification of the objective **LoS** for pedestrians environments by assigning labels to crowd densities. These **LoS** classifications are provided for various types of pedestrian environments such as walkways, queueing, and stairs. Here we are interested in walkways. [Fruin 1971] provides the original measures and descriptions used to classify **LoS** for pedestrian walkway environments.

3.2 Steering Algorithms

We perform our analysis using the following steering algorithms as they represent a range of different local agent-based steering approaches in large crowd simulation: (a) **ORCA**: a reciprocal velocity obstacles based approach [van den Berg et al. 2011], (b) **PPR**: a hybrid, rule based approach combining reactions, predictions, and planning [Singh et al. 2011], and (c) **SF**: a social forces based approach [Helbing et al. 2000]. For each algorithm the default parameters, as suggested by the algorithm's developers, are used.

3.3 Environment Configuration

The environment configuration refers to the starting arrangement or design of obstacles and agents in an environment or level. Formally we can define an environment configuration similar to [Berseth et al. 2013], as $s = \langle \mathbf{O}, \mathbf{A} \rangle$, where \mathbf{O} , \mathbf{A} are the sets of static obstacles and agents in the scenario respectively. For our experiments, an obstacle $o \in \mathbf{O}$ is either a cylindrical pillar or a rectangular bounding box, but may be extended to arbitrary polygonal objects. An agent $a \in \mathbf{A}$ is defined as $a = \langle \mathbf{x}, r, \mathbf{g} \rangle$, where \mathbf{x} is the current position, r is the collision radius, and \mathbf{g} is the goal position of the agent.

Our work focuses on the crowd flow of agents through corridors during building evacuation w.r.t **LoS** crowd density conditions (agents per square meter). We first define a set of agent regions, within which we can randomly distribute agents based on uniform sampling. Agent regions enforce the **LoS** crowd density conditions, described above. The goal region and/or the desired velocity of each agent is set to a region or a direction that effects desired interactions with the other group(s) of agents, and obstacles of interest.

For all experiments, agents and pillar obstacles are represented by disks with a radius of $0.23m$ corresponding to an average human shoulder width. We study crowd flow in various crowd densities conditions corresponding to known **LoS** described in [Fruin 1971]. These discrete crowd densities, used in the optimization experiments of Section 4, are 0.2, 0.4, 0.6, 0.9 and 2 agents/ m^2 for **LoS** A, B, C, D and E respectively. For evaluation of **LoS** as it relates to steering algorithms, we densely sample across a spectrum of crowd densities ranging from **LoS** A-E.

Uni-directional Hallway Egress: In this configuration, 100 agents are randomly placed in regions to accommodate crowd densities across our examined **LoS**. These regions correspond to **LoS**: A ($6m \times 83.33m$), B ($6m \times 41.66m$), C ($6m \times 27.77m$), D ($6m \times 18.51m$) and E ($6m \times 8.33m$). Each agent has a target location in the goal region outside of the hallway Fig. 2.

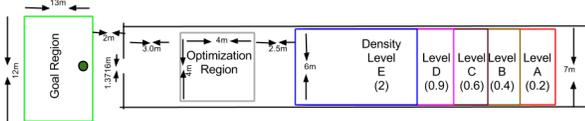


Figure 2: Uni-directional hallway scenario.

Bi-directional Hallway: This configuration, as illustrated in Fig. 3 has two groups of agents travelling in opposite directions in the hallway. There are 50 agents in each group placed in regions to accommodate crowd densities across our examined **LoS**. These regions correspond to **LoS**: A ($6m \times 41.66m$), B ($6m \times 20.83m$), C ($6m \times 13.88m$), D ($6m \times 9.25m$) and E ($6m \times 4.16m$). Each agent has a target location in the goal region outside of the hallway.

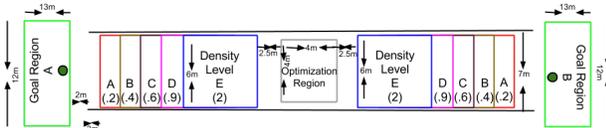


Figure 3: Bi-directional hallway scenario.

3.4 Crowd Flow

Crowd flow has been defined in a few of ways [Johansson et al. 2008; Helbing et al. 2007]. For our study, we define crowd flow as

the rate at which agents reaching their goal position:

$$f(\mathbf{p}) = \frac{|A_c|}{t_c}, t_c = t_l - t_0 \quad (1)$$

where t_0 and t_l are the completion times for the first and last agents to reach their goals respectively.

4 Optimizing Level of Service

Placing pillars at appropriate locations in an environment has been shown to often improve crowd flow during evacuations. In this section we study this important problem, using the concepts described in the previous sections. In particular, we perform two experiments. First, we identify the optimal locations for a number of pillars in our test scenarios for each **LoS** condition. Then we study how the placements of pillars optimized for each **LoS** condition, perform for all other **LoS** conditions in each of the test scenarios. Below we describe our optimization formulation, and then our experiments.

4.1 Optimization Formulation

Here we describe our optimization process for generating environment configurations with optimal parameters. We formulate a minimization of our crowd flow objective and a penalty function within the context of an environment subspace. This section describes our objective, penalty function and optimization algorithm.

Environment Subspace. By parameterizing an environment configuration, s , we define a configuration space of an environment from which we can draw arbitrary samples, \mathcal{S}_{sub} . We construct a subspace by extending the notion of an environment configuration with parameter bounds, as depicted by the grey optimization regions of Fig. 2 & 3. We indicate the parameters which the optimization can manipulate as \mathbf{p} , and the associated bounds, or constraints, $c \in \mathbf{C}$.

We optimize our objective subject to constraints \mathbf{C} , on free parameters $\mathbf{p} \in \mathcal{P}$ for a given environment subspace \mathcal{S}_{sub} , as follows:

$$\mathbf{p}^* = \arg \min_{\mathbf{p} \in \mathcal{P}} (-f_r(\mathbf{p}) + g(\mathbf{p})) \quad (2)$$

s.t. $c(\mathbf{p}) = \text{TRUE} \forall c \in \mathbf{C}$

where $f_r(\mathbf{p})$, $g(\mathbf{p})$ are defined as follows in Eqn. 3, 4, and 5.

Objective function. We define our objective as the opposite of relative crowd flow, similar to [Berseth et al. 2015]. Relative crowd flow is defined as the ratio of the number of agents that successfully reached their destination $|A_c|$ to the average agent completion time t_{avg} . Crowd flow for a specific parametrization of an environment is computed as:

$$f(\mathbf{p}) = \frac{|A_c|}{t_{avg}}, t_{avg} = \frac{\sum_{a \in A} t_a}{|A|}, \quad (3)$$

where t_a is the simulation completion time of agent a , A the set of all agents, and $|A|$ the cardinality of set A . The environment configuration is constructed from the vector of parameters \mathbf{p} . Finally, relative crowd flow is defined with respect to a reference or default parametrization \mathbf{p}_d of the subspace, defined as:

$$f_r(\mathbf{p}) = f(\mathbf{p}) - f(\mathbf{p}_d). \quad (4)$$

Penalty function: To prevent the optimization from placing a pillar that overlaps with other obstacles in the scene, we formulate an

overlap penalty function as follows:

$$g(\mathbf{p}) = \sum_{\forall (o_1, o_2) \in \mathbf{O} \times \mathbf{O}} g_{ov}(o_1, o_2), \quad (5)$$

where

$$g_{ov}(o_1, o_2) = (ov(o_1, o_2) + 1)(1 - f_r(\mathbf{p}))^2 \quad (6)$$

CMA-ES algorithm. We chose to use the Covariance Matrix Adaptation Evolutionary Strategy (CMA-ES) [Hansen and Ostermeier 1996] for our optimization experiments. CMA-ES is robust for non-convex problems such as ours, handles noise and has good convergence properties. Fig. 4 illustrates the convergence of optimal environment parameters for ORCA in the bi-directional hallway benchmark.

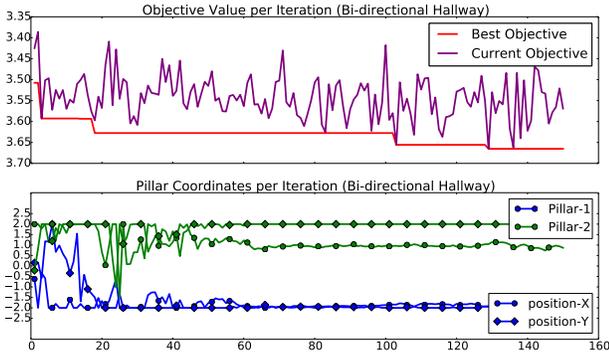


Figure 4: Parameter selection process for ORCA with 2 pillars in the bi-directional hallway environment.

4.2 Optimizing Pillar Placements

Experiment Design. We optimize the placement of 1 – 4 pillars for ORCA, SF, and PPR in the uni-directional and bi-directional hallway benchmarks. Fig. 2 and 3 illustrate the optimization regions (shown in grey) for pillar placements. Our CMA optimizer required 200 iterations with each having 6 evaluations of the objective (flow rate, $f(\mathbf{p})$) to converge. Optimizations were run on two separate machines. Linux, AMD FX(tm)-8320, having 8 Cores with 8GB memory and OSX, Intel Xeon 2.4 GHz, having 12 Cores with 12GB memory. The completion time for a single optimization for ORCA and SF was around 30 minutes, whereas for PPR, it was around 60 minutes.

Results. Optimal pillar placements produce interesting patterns across steering simulators for different LoS conditions. Fig. 5 illustrates the optimal placement of 4 pillars for ORCA, PPR, and SF in the uni-directional hallway across LoS A – E. ORCA maximizes its crowd flow by generally placing the pillars along the boundaries of the optimization region, forming wall-like structures for both the uni-directional and bi-directional hallways. PPR placements show interesting, lane-forming structures such as blocks or funnels, which become less regular at higher density conditions. The SF pillar placements show a tendency towards both wall-like and lane forming structures.

4.3 Level of Service Analysis for Optimized Environments

Experiment Design. In this experiment, we examine the influence of the optimized placement of pillars on the flow-density re-

lationship across all levels of service. For each of the three algorithms (ORCA, SF, PPR), we measure the optimal crowd flow values across all LoS (A – E) in the uni- and bi-directional hallways. Experiments are repeated using 1 – 4 optimal pillar placements, and also compared to the default benchmark, with 0 pillars. To account for sensitivity to initial configurations, each experiment was repeated 200 times and the mean flow was calculated.

Results. We report the results of our experiments for both the uni-directional and bi-directional hallway respectively.

Uni-Directional Hallway Egress. Figure 6 illustrates the flow-density relationships for ORCA, SF and PPR for 0 – 4 optimal pillar placements, in the uni-directional hallway. ORCA consistently produces the greater flow – reaching critical density values near LoS-C with $\rho_c = 0.6$ agents/m². This effectively increases the level of service for higher density crowds. SF and PPR do not exhibit a clear critical density but still we can see some increase near LoS-C and then the curve is smooth and gradually increasing. Overall, ORCA seemed to have the highest flow rate, and conformed to the typical flow-density relationship observed in real crowds. In contrast, SF and PPR produce atypical results where the flow is seen to monotonically increase across all levels of service.

Bi-Directional Hallway. In the bi-directional hallway, for the flow-density relationships for ORCA, SF and PPR for 0 – 4 optimal pillar placements, SF produces the highest flow rate for all pillar configurations, reaching its critical density near LoS-D. SF and PPR produce very interesting results. The flow-density relationship of SF is inverted, in comparison to the typical LoS curve, for all pillar configurations near LoS-C and LoS-D. It is also interesting to observe that ORCA shows an anomalous rise in flow at LoS-D, before dropping again, for 3 pillars. In a sense, PPR also reached a critical density between LoS-B and LoS-C but the curve rapidly increased afterwards. Overall, in the bi-directional scenario, SF has the highest flow rate and followed the typical flow-density curve. ORCA and PPR increased along the flow-density curve after LoS-B and did not reach critical densities in the chosen LoS range.

4.4 Flow-Density Relationships in Environments Optimized for each LoS Condition

Experiment Design. In this experiment we study what happens when we use the optimal placement of pillars obtained for one LoS condition, for crowds that exhibit other LoS conditions. As a first step, we optimized both scenarios under LoS conditions A – E, for all three steering algorithms (ORCA, SF and PPR) to get the corresponding optimal pillar positions with the highest flow rate. We then simulated each LoS scenario with each optimal pillar position to observe how a scenario with a specific LoS performed when simulated using the optimal pillar placement for another level of service. We repeat this for all environment benchmarks, configurations, algorithms, and levels of service for a total of 600 combinations with 200 simulation runs each, in this experiment.

Results. We describe our results for both environment benchmarks below.

Uni-Directional Hallway Egress. From the results of experiment in the uni-directional hallway, all algorithms, including PPR, performed similarly to their optimal LoS conditions, showing no significant difference when simulated across different crowd densities. ORCA outperforms both SF and PPR and produces consistently higher flow. The best results appear to be optimizations of 1 – 3 pillars with LoS B – C conditions. SF, on the other hand, formed the critical density, ρ_c , at LoS-D (0.9 agents/m²). Overall, SF exhibits lower crowd flow across all LoS as compared to ORCA. PPR pro-

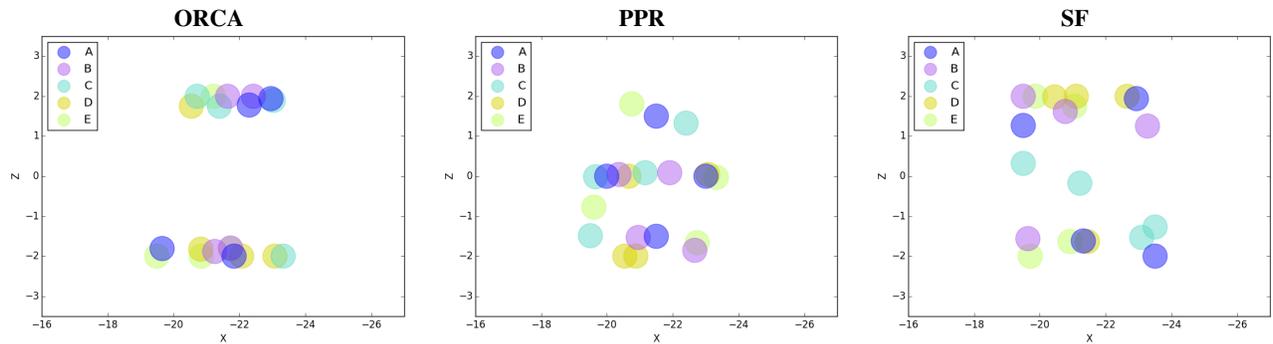


Figure 5: Optimal pillar placements across LoS conditions for ORCA, PPR and SF with 4 pillars in the uni-directional hallway.

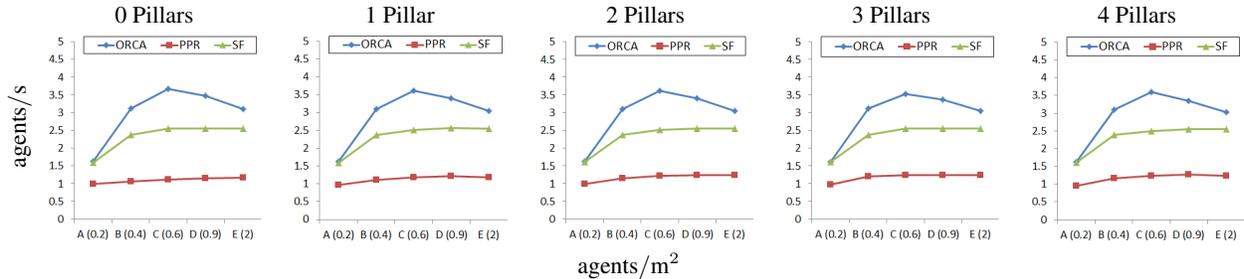


Figure 6: The optimal crowd flow values $f(p)$, across all levels of service (A–E), for ORCA, SF, and PPR, in the uni-directional hallway.

duces a monotonically increasing flow-density curve with no critical density value.

Bi-Directional Hallway. The pattern of our results change in the bi-directional scenario. SF outperforms the other two steering algorithms for this benchmark, with the highest flow values. It exhibits a typical flow-density curve with a ρ_c within the LoS C-D range ($\rho_c = 0.6$ to 0.9 agents/m²). PPR plateaus and produces consistently low flow in this particular scenario. As illustrated in Table 1, the flow-density relationship for ORCA steadily increases across all pillar configurations, as seen in the dense sampling of the default no pillar scenario. However, in the LoS E optimal scenario, the one pillar configuration exhibits an increase in flow rate across all LoS conditions.

Conclusions. We observe that optimizing for an environment under particular LoS conditions may produce interesting results under other LoS conditions. For example, optimizing ORCA for a single pillar under LoS E conditions produced a higher flow rate across LoS conditions. Similarly, for SF across all number of pillars, optimizing under higher density LoS conditions led to higher flow rates across LoS conditions. The same may be true for PPR though flow rates were not significantly increased. Though further investigation is required, it appears that optimizing under high density conditions may produce optimal results that generalize better across different conditions.

5 Conclusion

This paper presents an analysis of LoS for agent-based synthetic crowds in evacuation scenarios. We optimize the placement of pillars in these environments and observe their impact on LoS. Several interesting insights are revealed. The critical density of crowd simulators can effectively increase due to the optimal placement of 1 – 4 pillars in the environments. However, this behavior is not uniform across simulators and environment benchmarks. Different simula-

tors optimize for different crowd patterns in an effort to maximize the effective crowd flow. Force-based approaches and hybrid solutions lead to central placements in hallways and in front of the egress point, producing emergent lanes, while predictive methods produce wall-like combinations of pillars that funnel agents toward their goal. Some simulators do not reach their critical density at LoS E, and their flow-density curves do not conform to those found with real crowds. These experiments serve as an important precursor to computational crowd optimization and management and motivate the need for further study using additional real and synthetic crowd datasets across a larger representation of environment benchmarks.

References

- BERSETH, G., KAPADIA, M., AND FALOUTSOS, P. 2013. Steerplex: Estimating scenario complexity for simulated crowds. In *Proceedings of Motion on Games*, ACM, New York, NY, USA, MIG '13, 45:67–45:76.
- BERSETH, G., HAWORTH, M. B., KAPADIA, M., AND FALOUTSOS, P. 2014. Characterizing and optimizing game level difficulty. In *Proceedings of the Seventh International Conference on Motion in Games*, ACM, 153–160.
- BERSETH, G., KAPADIA, M., HAWORTH, B., AND FALOUTSOS, P. 2014. SteerFit: Automated Parameter Fitting for Steering Algorithms. In *ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, ACM, New York, NY, USA, SCA '14.
- BERSETH, G., USMAN, M., HAWORTH, B., KAPADIA, M., AND FALOUTSOS, P. 2015. Environment optimization for crowd evacuation. *Computer Animation and Virtual Worlds* 26, 3-4, 377–386.
- BEST, A., NARANG, S., CURTIS, S., AND MANOCHA, D. 2014. Densesense: Interactive crowd simulation using density-

	A	B	C	D	E		A	B	C	D	E		A	B	C	D	E		A	B	C	D	E
A	1.67	2.12	2.36	2.70	3.59		1.63	2.07	2.32	2.70	3.62		1.67	2.17	2.22	2.63	3.52		1.66	2.03	2.35	2.71	3.53
B	1.61	2.37	2.41	2.85	3.67		1.69	2.01	2.31	2.70	3.50		1.61	2.15	2.43	2.77	3.64		1.65	2.02	2.23	2.69	3.45
C	1.62	2.08	2.33	2.77	3.61		1.66	2.19	2.19	2.78	3.60		1.67	2.18	1.89	2.73	2.63		1.69	2.01	2.25	2.67	3.42
D	1.59	2.14	2.32	2.73	3.56		1.60	2.11	2.37	2.75	3.55		1.59	2.10	2.35	2.73	3.60		1.63	2.14	2.64	2.85	3.68
E	3.04	4.26	4.64	4.58	4.47		1.64	2.13	2.36	2.74	3.64		1.67	2.18	1.89	2.73	2.63		1.71	2.12	2.33	2.71	3.50
	1 Pillar						2 Pillars						3 Pillars						4 Pillars				

Table 1: The optimal crowd flow values $f(\mathbf{p})$, for each crowd density condition with all LoS optimized environments in the bi-directional scenario with ORCA. For every heatmap, each cell value is the mean crowd flow in the environment optimized for a specific LoS (rows) at a specific crowd density (columns).

- dependent filters. In *Symposium on Computer Animation*, 97–102.
- FRUIN, J. J. 1971. Pedestrian planning and design. Tech. rep.
- GUY, S. J., CHHUGANI, J., KIM, C., SATISH, N., LIN, M., MANOCHA, D., AND DUBEY, P. 2009. Clearpath: highly parallel collision avoidance for multi-agent simulation. In *Proceedings of the 2009 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, ACM, 177–187.
- HANSEN, N., AND OSTERMEIER, A. 1996. Adapting arbitrary normal mutation distributions in evolution strategies: the covariance matrix adaptation. In *IEEE International Conference on Evolutionary Computation*, 312–317.
- HELBING, D., FARKAS, I., AND VICSEK, T. 2000. Simulating dynamical features of escape panic. *Nature* 407, 6803, 487–490.
- HELBING, D., JOHANSSON, A., AND AL-ABIDEEN, H. Z. 2007. Dynamics of crowd disasters: An empirical study. *Physical review E* 75, 4, 046109.
- JIANG, L., LI, J., SHEN, C., YANG, S., AND HAN, Z. 2014. Obstacle optimization for panic flow-reducing the tangential momentum increases the escape speed. *PLoS one* 9, 12, e115463.
- JOHANSSON, A., HELBING, D., AL-ABIDEEN, H. Z., AND AL-BOSTA, S. 2008. From crowd dynamics to crowd safety: a video-based analysis. *Advances in Complex Systems* 11, 04, 497–527.
- KAPADIA, M., WANG, M., SINGH, S., REINMAN, G., AND FALOUTSOS, P. 2011. Scenario space: characterizing coverage, quality, and failure of steering algorithms. In *Proceedings of the 2011 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, ACM, 53–62.
- KARAMOUZAS, I., HEIL, P., VAN BEEK, P., AND OVERMARS, M. H. 2009. A predictive collision avoidance model for pedestrian simulation. In *Motion in Games*. Springer, 41–52.
- KARAMOUZAS, I., SKINNER, B., AND GUY, S. J. 2014. Universal power law governing pedestrian interactions. *Physical review letters* 113, 23, 238701.
- LERNER, A., CHRYSANTHOU, Y., SHAMIR, A., AND COHEN-OR, D. 2009. Data driven evaluation of crowds. In *Motion in Games*. Springer, 75–83.
- LERNER, A., CHRYSANTHOU, Y., SHAMIR, A., AND COHEN-OR, D. 2010. Context-dependent crowd evaluation. In *Computer Graphics Forum*, vol. 29, Wiley Online Library, 2197–2206.
- NARAIN, R., GOLAS, A., CURTIS, S., AND LIN, M. C. 2009. Aggregate dynamics for dense crowd simulation. In *ACM Transactions on Graphics (TOG)*, vol. 28, ACM, 122.
- NARANG, S., BEST, A., CURTIS, S., AND MANOCHA, D. 2015. Generating pedestrian trajectories consistent with the fundamental diagram based on physiological and psychological factors.
- ONDŘEJ, J., PETTRÉ, J., OLIVIER, A.-H., AND DONIKIAN, S. 2010. A synthetic-vision based steering approach for crowd simulation. In *ACM Transactions on Graphics (TOG)*, vol. 29, ACM, 123.
- REYNOLDS, C. W. 1987. Flocks, herds and schools: A distributed behavioral model. In *ACM Siggraph Computer Graphics*, vol. 21, ACM, 25–34.
- RODRIGUEZ, S., ZHANG, Y., GANS, N., AND AMATO, N. M. 2013. Optimizing aspects of pedestrian traffic in building designs. In *Intelligent Robots and Systems (IROS), 2013 IEEE/RSJ International Conference on*, IEEE, 1327–1334.
- SEYFRIED, A., BOLTES, M., KÄHLER, J., KLINGSCH, W., PORTZ, A., RUPPRECHT, T., SCHADSCHNEIDER, A., STEFFEN, B., AND WINKENS, A. 2010. Enhanced empirical data for the fundamental diagram and the flow through bottlenecks. In *Pedestrian and Evacuation Dynamics 2008*. Springer, 145–156.
- SINGH, S., KAPADIA, M., FALOUTSOS, P., AND REINMAN, G. 2009. Steerbench: a benchmark suite for evaluating steering behaviors. *Computer Animation and Virtual Worlds* 20, 5-6, 533–548.
- SINGH, S., KAPADIA, M., HEWLETT, B., REINMAN, G., AND FALOUTSOS, P. 2011. A modular framework for adaptive agent-based steering. In *Proceedings of I3D*, ACM, 141–150.
- STILL, G. K. 2007. Review of pedestrian and evacuation simulations. *International journal of critical infrastructures* 3, 3-4, 376–388.
- VAN DEN BERG, J., GUY, S. J., LIN, M., AND MANOCHA, D. 2011. Reciprocal n-body collision avoidance. In *Robotics Research*, vol. 70, Springer, 3–19.
- WOLINSKI, D., GUY, S. J., OLIVIER, A.-H., LIN, M. C., MANOCHA, D., AND PETTRÉ, J. 2014. Optimization-based pedestrian model calibration for evaluation. *Transportation Research Procedia* 2, 228–236.