

Menge: A Modular Framework for Simulating Crowd Movement

Sean Curtis¹ · Andrew Best¹ · Dinesh Manocha¹

¹ University of North Carolina at Chapel Hill E-mail: seanc,best,dm@cs.unc.edu

Received: date1 / Accepted: date2

DOI XX.XXXX/XXX / Online at collective-dynamics.eu

Abstract We present Menge, a cross-platform, extensible, modular framework for simulating pedestrian movement in a crowd. Menge’s architecture is inspired by an implicit decomposition of the problem of simulating crowds into component subproblems. These subproblems can typically be solved in many ways; different combinations of subproblem solutions yield crowd simulators with likewise varying properties. Menge creates abstractions for those subproblems and provides a plug-in architecture so that a novel simulator can be dynamically configured by connecting built-in and bespoke implementations of solutions to the various subproblems. Use of this type of framework could facilitate crowd simulation research, evaluation, and applications by reducing the cost of entering the domain, facilitating collaboration, and making comparisons between algorithms simpler. We show how the Menge framework is compatible with many prior models and algorithms used in crowd simulation and illustrate its flexibility via a varied set of scenarios and applications.

Keywords Crowd Simulation · Open Source · Software System

1. Introduction

Whether for interactive graphics, special effects, or engineering applications, crowd simulation – the simulation of a large number of independent entities acting and moving through a shared space – relies on the solution to many subproblems: determining what an agent wants to do, how it will achieve its purpose, how it responds to unforeseen challenges, and, for visual applications, determining how its virtual body moves. These subproblems are manifest in computer graphics, robotics, animation, psychology, pedestrian dynamics, and biomechanics literature, where significant work has been performed to provide increasingly superior solutions. A full crowd simulator can be regarded as the union of solutions to each of these subproblems.

1 Each of these subproblems typically admits various solutions. For example, the prob-
2 lem of determining how an agent reaches its goal can be mapped to global motion plan-
3 ning. To solve this subproblem, one could use algorithms including, but not limited to,
4 potential fields [31], road maps [37], navigation meshes [58], or corridor maps [14]. Se-
5 lecting one is a non-trivial choice. First, each of these approaches has its own strengths
6 and weaknesses – there are some problem domains for which a particular approach may
7 be better suited than others. Second, implementing one approach may be more complex
8 than another. Third, while each approach will solve the subproblem, the solutions may
9 not be the same; the choice of how a subproblem is solved can have an impact on the
10 resulting agent behavior.

11 The inherent complexity of creating a functional crowd simulator can also serve as an
12 obstacle to researchers and developers. Developing a full system is complex and time
13 consuming. Even if a researcher is interested in a single aspect of crowd simulation,
14 proper evaluation of a novel technique requires the greater context of a full simulation
15 system. Every researcher who implements an *ad hoc* crowd simulator, for the express
16 intent of testing one component, spends time and effort only tangentially related to their
17 core research. Worse yet, this effort is duplicated across independent research groups.

18 In addition, each time an entire crowd movement simulator is created to support the
19 creation of a single component, the task of performing meaningful comparisons between
20 novel and pre-existing approaches becomes increasingly difficult. Currently, the best
21 common practice is a straightforward implementation of published models for compar-
22 ison. But in these cases, a reimplementaion of a paper is unlikely to be the same as the
23 author’s original, rendering the significance of the comparison uncertain.

24 Research in and development of crowd simulation applications would benefit from a
25 common framework. This common framework would be architected with a view of the
26 various subproblems in mind; each subproblem would be encapsulated within an appro-
27 priate interface. Novel solutions to subproblems could be incorporated with other solu-
28 tions drawn from a library. A common framework would contribute to the science of
29 crowd research, not through novel models or algorithms, but by facilitating subsequent
30 research. We posit that such a framework would have multiple important benefits:

- 31 • Low-cost entry: Researchers would not be obliged to create a simulator from scratch.
32 Researchers first entering the domain could focus on one aspect, but still evaluate it
33 in a complex context by exploiting the frameworks built-in implementations.
- 34 • Focused development: Researchers could focus on a single subproblem, while ex-
35 ploiting shared implementations of solutions for the surrounding context. This
36 would reduce the initial cost of performing research in crowd simulation.
- 37 • Efficient dissemination: Novel solutions to subproblems could be released (either
38 in code or in binary form) into the common framework, allowing other users to
39 make use of the novel models, exploiting their improved properties.
- 40 • Meaningful comparison: As users release their novel subproblem solutions back to
41 the framework, other users of the framework could make direct, meaningful com-




Figure 1 An abstraction of crowd simulation based on subproblems. First, a goal is selected. Second, a base plan to reach that goal is computed. Third, the plan is adapted to local, dynamic conditions. Finally, motion is synthesized in support of the realized plan. Each subproblem can make queries into the environment to support its computation. Only those elements in grey are included in Menge, although Menge is capable of propagating complex agent state to the motion synthesis stage.

1 parisons with previous results because they are running the original implementation
 2 in its original context.

- 3 • Bespoke functionality: Custom components could be introduced according to the
 4 needs of a particular simulation problem.
- 5 • Flexible specification: In order to simulate varied, complex, real-world scenarios,
 6 the framework would be able to define simulation scenarios efficiently.

7 To that end, we present Menge, a modular, open-source, cross-platform framework for
 8 simulating crowd movement, explicitly designed to realize all of the desired beneficial
 9 properties outlined above. Moreover, we argue that Menge’s ability to provide this broad
 10 set of benefits is unique among the various simulation applications which have been re-
 11 leased by the crowd simulation community.

12 The discussion in this paper provides the evidence in support of this position – that
 13 Menge’s properties make it uniquely capable of providing the benefits outlined above.
 14 We discuss Menge’s underlying paradigm and show that many broad categories of crowd
 15 research work implicitly fit this paradigm in Section 2. In Section 3, we present the ar-
 16 chitecture designed to realized the targeted properties. We provide examples of Menge
 17 applied to meaningful research problems in Section 4, illustrating the research benefits of
 18 the architecture. Section 5 summarizes Menge’s unique capacity to serve as a common
 19 simulation framework by comparing and contrasting its benefits with those of other, pub-
 20 licly available crowd simulation systems. Finally, we offer our concluding thoughts in
 21 Section 6.

22 2. Simulating Crowds

23 Menge realizes a particular abstraction of crowd movement simulation. The abstraction
 24 is a decomposition of the problem into related *subproblems*: goal selection, plan compu-
 25 tation, plan adaptation, and spatial queries (see Figure 1). This is not a novel abstraction;

1 it has been referred to in previous work [13,61] and is well represented in the crowd sim-
2 ulation literature. In this section, we discuss representative work in crowd simulation in
3 the context of this abstraction.

4 2.1. Goal Selection

5 The first subproblem, *goal selection*, involves determining what each pedestrian wants to
6 achieve. Generally, decisions of this type can incorporate diverse factors, e.g. psychology,
7 world knowledge, etc. What the pedestrian wants to achieve can change with respect to
8 time and conditions. The complexity required depends on the simulation scenario and can
9 range from simple (flow down a corridor) to complex (populating a train station).

10 The problem of determining what an agent *wants* to do has been extensively explored.
11 Shao and Terzopoulos [54] used *situation calculus* to author a complex train platform sce-
12 nario . Uliny and Thalmann [61] computed high-level behaviors with a combination of
13 rules and behavior finite state machines. Similarly, Bandini et al. [1] used finite state ma-
14 chines to model complex behaviors with a cellular automata pedestrian model. Paris and
15 Donikian use a hierarchical finite state machine to determine high-level agent behaviors
16 (although it is used to determine sub-tasks selected to reach the pre-defined, ultimate goal)
17 [46]. Generally, this domain is solved using some form of decision or network graph. The
18 product of this stage, a “goal”, is provided to the next stage as input.

19 2.2. Plan Computation

20 The second subproblem, *plan computation*, seeks to create a *static* plan to achieve the
21 goal. This is most typically associated with *motion planning* [37]. In crowd simulation,
22 if the goal requires the agent to perform an action at its current location, the motion
23 planning consists of motion synthesis of the pedestrian’s visual representation¹. If the
24 goal requires the agent to traverse the simulation domain, then the problem combines
25 “path planning” and motion synthesis. The path is an abstract concept. An agent’s path
26 defines an agent’s *preferred velocity* – the velocity the agent would take at any given
27 moment to make progress towards its goal.

28 There are multiple approaches for computing paths. Many of them are predicated on
29 discretizing the traversable space into connected primitives. The connected primitives
30 imply a graph which can be searched using standard algorithms (e.g., A*). These graph-
31 based algorithms include: road maps [37], navigation meshes [44,58], delaunay triangu-
32 lation [36], and corridor maps [14]. These data structures have traditionally been applied
33 to traversable space with respect to static obstacles, but work has also been performed to
34 adapt them to dynamic changes to traversable space (e.g., [25,27,63]).

35 Another common approach uses potential fields. The simulation domain is discretized
36 and a field is computed that is the gradient of a cost function [31,35]. No path is explic-
37 itly computed. Instead, the resultant vector field provides a direction of “optimal” travel

¹As indicated in Figure 1, Menge does not include motion synthesis, but its simulation output is compatible with off-line synthesis.

1 toward the goal. Plan computation’s ultimate product, preferred velocity, serves as input
2 to the next subproblem.

3 **2.3. Plan Adaptation**

4 Typically, computed plans only consider static obstacles and low-frequency phenomena.
5 This gives rise to the third subproblem, *plan adaptation*. Rather than recomputing a
6 plan each time the simulation environment changes, the plan is adapted to handle local,
7 dynamic obstructions as needed. This subproblem has many names: “pedestrian model”,
8 “local navigation”, “steering”, etc. Essentially, the solution to this subproblem transforms
9 the ideal, preferred velocity into a *feasible* velocity.

10 There are a large number of models which adhere to this paradigm. Such approaches
11 include, cellular automata [53], social forces [22, 30, 47], vision-based [45], continuum-
12 based [42], velocity-obstacle-based [3, 48], and rule based [50]. *All* of these models are
13 compatible with the plan adaptation abstraction. This list is meant to be representative of
14 classes of simulation paradigms; for a more thorough discussion, please see contemporary
15 surveys [11, 65].

16 It is worth noting that there are crowd models which use a different paradigm (e.g., [29,
17 60]). In these problems, the plan computation and adaptation are collapsed into a single
18 problem; the plan computation considers the full domain, rendering adaptation largely
19 unnecessary. Even with these differences, they could still be implemented in Menge;
20 in this case, all of the work would be performed during plan computation, and the plan
21 adaptation would be an identity operation.

22 **2.4. Motion Synthesis**

23 For visual applications, it is necessary to compute physical character motion consistent
24 with the activity computed by the previous stages. There has been a great deal of work
25 in this field including procedural methods [6, 59], data-driven methods [20, 34, 38, 41],
26 and, for locomotion, foot-step driven methods [2]. In its current release, Menge does not
27 directly address this issue².

28 **2.5. Environmental Queries**

29 Finally, the various subproblems typically need to perform spatial queries in the environ-
30 ment. For example, it is reasonable to limit the effect of the environment on an agent to
31 those factors which are in the line of sight to the agent (visible) or near the agent (proxi-
32 mal). To support this type of operation, we require the ability to perform spatial queries
33 such as visibility queries or proximity queries. For details on the many solutions to these
34 types problems, we refer the reader to the following resources for visibility queries [9]
35 and proximity queries [52].

²The visualizations shown in Section 4 have been produced by a proprietary visualizer using Menge’s output data.

2.6. Crowd Systems

There is also research in full crowd simulation systems. Autonomous Pedestrians, in part inspired by Newell’s [43] Unified Theories of Cognition, expresses the crowd simulation problem as a composition of conceptual layers [54]. These conceptual layers correspond well to Menge’s abstraction of goal selection, plan computation, and plan adaptation. Other open-source simulation systems have been released, e.g., SteerSuite [56], ADAPT [28], etc. We provide a detailed comparison with these systems in Section 5.

3. Menge’s Architecture

In this section, we discuss the design philosophy and architecture of Menge. We analyze how this architecture realizes the benefits of a common simulation framework in Section 5.

3.1. Mathematical Realization

Menge’s architecture is primarily focused on facilitating the simulation of agents *moving* through a shared space.³ The problem of computing agent trajectories can be thought of as an initial value problem (IVP):

$$\dot{\mathbf{x}}_i(t) = \mathbf{v}_i(t) = \mathbf{V}_i(t, \mathbb{S}(t)), \quad (1)$$

where $\dot{\mathbf{x}}_i(t)$ or $\mathbf{v}_i(t)$ is the instantaneous velocity of agent i at time t , $\mathbb{S}(t)$ is the simulator *state*, likewise at time t , and \mathbf{V}_i is a function that determines the agent’s instantaneous velocity. By solving for $\mathbf{x}_i(t)$, we determine the position of the agent with respect to time.

The simulator state \mathbb{S} is the union of all entities in the scene, including the features of the simulation domain (e.g., obstacles) and the full crowd state space. The crowd state space $\mathbb{X} = \bigcup_i \mathcal{X}_i$ is the union of each agent’s state space. The minimum agent state space necessary to satisfy the differential equation is $\mathcal{X}_i = [\mathbf{x}_i \ \mathbf{v}_i]^T$, where \mathbf{x}_i and $\mathbf{v}_i \in \mathbb{R}^2$. Menge assumes that simulation is performed in a two-dimensional domain⁴. In practice, particular solutions to the initial value problem require additional per-agent properties which extend the agent state.

Ultimately, the properties of the crowd simulator, and the behaviors its agents exhibit, is dominated by the agent state and, more particularly, the velocity function \mathbf{V}_i .

3.2. Conceptual Abstraction as Functions

We can easily map each of the conceptual subproblems into functions. Furthermore, we can compose those functions to define the velocity function \mathbf{V} . The IVP abstraction may

³Menge’s architecture can also account for simulation in which agents remain stationary but nevertheless have changing relationships with respect to each other and their environment (see Section 3.3.)

⁴Although allowances are made for three-dimensional simulation domains that are only *locally* two-dimensional.

1 admit other mappings, but this mapping supports the modular formulation which is one
2 of Menge’s design goals.

3 The goal selection subproblem would be: $G_i : t \times \mathbb{S} \rightarrow \mathbb{R}^2$. For a single agent i , this
4 function maps time (t) and simulation state (\mathbb{S}) into a two-dimensional goal position⁵.

5 The plan computation becomes *path* computation and its corresponding function, $P_i : t \times \mathbb{S} \times \mathbb{R}^2 \rightarrow \mathbb{R}^2$, maps time, simulation state, and the agent’s goal position into an in-
6 stantaneous preferred velocity.
7

8 Finally, the plan adaptation function, $A_i : S_i \times \mathbb{R}^2 \rightarrow \mathbb{R}^2$, maps the preferred velocity
9 and *local* simulation state into a feasible velocity. Generally, the adaptations are assumed
10 to have limited temporal validity, so in this case, “local simulation state” refers to the
11 simulation features near the agent i .

12 The simulation state serves as a parameter to all three functions. By assuming that im-
13 plicitly, the functions simplify to: $G_i : t \rightarrow \mathbb{R}^2$, $P_i : \mathbb{R}^2 \rightarrow \mathbb{R}^2$, and $A_i : \mathbb{R}^2 \rightarrow \mathbb{R}^2$. The instan-
14 tantaneous velocity of an agent is the composition of these functions: $\mathbf{V}_i(t) = A_i(P_i(G_i(t)))$
15 and can be substituted into Eq. 1 as:

$$\mathbf{v}_i(t) = A_i(P_i(G_i(t))). \quad (2)$$

16 Menge implements this abstraction. Each subproblem function is implemented by a
17 set of one or more orthogonal *elements*. A particular crowd simulator can be instanti-
18 ated by specifying particular elements and their relationships. For example, configuring
19 two different simulators such that they use the same solutions to the goal selection and
20 path planning subproblems, but different path adaptation solutions is trivial; one simply
21 changes the reference to the path adaptation module in the Menge project file (see Sec-
22 tion 4.1 for specific examples). This is how one would perform comparisons between two
23 or more steering algorithms.

24 3.3. Stationary Agents

25 Any crowd simulation system which is primarily focused on *moving* agents would seem
26 to inherently consider all stationary agents to be equivalent. In reality, two stationary
27 agents could still have significantly different properties, goals, and relationships with their
28 surroundings. Menge’s architecture makes it possible to distinguish between two agents
29 which may otherwise have identical trajectories (e.g., standing still) via its Behavioral
30 Finite State Machine (BFSM). Two stationary agents could occupy different states in the
31 BFSM, representing different activities or mental conditions. The trade show example in
32 Section 4.1 illustrates just this distinction.

33 3.4. Architectural Elements

34 Menge’s modular architecture is based on the concept of *elements*. An element type
35 defines a particular aspect of a subproblem. The element type defines an interface that

⁵A goal point in \mathbb{R}^2 is a common simplification; goals could be regions. But for many applications, this simplification is sufficient.

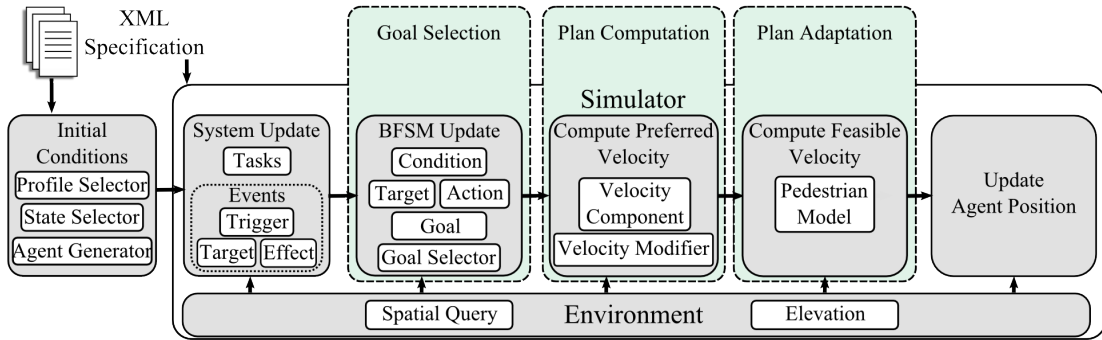


Figure 2 Menge's computation pipeline. The modular *elements* are shown in white boxes. The green boxes show how the elements relate to the conceptual subproblems. The simulator definition (including initial conditions and BFSM) is given as an XML specification. At each time step, the system updates event state and task state. Then the BFSM is updated for each agent. Next, the preferred velocity for each agent is computed. The pedestrian model is used to compute a feasible velocity. Finally, the agent position is updated.

1 can be implemented to provide a particular solution. Each element type can have an
 2 arbitrary set of implementations. The implemented elements are explicitly instantiated
 3 via the XML specification. In its initial release, Menge includes a set of representative
 4 implementations for each element type.

5 The elements are grouped by functional purpose. Conceptually, we ascribe selecting
 6 a goal and planning a path to an agent's *behavior*. We model agent behavior and how it
 7 changes with respect to time with a Behavioral Finite State Machine (BFSM). As such,
 8 the *Goal Selection* and *Plan Computation* problems are solved by elements which belong
 9 to the BFSM. The *Plan Adaptation* domain belongs to the pedestrian models element.
 10 Menge's underlying system is exposed via the system elements and, finally, initial sce-
 11 nario conditions are defined by a set of appropriate elements. We will now discuss the
 12 seventeen elements which make up Menge's modular architecture, as illustrated in Fig-
 13 ure 2. An example scenario specification can be seen in the Appendix.

14 **Agent state:** We have previously introduced the agent state as the vector $[\mathbf{x}_i \ \mathbf{v}_i]^T$. These
 15 properties are sufficient to express the initial value problem, but in practice, for a particular
 16 agent model, more parameters are required. We refer to the agent state vector consisting
 17 of position and velocity as the *agent state*, or a-state. The set of additional properties
 18 (e.g., radius of disk, response time, etc.) will be called the *behavioral state* or b-state;
 19 modifying these properties changes how the agent's trajectory is computed, leading to a
 20 different agent behavior. Finally, to avoid confusion, we will refer to a state in the BFSM
 21 as an agent's FSM-state.

22 3.5. Behavioral Finite State Machine Elements

23 The Behavioral Finite State Machine encapsulates the core of agent behaviors. Each state
 24 in the BFSM governs what goal the agent seeks, how it intends to achieve that goal, and
 25 can even influence the agent's fundamental characteristics, modeling changes in mood and
 26 thought. The transitions from one state to another govern changes in the agent's behavior.

1 Finite state machines have been shown to be quite effective for this purpose [1, 61]. The
2 specification of a particular BFSM defines how the agents interact with their environment
3 and each other and how those relationships change with time.

4 **Condition**

5 The `Condition` element, in conjunction with the `Target` element, defines a BFSM
6 transition. The `Condition` provides a boolean test which determines if a pre-defined
7 condition is satisfied. If so, the transition is activated and the agent exits its current FSM-
8 state and moves to the FSM-state defined by the transition's `Target` element. An FSM-
9 state can be connected to multiple out-going transitions. These transitions are prioritized
10 and the condition of each transition is evaluated in priority order; the first transition whose
11 condition is met is taken.

12 The `Condition`'s boolean test can consist of arbitrary logic. Menge's default imple-
13 mentation contains implementations which depend on temporal, spatial, and stochastic
14 parameters. For example, in simulating passengers disembarking an airplane, an agent
15 might wait to leave its seat until the aisle is empty; this would be realized with a cus-
16 tom `Condition`. Pre-existing conditions can be combined or new implementations can
17 be introduced to the system via its plug-in architecture to achieve desired results. See
18 Section 4.2 for examples.

19 **Target**

20 The `Target` element determines which FSM-state an agent moves to when the corre-
21 sponding `Condition` is satisfied. In a strictly-defined finite state machine, a transition
22 would connect one source FSM-state to one destination FSM-state. When defining agent
23 behavior via the BFSM, it can be convenient to model the behavior that a single con-
24 dition could lead to one of a set of new FSM-states, based on some additional criteria.
25 The `Target` element makes this possible in a compact manner. Menge includes targets
26 which allow transitions to a single FSM-state, transition to a randomly selected member
27 of a set of FSM-states, or an automatic return to the FSM-state preceding the current state.
28 Simulating a train station would provide a simple example; following a "buying-ticket"
29 FSM-state an agent might proceed to concessions or their train platform. The probabilis-
30 tic target will allow for a controlled distribution of behaviors. As with all elements, new
31 target implementations are easily introduced.

32 **Action**

33 The `Action` element allows an FSM-state to directly make changes to an agent's a-
34 state or b-state; the BFSM *acts* on the agent and not, as the name may suggest, an action
35 taken *by* the agent. `Actions` are executed on an agent when the agent enters the FSM-
36 state and can be configured to undo the change when the agent leaves the FSM-state or
37 not, as appropriate for the simulation. These actions can be used to varying effect. For
38 example, stress can be modeled by an agent successively entering an "increased stress"
39 FSM-state where each time, an `Action` modifies the agent's b-state properties to repre-
40 sent a heightened response to stress (see Section 4.2). An `Action` element can also be
41 used for reasons of convenience. For example, a simple scenario with periodic bound-
42 aries can be simulated by including an `Action` which teleports agents from their current
43 position back to the beginning of a straight hallway (demonstrated in the Appendix).

3.6. BFSM Goal Selection Elements

In simple scenarios, goal selection can be defined externally to the simulator and remain constant for the simulation duration (e.g., flow down a corridor). In complex scenarios, the agent’s goal can change from moment to moment. These changing goals are modeled using the FSM-states. Upon entering an FSM-state, an agent is assigned a `Goal` using the `Goal Selector` element associated with that FSM-state.

Goal

The `Goal` element is the basic primitive for defining the space the agent wants to reach. As previously indicated in Section 3.2, an agent’s goal is a region in two-dimensional space. Menge’s default implementation contains a number of simple, convex regions (a point, a circle, an axis-aligned box, and an oriented box). At any given moment, the agents seek to move toward the nearest point in the region. By defining `Goals` as two-dimensional regions, `Goals` can be efficiently shared by multiple agents without causing artificial queuing arising from agents waiting to access, what would otherwise be, a point goal. Regions inherently have a greater capacity to accommodate multiple agents. Menge `Goals` can have finite “capacity”, meaning that there is a limit on the number of agents which can simultaneously share that goal. Each agent pursuing that goal consumes a portion of the capacity; when all capacity is taken, no more agents can be assigned that goal.

Goal Selector

The `Goal Selector` element is the primitive which defines the basis for assigning an agent a `Goal`. When an agent enters a state, its `Goal Selector` is evaluated and a `Goal` is assigned to the agent. When the agent leaves the FSM-state, the goal is “released”. This behavior is configurable; the `Goal Selector` can be made “persistent”, meaning that the `Goal` assigned the first time the `Goal Selector` is evaluated is not freed up when the agent leaves the state. This allows the agent to return to the state and return to its original goal. It also means that the capacity of that goal is not freed up. Furthermore, this persistent goal can be shared across multiple states via “goal sharing.” Menge includes a wide range of goal selectors including: a single, pre-defined `Goal`, a uniform or weighted random selection from a set, the nearest or farthest to the agent’s current position in the set (based on Euclidian distance), the nearest or farthest based on path length through a navigation mesh, and more. Ultimately, a novel `Goal Selector` could include arbitrary algorithms for selecting a `Goal`. For example, pedestrian simulation was used in the redesign of the London Bridge Station. Surveys of passenger behaviors were used to build a statistical model for assigning destinations [23]. This statistical model could serve as stochastic weights on a set of `Goals`.

3.7. BFSM Plan Computation Elements

Solutions to the plan computation subproblem must provide an instantaneous *preferred velocity*; at any given time, an agent should “know” which direction and at what speed it wants to travel. The relationship between agent and goal can range from trivial (standing still) to complex (navigating a maze). Menge is architected in such a way as to easily

1 specify what type of implemented solution to use in any context; one simply references
2 the desired element in the Menge project file. The elements used for plan computation
3 are the `Velocity Component` and `Velocity Modifier`.

4 **Velocity Component**

5 The `Velocity Component` element is responsible for computing the agent's preferred
6 velocity; each FSM-state contains one `Velocity Component`. As such, the
7 manner in which a preferred velocity is computed for an agent in one FSM-state can be
8 completely different from that computed for the same agent in a different state. For ex-
9 ample, in simulating a train station, a pedestrian would travel to the train platform and
10 then stand and wait for the train. The FSM-state that corresponds to the traversal of the
11 train station would use a `Velocity Component` that can find a path through the com-
12 plex environment. But the waiting FSM-state can simply produce a preferred velocity
13 sufficient to maintain its position.

14 Generally the `Velocity Component` implementations primarily define the *direc-*
15 *tion* of preferred velocity and rely on the agent's own preferred *speed* to specify the mag-
16 nitude of the preferred velocity vector. However, the interface also allows for a velocity
17 component to arbitrarily deviate from the agent's preferred speed.

18 Following Curtis et al. [10], preferred velocity is represented by an *arc* of velocities
19 rather than the single vector traditionally used. The arc represents a space of velocities
20 all of which would lead the agent to travel through a space of topologically equivalent
21 paths. The arc is coupled with a function defined over the domain of the arc to distin-
22 guish a single "most-preferred" velocity from the space. This preferred velocity arc is
23 the output of the `Velocity Component` and acts as input to the plan adaptation layer.
24 For algorithms which cannot generate such a velocity arc, an arc with a zero-radian span
25 is sufficient. Similarly, if a pedestrian model cannot make use of an arc of velocities, it
26 can operate strictly on the most-preferred velocity from the arc, maintaining the broadest
27 compatibility.

28 Menge includes many default `Velocity Component` implementations including
29 graph searches on road maps or navigation meshes, straight-to-goal computation, guid-
30 ance fields, and constant velocities, allowing for the creation of complex scenarios and
31 facilitating the efficient creation of simple scenarios.

32 **Velocity Modifier**

33 The `Velocity Modifier` element serves as an interface between the plan compu-
34 tation and plan adaptation modules. The `Velocity Component` is typically an imple-
35 mentation of a global path-planning algorithm concerned with minimizing a property of
36 the path (e.g., length or travel time). The path adaptation uses purely local information to
37 transform the preferred velocity into a feasible velocity. However, this paradigm may be
38 insufficient for modeling behaviors that are dependent on temporal or spatial scopes that
39 lie outside of the global or local path planner. The `Velocity Modifier` element pro-
40 vides a mechanism for introducing additional *layers* of velocity computation. The element
41 receives a preferred velocity as input and transforms the preferred velocity based on its
42 intrinsic algorithm to output a new, modified preferred velocity. A series of `Velocity`
43 `Modifier` elements can be composed to produce the final preferred velocity used by the
44 plan adaptation stage.

1 For example, a `Velocity Modifier` element can be used to perform mid-range
2 collision avoidance (e.g., [16,19]); the basic direction of travel to reach the ultimate global
3 goal can be modified according to the presence of other agents beyond the planning hori-
4 zon of the local collision avoidance. Menge includes modifiers for modeling formations,
5 moving on uneven terrain, and modeling pedestrian density sensitivity (see Section 4.2).

6 **3.8. Plan Adaptation**

7 The preferred velocity computed in the previous section reflects a *static* plan. Dynamic
8 features, such as other agents, may interfere with the execution of that plan. Thus, the
9 preferred velocity needs to be transformed to the next best *feasible* velocity. The definition
10 of “best” and how it is evaluated can be arbitrary. As shown in Section 2.3, there already
11 exist many different models which adhere to this paradigm and, therefore, are compatible
12 with the Menge framework.

13 **Pedestrian Model**

14 At its core, a novel `Pedestrian Model` element need only define a single function:
15 the function mapping preferred velocity to feasible velocity. In practice, novel models
16 require their own parameters. As with all other elements, part of the design includes an
17 interface to automatically extend the XML simulator specification to parse and validate
18 required model parameters. Menge’s initial release includes several models including
19 two velocity-obstacle-based models and several force-based models. Additional models
20 are forthcoming, including continuum and cellular automata.

21 **3.9. System Elements**

22 The previous elements provide the core behavioral functionality of a Menge simulation.
23 In contrast, the system elements encapsulate the elements which support behavioral com-
24 putation. This includes the `Spatial Query`, `Elevation`, `Task`, and `Event`-related
25 elements.

26 **Spatial Query**

27 The `Spatial Query` element provides an interface to perform visibility and prox-
28 imity queries. Implementations of novel spatial query algorithms and data structures can
29 be incorporated in Menge via the plug-in architecture. Menge includes two different im-
30 plementations: a navigation-mesh centric query class and a kd-tree-centric class. Other
31 spatial queries can be introduced as simulation needs present themselves. For example, it
32 is easy to imagine that in some cases, a simple grid-based solution may be best.

33 **Elevation**

34 Menge performs its simulation in two dimensions. Strictly speaking, it can be consid-
35 ered to be a local, two-dimensional manifold in a larger, complex domain. Menge pro-
36 vides the `Elevation` element to provide a mapping from the local 2D planning plane
37 to a complex topology. The `Elevation` element defines the height and the gradient of
38 the domain at an agent’s position. Menge’s default release includes two `Elevation`
39 implementations: a 2.5D height field and a navigation mesh (which allows for complex,
40 non-planar topologies; see Section 4.1 for an example).

1 Task

2 The `Task` element is the mechanism by which Menge allows for the insertion of arbitrary, user-defined blocks of work into the simulation pipeline. `Tasks` are evaluated serially in the update stage (as shown in Figure 2). The `Task` can be explicitly instantiated in the simulator specification, or implicitly instantiated in support of another element. For example, algorithms which use a navigation mesh require accurate knowledge of where on the mesh an agent is located. The work to update this information is encoded in a task and executed at the beginning of the pipeline cycle. Even if multiple, independent elements require this work to be done, the shared task guarantees the work is only performed once. Alternatively, a `Task` can be explicitly instantiated by the user in the XML specification.

12 Event Triggers, Targets, and Effects

13 Menge provides the basis of a complex event system. An event is uniquely defined by three elements: `Event Trigger`, `Event Effect`, and `Event Target`. An event is *triggered* by some specified condition being met. In response its corresponding *effect* is applied to the indicated *target*. The event system has been decomposed in this way to maximize re-use of conceptual blocks. Events complement the BFSM for changing the simulation with respect to time. The BFSM changes the agents behavior based on the agents internal state (e.g., reaching a goal, running out of time, etc.) The event system allows changes to an agent due to factors external to the agent (e.g., a fire happening at a random time).

22 `Event Triggers` define the conditions for an event to be emitted. The conditions can be defined with respect to any subset of the simulator state. This can include simple timers (such as traffic signals), region population, user actions (in an interactive context), or an `Event Trigger`'s arbitrary internal state. `Event Targets` specify the Menge components upon which the event operates. Events can affect agents, states, or other elements of Menge; one could use an event to dynamically “re-wire” the BFSM. `Event Effects` encode the actual effect of the event when triggered. `Event Effects` can include changing b-State parameters of agents, disabling transitions, terminating the simulation, dynamically blocking pathways, etc.

31 Finally, Menge's architecture assumes that agents are independent entities. This admits the possibility of extensive, simple parallelization of the algorithms on shared-memory systems. The major stages in the simulation pipeline (such as computing preferred velocity, computing feasible velocity, updating agent state, etc.) are performed in parallel and the pipeline is synchronized at the end of each stage. This gives Menge the potential to be very scalable for many agents on many cores (see Section 4.1 for details).

37 3.10. Scenario Specification Elements

38 Menge also provides elements for specifying the initial conditions of the simulation as well as the BFSM. To define the initial conditions of a simulation, each agent's a-state, b-state, and FSM-state are initialized by the `Agent Generator`, `Profile Selector`, and `State Selector` elements, respectively. A group of agents is defined by a triple

1 consisting of an instance of each of those elements. The impassable obstacles in the scene
2 are defined by the `Obstacle Set` element.

3 **Agent Generator**

4 The `Agent Generator` is responsible for generating a number of agents and assign-
5 ing them initial positions and velocities (the agent's a-state). To facilitate the construction
6 of simulation scenarios, Menge provides several implementations ranging from explicit
7 lists of agent positions to abstractions of two-dimensional arrays of agents. Using para-
8 metric generators makes experimenting with the simulator simple; one can simply modify
9 the parameters to scale the number of agents in the simulation.

10 **Profile Selector**

11 An agent's b-state is defined by an *agent profile*. The agent profile consists of collec-
12 tions of values for b-state parameters. For a given property, the profile can define a value
13 as a global value or drawn from a distribution of values. For example, one could model a
14 crowd of average pedestrians by defining an agent's preferred *speed* with a normal distri-
15 bution⁶ (mean: 1.3 m/s, standard deviation 0.1 m/s). A `Profile Selector` assigns
16 a user-defined profile to each agent. The assignment criteria can be, as with all Menge
17 elements, based on arbitrary user-defined principles. They could be based on initial po-
18 sition in the simulation, count, round-robin assignment, random assignment, etc. Profiles
19 and `Profile Selector` elements permit the user to efficiently create heterogeneous
20 crowds. The populations can easily be varied to facilitate experimentation.

21 **State Selector**

22 The `State Selector` is similar to the `Profile Selector`. The `State Selector`
23 assigns an initial state in the BFSM to each agent's FSM-state. As with previous elements,
24 the assignment criteria can be arbitrary. This is particularly important because the BFSM
25 can consist of connected components; not every state may be reachable from an arbitrary
26 start state. These connected components inherently segregate the agents based on behav-
27 ior. Each connected component defines a unique *category* of agent (e.g, police, pedestrian,
28 etc.) To refer again to the train station, agents can easily be partitioned into initial states
29 which represent having a ticket or not through the use of a `State Selector`. The
30 `State Selector` facilitates the creation of behavioral categories.

31 **Obstacle Set**

32 `Obstacle Sets` specify the impassable walls in the simulation. These may be the
33 boundaries of an office building, or hazards which are activated dynamically. `Obstacle`
34 `Sets` allow for the explicit instantiation of obstacles through vertex lists, or more complex
35 obstacle generation such as capturing obstacles from a navigation mesh or from a geom-
36 etry file. Novel implementations could create obstacles from any arbitrary construct.

37 **3.11. Extensible XML-based Specification**

38 The XML specification facilitates realizing the design goal of a framework that provides a
39 *low-cost entry*; novel simulation scenarios can be created without writing any C++ code.

⁶In fact, Menge uses an approximate normal distribution. Values are limited to the range: $[\mu - 3\sigma, \mu + 3\sigma]$. This prevents unlikely but possibly catastrophic values from being generated.

1 Menge’s XML-based specification language is used to instantiate a simulator session,
2 define the simulation environment, initial conditions, behaviors, and more. However,
3 Menge simultaneously seeks to offer *bespoke functionality* by allowing users to introduce
4 novel elements into the system. These novel elements must also be accessible via the
5 XML specification.

6 To that end, Menge includes a set of utilities to facilitate the extension of the XML-
7 based specification. Each element definition includes a straightforward interface for com-
8 municating to Menge the parameters the element requires and how they should be repre-
9 sented in the XML specification. At run-time, Menge refers to this data to parse the XML
10 and provide the novel element implementation the required data (including detecting if
11 the data is incomplete or incorrectly formatted). In most cases, it is completely unneces-
12 sary for a researcher to deal with XML parsing in order to instantiate novel elements from
13 the XML specifications. Alternatively, the element abstraction also provides an alternate,
14 advanced interface which allows for arbitrarily complex XML sub-trees to be parsed by
15 the plug-in. If a complex XML specification is required, the plug-in writer can take the
16 responsibility for parsing it.

17 4. Application and Evaluation

18 In this section we examine specific examples which illustrate the Menge’s efficacy as
19 a research framework. We focus this discussion on the attached video. We begin with
20 the examples which illustrate the unique benefits of Menge. Then we examine other
21 pedestrian research. We show how various subproblems in pedestrian research can be
22 implemented in Menge. Finally, by implementing these independent works in the Menge
23 framework, we produce a scenario which effectively makes use of otherwise independent
24 research results.

25 4.1. Illustrative Examples

26 In this section, we draw attention to some of the examples in the accompanying video and
27 show how they illustrate the benefits of Menge. The actual simulated results are available
28 on the Menge website ⁷.

29 **Cross Flow:** The cross flow experiments illustrates a common experiment for pedes-
30 trian simulation; two groups of agents move through intersecting, perpendicular hallways
31 (shown in Figure 3(a)). In this example, we vary the `Pedestrian Model` implemen-
32 tation between a velocity-obstacle model [3], a simple social-force model [21], and a
33 predictive social-force model [30]; all other aspects of the simulation are fixed. The dif-
34 ferences in behavior due to the `Pedestrian Model` are clear (as illustrated by the
35 sample trajectories shown in Figure 4).

36 **Obstacle Course:** The obstacle course experiment compares global planning algo-
37 rithms. The agents shown in Figure 5 must traverse the scene from top to bottom. This

⁷<http://gamma.cs.unc.edu/Menge/>




Figure 3 Images from a subset of the various prototype scenarios included with Menge. (a) Cross flow highlighting pedestrian model comparisons. (b) A benchmark translated from SteerBench XML. (c) Airplane loading using random goal selection. (d) Agents work at desks and perform other activities in a three-story office building. (e) General Adaptation Syndrome algorithm simulation. (f) A battle scene showing 32,000 agents moving across complex terrain at interactive simulation rates. (g) The trade show scene demonstrating agents moving to and judging exhibits. (h) Agents (green) waiting for the aisle to clear using a custom transition in an airplane.




Figure 4 Trajectories plotted for three different pedestrian models in the Cross Flow scenario: a simple social force based model (SF) [21], a velocity-obstacle model (ORCA) [3], and a predictive forces model (Predictive) [30]. With all other simulation elements the same, these trajectories illustrate differences in the model behaviors.

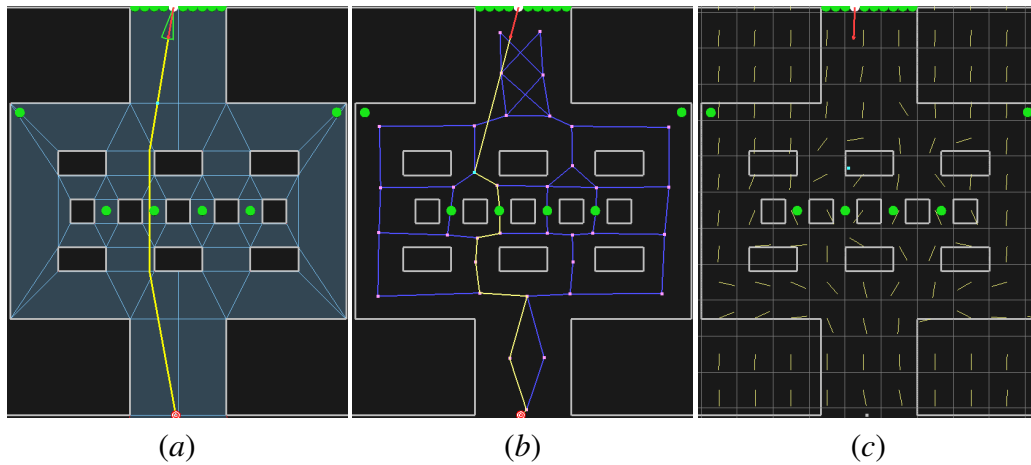


Figure 5 Visualization of three different global navigation methods applied to the Obstacle Course scenario. The green discs are agents; the yellow line represents the path computed for a single agent by each algorithm. In the case of the guidance field, each cell's direction vector is shown in yellow. (a) the navigation mesh, (b) the roadmap, and (c) the guidance field. These navigation structures can be swapped by changing a single line of XML, the `VelocityComponent`

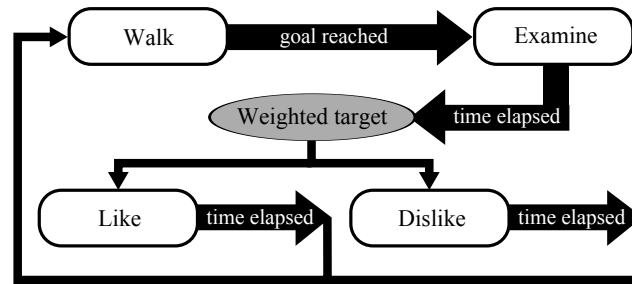


Figure 6 An illustration of the BFSM used in the trade show scenario. The white boxes represent FSM-states, the black arrows represent transition Conditions, the grey circle is a transition Target. Agents walk to an exhibit. When they reach the exhibit they enter the “Examine” state and stay there for a random amount of time after which they randomly enter the “Like” or “Dislike” state based on weighted probabilities. Finally, after a random amount of time in those states, they select and move to a new exhibit.

1 time, the Pedestrian Model is fixed and the Velocity Component changes.
 2 We compare a road map, navigation mesh, and guidance field. This experiment, in con-
 3 junction with the cross flow experiment, illustrate how Menge facilitates contrasting and
 4 comparing algorithms. Menge’s formulation of a crowd simulator as a composition of
 5 *elements* makes this possible.

6 **SteerBench:** The SteerBench scenario illustrates the ease with which scenarios can be
 7 defined in Menge’s specification language. SteerBench is a set of scenarios designed to
 8 evaluate steering algorithms [57]. Each benchmark explores a particular task of crowd
 9 navigation and offers a score for an algorithm based on several extensible criteria. The
 10 environments, behaviors, and initial conditions of SteerBench are all well expressed in
 11 Menge; we use a conversion script to translate from SteerBench XML to Menge’s XML
 12 specification.

13 **Trade Show:** The trade show demo illustrates the principle discussed in Section 3.3 –
 14 modeling changes in agent mental state without changes in movement. In this example,
 15 we are simulating the behavior of exhibition attendees on the exhibition floor. Agents
 16 approach exhibits, examine them briefly, and then decide whether they “like” the exhibit
 17 or not. The examination and decision are stationary activities but these activities are
 18 encoded as different FSM-states in the BFSM for the agent (shown in Figure 6). In turn,
 19 we can use this FSM-state information to visualize their mental state. In the video, we
 20 illustrate the examination, approval, and rejection of an exhibit via an icon floating above
 21 the agent’s head (a question mark, happy face, and angry face, respectively.) This simple
 22 visualization hints at what a more sophisticated visualizer could do with the behavior
 23 FSM-state information, synthesizing custom behavioral animation that extends beyond
 24 mere locomotion.

25 **Battle:** The battle scenario demonstrates Menge’s scalability and features the `Elevation`
 26 and `VelocityModifier` elements. Menge’s crowd simulation is not limited to simple
 27 planes. Menge agents can move along height fields and, in turn, be affected by those
 28 height fields. In this scene, an army of approximately 8,000 agents flee from a pursu-
 29 ing army of 24,000 agents. The terrain is defined by a height field. The `Elevation`




Figure 7 The results of a scalability experiment for Menge. The Battle scene was simulated using 16,000, 32,000, and 64,000 agents, respectively. The simulation used direct to goal navigation with a terrain sensitive `VelocityModifier` and ORCA for local navigation. The average frame computation time was measured based on the number of threads. The speed up over a single thread is shown. The use of a rectangle placement `Agent Generator` makes this experiment simple; we only change the agent count in the XML specification to change the initial conditions.

1 element places the agents at the appropriate elevation on the terrain. The agents use a
 2 simple `Velocity Component` pointing toward a distant goal. However, we have in-
 3 troduced a novel `VelocityModifier` which causes the agents to avoid steep inclines.
 4 Together, the agents move towards their goal while adapting to the terrain; agents flow
 5 toward valleys and avoid peaks.

6 This scenario contains the largest population and provides an opportunity to show how
 7 Menge scales with population. Figure 7 reports the performance as we varied the popula-
 8 tion. In its current state, Menge uses primitive locks to maintain safe, concurrent execu-
 9 tion. Future versions will include more sophisticated mechanisms and improve Menge’s
 10 scalability.

11 **Stadium:** In this scenario, we reproduce an experiment performed with human sub-
 12 jects: exiting a soccer stadium. This illustrates one way Menge can be used for simulating
 13 real-world scenarios. Furthermore, it highlights Menge’s ability to perform simulation in
 14 complex, three-dimensional scenarios with non-planar topology (illustrated in Figure 8).
 15 In this case, the simulation makes use of a navigation mesh structure as part of implemen-
 16 tations of a `Velocity Component`, `Elevation`, and `Spatial Query` elements.

17 **Office:** The office scenario demonstrates the most complicated BFSM in the set of ex-
 18 amples, and shows a practical alternative to simulating complex topologies. Behaviorally,
 19 each agent in the scene engages in one of several actions: working at a desk, using the
 20 restroom, getting refreshments, leaving the building, and visiting the copy room. To per-
 21 form the activity, the agent must move to the activity location. Agents can plan across
 22 floors to reach the activity location. However, instead of representing the three-story
 23 office block literally (i.e., in three dimensions using a complex navigation mesh), we im-
 24 prove the visual clarity of the simulation by laying each floor out on a single plane. This
 25 physically disconnects the stairs, but we can account for this by using a `teleport Action`




Figure 8 Images from the stadium experiment. Pedestrians walk down the aisles and exit through the stadium tunnel. (a) A photo from the original data collected by [7]. (b) A rendered screenshot of the experiment replicated in Menge. (c) The 3D navigation mesh used for the Stadium scenario. The navigation mesh provides elevation information, navigation, and provides a spatial query structure.

1 to seamlessly move agents traversing the stairs across the discontinuity. We use a road
 2 map in the scene and explicitly connect nodes across the disconnected regions. Concep-
 3 tually, the agents behave the same as if the three floors were stacked on top of each other.
 4 This scenario illustrates a combination of goal-choice mechanisms, actions, transitions,
 5 and states. It demonstrates Menge’s ability to represent populations of agents performing
 6 different tasks, with different goals and different strategies all in a single simulation.

7 4.2. Novel Models in Menge

8 The previous section illustrates Menge’s flexibility in general; abstract scenarios exercise
 9 straight-forward algorithms. But Menge can serve as an effective platform for future
 10 research as well. To illustrate this, we discuss several bodies of work – some pre-date
 11 Menge and we have implemented them in the Menge framework and others have in fact
 12 been developed on top of the Menge framework. These examples underscore how flexible
 13 the Menge framework is. Finally, we show that by implementing otherwise disparate
 14 research in a common framework, we can easily combine them to model never-before
 15 seen scenarios.

16 **General Adaptation Syndrome:** The work on modeling General Adaptation Syn-
 17 drome (GAS) by Kim et al. models how humans respond to stress [32]. Essentially, as
 18 stress accumulates, people respond by exhibiting more aggression-like behaviors. The au-
 19 thors modeled the accumulation of stress and used work by Guy et al. to model personality
 20 changes [18]. Guy et al. performed user studies to correlate agent b-state parameter space
 21 with perceptions of personality characteristics. This study was able to suggest a *displace-*
 22 *ment* vector in b-state parameter space which was the direction of increased aggression.
 23 We implemented this in Menge with a custom `Action` which applies the so-called ag-
 24 gression displacement on agents. We assign the `Action` to a stress-inducing FSM-state
 25 and include a transition which causes the agent to periodically re-enter the state – shorter
 26 periods model a higher rate of stress accumulation, longer periods, a slower accumulation
 27 rate. We reproduced one of Kim et al.’s simulation experiments: two groups of agents




Figure 9 Images from the formation experiments in the video. (a) and (b) Two formations out of a sequence created by a group of agents moving through space. (c) A dense formation of agents navigating around obstacles.

1 moving in anti-parallel directions in a wide corridor (see Figure 3e. As with the original
 2 results, as stress increases, the agent performance (as measured by flow in the corridor)
 3 initially improves before eventually breaking down under an excess of stress.

4 **Formations:** Although Menge’s agents are fundamentally modeled independently, this
 5 does not preclude complex, coordinated group behaviors such as those shown in [17,
 6 26, 64]. As a representative sample, we implemented the approach of Gu and Deng to
 7 illustrate how easily formations can be introduced into Menge [17]. This approach defines
 8 a formation via a canonical collection of *prioritized* points – transform-invariant positions
 9 which define the formation. At each time step, the canonical points are mapped to world
 10 space and agents are assigned to formation points, in a prioritized manner. See the original
 11 authors’ work for the exact details [17].

12 We reproduce this in Menge by introducing two new elements: a `Task` and a `VelocityModifier`.
 13 The `Task` is responsible for transforming the canonical formation and mapping agents to
 14 formation positions. It executes once per time step, populating a data structure used by
 15 the `VelocityModifier`. Agents in a common formation are affected by a common
 16 `VelocityModifier`. After each agent computes its own preferred velocity (presum-
 17 ably to the same goal) the `VelocityModifier` modifies it so that it will cause the
 18 agent to converge towards its position in the formation. Figure 9 illustrates some of the
 19 results using these new elements. In the video, we show one example in which a single
 20 group of agents changes formation as it traverses through space and a second example in
 21 which a larger formation navigates around obstacles.

22 **Ped-Air:** Ped-Air, a simulator described by Best et al. [4], uses Menge to simulate pas-
 23 senger loading, unloading, and evacuation behaviors in aircraft. Simulating passengers on
 24 aircraft is challenging for several reasons: passengers can span a broad space of physical
 25 and psychological types, they often are pursuing simultaneously contradictory objectives,
 26 and they must act in an extremely constrained environment.

27 Ped-Air exploits Menge’s `GoalSelector` element to model passenger seat assign-
 28 ment and to experiment with boarding strategies. The `GoalSelector` defines which
 29 seat an agent is heading towards (i.e., its seat assignment). By simply changing the param-
 30 eters of the `GoalSelector` element, Ped-Air can simulate back-to-front, front-to-back,
 31 random, and zone-based seating assignments.

32 A `GoalSelector` element is also used to model agents stowing luggage in bins. The

1 bin space is discretized into slots with fixed capacity. As each agent boards the plane,
2 it searches for a bin `Goal` near its seat with sufficient capacity for its luggage. The
3 `GoalSelector` easily determines a viable target bin, while constantly accounting for
4 capacity.

5 When disembarking an airplane, some passengers may remain in their seats until the
6 plane is mostly empty. Ped-Air models this behavior with a custom transition `Condition`.
7 A delaying passenger only transitions from its seated FSM-state to an exiting FSM-state
8 when the aisle forward of its seat is empty of passengers. Furthermore, in some cases,
9 such a passenger requires assistance to disembark. Ped-Air uses custom `Condition` and
10 `Goal` elements to achieve this. When the aisle to a waiting passenger is clear, an agent
11 representing a member of the flight staff moves to the waiting agent. The `Condition`
12 for the waiting agent to begin exiting is that the flight staff agent *reach* it. Then, when the
13 waiting agent begins the exit, the flight staff agent uses a custom `Goal` to accompany the
14 agent; in effect, the exiting agent defines a moving goal for the accompanying agent.

15 **Density-dependent Behaviors:** The Fundamental Diagram is a name given to a com-
16 monly observed phenomenon in crowd behaviors; as crowds get denser, they get slower
17 [62]. Best et al. propose an algorithm (DenseSense), based on Menge, which successfully
18 reproduces this behavior [5]. The approach works by modifying an agent’s preferred ve-
19 locity based on local density; it operates on the hypothesis that in dense environments,
20 pedestrians are less comfortable moving at high speed. The authors use a relationship
21 between various biomechanical and psychological factors and *preferred velocity* to model
22 this.

23 Like in the formation work, DenseSense uses a `VelocityModifier` to achieve its
24 goal. The `VelocityModifier` computes the density in an agent’s region and then uses
25 it to compute a “comfortable” velocity for the agent to take (see the paper for details). To
26 further optimize this task, it also introduces a new `Task`. At each time step, the custom
27 `Task` computes a density field in the simulation domain. The density field is shared for
28 all agents and the `VelocityModifier` can simply “look up” the density for the agent
29 in question.

30 **Formation Stress:** The Formation Stress example underscores Menge’s greatest ben-
31 efit: the simple combination of orthogonal research. In this example, we have combined
32 three separate research results into a single scenario: GAS, formations, and density-
33 dependent behaviors. Because they have been implemented in a common framework,
34 we can author a scenario that makes use of all three. In this scenario, a formation of
35 agents moves towards the entrance of a building. After a predetermined time, an alarm
36 sounds causing the agents to begin accumulating stress. The stress causes the agents to
37 leave their formation and run to the entrance in a chaotic manner, creating a bottleneck
38 at the entrance. After traveling through a short corridor, they enter a large hallway where
39 they must cross through a confused flow of agents, all the while exhibiting the hallmark
40 sensitivity to density seen in real pedestrians (see Figure 10).




Figure 10 Visualization of the Formation Stress scenario. This brief experiment illustrates the simplicity of combining previously unrelated algorithms to produce a novel simulation. This experiment combines results from stress modeling, formations, and Fundamental Diagram adherence [5, 17, 32]. (a) The agents travel in a three row formation towards the building. (b) After the alarm sounds, the agents run for the entrance, breaking formation. (c) When agents reach the corridor after the entryway, they must navigate through a cross-flow, respecting local density constraints on their velocity.

5. Menge’s Unique Realization of a Simulation Framework

Menge’s architecture realizes the desired framework benefits in the following ways:

- Low-cost entry: Menge is an “out-of-the-box” simulator; users can simulate novel scenarios without writing a line of code. New scenarios, from simple to complex, can be created using only the XML specification language. Menge’s behavior FSM allows for complex scenarios, illustrated by the many examples included in the release. The decomposition into elements and the many examples, serve as tutorials to new researchers.
- Focused development: Researchers can make use of Menge’s element-based design to focus their efforts on specific aspects of crowd simulation, relying on the remaining implementations to provide a robust context to test their models.
- Efficient dissemination: Novel crowd models can be released as Menge plug-ins, such as the formations and density-sensitivity functionality discussed above. Other Menge users can include the plug-ins in their own builds and immediately make use advances in the state of the art.⁸ Authors can host their plug-ins themselves, or new functionality can be rolled back into the Menge release, facilitating sharing.
- Meaningful comparison: Menge’s modular approach particularly facilitates comparisons. For a given simulation scenario, two competing algorithms can be compared, simply by changing the reference in the XML configuration. As shown in Section 4.1 with respect to pedestrian models and global navigation. Alternatively, in many cases, both algorithms could be used in a single simulation scenario.

⁸Although binary distributions are strictly possible, due to compatibility issues in C++ runtime libraries, release of the source code is preferred.

- 1 • Bespoke functionality: The elements in Menge have interfaces designed with the
2 intent to be as simple as possible – the simpler the interface, the simpler the in-
3 tegration. Furthermore, Menge includes utilities to ease the task of extending the
4 XML specification to include new element implementations.
- 5 • Flexible specification: Without compiling Menge, a user can make significant changes
6 in what is simulated, how it is simulated, and how the results are stored. Menge’s
7 XML simulation scenario specification includes more than just the utilities for test-
8 ing simulation models. It includes mechanisms for efficiently defining initial con-
9 ditions, specifying simulation parameters such as time step and duration, caching
10 simulation results, and controlling the visualization.

11 We have shown how Menge can serve as an effective framework for the crowd simu-
12 lation community. However, Menge is not the only crowd simulation software available.
13 Even putting aside the many commercial simulation packages, there are still many open-
14 source software package targeted towards academia. We have examined a number of such
15 tools and illustrate that while they have many strengths – some quite unique and valuable
16 – none of them can serve as a framework to the same degree as Menge. Although we
17 have made a concerted effort to provide the most complete survey possible, we limit this
18 discussion to those applications targeted towards the research community, are currently
19 available, and are “extensible”, in some sense.

20 All of these applications provide a *reduced entry cost* to varying degrees; simply having
21 access to pre-existing software provides a new researcher an advantage. The utility of this
22 starting point, however, depends on what the researcher intends to do with the application,
23 and how compatible the application is with that intent. This ability to introduce new
24 behavior into the simulation application is what we mean by “extensible”. Two ways for
25 an application to be extensible is to be open-source or to allow plug-ins (or both).

26 Open-source applications are extensible because users can modify the code to suit their
27 purposes. However, the architecture of the application will necessarily render some mod-
28 ifications easier than others. If the original authors intend a particular feature to be modi-
29 fied, then the code will facilitate this action. However, the converse is likely true – aspects
30 of the system the original authors expected to go untouched are likely to be tightly coupled
31 and replacing them will be more challenging. This type of extensibility, while empower-
32 ing isolated research groups, is less effective for dissemination of results. Two research
33 groups may go about modifying the original system in two different ways to incorporate
34 their novel models. Merging such results into a common framework may prove to be
35 problematic.

36 A plug-in-based application has strongly formalized extensibility. It has the same pit-
37 falls as “open-source extensibility”; only those aspects the original authors anticipated as
38 needing to be modified would be included in the plug-in interface. However, the plug-in
39 architecture greatly facilitates result dissemination. The plug-in interface defines an iso-
40 lating layer – on one side lies the application, on the other side lies the novel contributions.
41 The plug-in can be released either as code or as binaries. This is one reason why Menge
42 has a plug-in architecture and dozens of components are exposed in this interface.

1 We have found five open-source crowd simulation applications targeted towards the
2 research community. Below we give a brief overview of each application and our best
3 judgment in how well it provides the desired benefits of a common simulation framework.

4 **PedSim** is the simplest of the available applications [15] discussed here. At its core,
5 PedSim is a small C++ library consisting of a social-force-based steering model and def-
6 initions of paths and obstacles. It has no global planning ability, no high-level behaviors,
7 no events, and no goals. It has a simple application which can parse a lightweight speci-
8 fication file to initialize the simulation; agents are defined in place along with their paths.
9 Its social force model includes more features than a “basic” social force model, such as
10 a “look ahead” and a “follow” force. PedSim does not facilitate *focused development*
11 because it is so simple; one cannot rely on components that do not exist. That same sim-
12 plicity limits *efficient dissemination* and *meaningful comparisons* because so much would
13 have to be invented to mature it into a fully-fledged crowd simulator that it is unlikely that
14 any two groups would do this work in a compatible way. However, it is worth noting that
15 the limited functionality is fully accessible through its scenario configuration.

16 **OpenSteer** is a C++ application for exploring steering behaviors [51]. OpenSteer uses
17 a plug-in architecture to introduce various scenarios such as, capture the flag, multiple
18 pursuit, boids [50], waypoint following, soccer, and, curiously, pedestrians. The ap-
19 plication is not specifically designed to simulate human pedestrians and is reflected in
20 the architecture – a pedestrian derives from a “simple vehicle” class. Like PedSim, the
21 framework omits a number of capabilities: no behaviors, global planning, events, etc.
22 It is completely focused on evaluating steering behaviors for “vehicles”. The compile-
23 time plug-ins encode a particular steering behavior and two behaviors cannot co-exist in
24 a single simulation. As with PedSim, its simplicity precludes its ability to serve as a
25 framework for general pedestrian simulation rendering it unusable from a *focused devel-*
26 *opment*, *efficient dissemination*, or *meaningful comparison* perspective. Furthermore, it
27 has no external specification mechanism; simulation scenarios must be hard-coded into
28 the plug-in.

29 **ADAPT**'s primary focus is on the final stage of simulation: pedestrian visualization and
30 motion synthesis [28]. It uses Rekast, an open-source navigation mesh and steering algo-
31 rithm as the underlying planner and pedestrian model [40]. ADAPT's unique contribution
32 is in its behavior tree for controlling articulated pedestrian visualizations through, what
33 the authors term, *coordinators* and *choreographers*. The final, articulated pedestrians
34 can exhibit fine-detailed animations such as sitting, reaching, upper-body gestures, etc.
35 ADAPT's framework relies heavily on the Unity game engine to handle the visualization
36 and motion synthesis as controlled by their behavior tree. As such, its core functionality is
37 authored in C#. *Focused development* is problematic; ADAPT's focus is on the visualized
38 motion synthesis. There is excellent infrastructure for introducing new choreographers,
39 but no allowance for other aspects of crowd simulation, e.g., alternative steering or plan-
40 ning algorithms. As long as researchers are focused on motion synthesis, ADAPT can be
41 effective in *disseminating* novel work. New classes which fit into the behavior framework
42 can be distributed directly and other researchers can compile it into their own version of
43 ADAPT. However, modifying other aspects of the simulator will require custom modifi-
44 cations which may not be compatible from one research group to another. ADAPT is not

1 well suited to model *comparisons*. Scenarios are defined in code, so to compare two dif-
2 ferent models, two different compilation paths must be maintained to build two different
3 binaries. Generally, when an application has been designed with tightly integrated com-
4 ponents, replacing those components becomes logistically awkward. And any endeavor
5 trying to compare and contrast those components will exhibit unwieldiness. Bespoke
6 functionality *can* be introduced into the existing framework but, again, except for motion
7 synthesis, there are no designed vectors for introducing other novel functionality. Finally,
8 ADAPT has no apparent external simulation configuration; particular simulation scenarios
9 must be defined in code to be included in compile time.

10 **JuPedSim** [33] is the successor to the former OpenPedSim [12] and shares many com-
11 mon features. JuPedSim is a C++ pedestrian simulator apparently targeted towards evacu-
12 ation scenarios. JuPedSim has no high-level behavior module; simulated scenarios consist
13 of the agents starting from an initial condition and moving toward a final goal. Navigation
14 is handled through 3D navigation mesh-like algorithm and the steering is handled by the
15 generalized centrifugal-force model [8]. Consistent with the apparent problem domain of
16 evacuation scenarios, the simulation supports dynamic environments (e.g., doors block-
17 ing and unblocking) and the concept of pedestrian *knowledge* and how it moves through
18 the crowd (useful for simulating evacuation in smoke-filled environments). JuPedSim
19 includes tools for analyzing the results of the simulation (e.g., density analysis and com-
20 puting the so-called “fundamental diagram”.) The various components of the simulator
21 are tightly coupled which limits the ability to perform *focused development*. For exam-
22 ple, the knowledge model is implemented directly into the pedestrian model. If a new
23 researcher wanted to investigate a different pedestrian model with the same knowledge,
24 or a different model of how knowledge is transmitted, it would require extensive coding to
25 pair new components in place of the current pairing. When it comes to *efficient dissemina-*
26 *tion, meaningful comparisons, and bespoke functionality*, JuPedSim exhibits many of the
27 same issues as ADAPT. Without a plug-in architecture, there is no designed interface for
28 introducing novel models, so new models would require a custom harness. Sharing novel
29 models predicated on different harnesses is only slightly better than sharing code across
30 independent frameworks. However, JuPedSim has an extensive XML-based specification
31 language for fully exercising its functionality, giving it *flexible specification*.

32 **SteerSuite** [56] and Menge have the most similar architectural designs. The C++ appli-
33 cation includes the ability to use run-time plugins to modify the behavior of the simulator.
34 These plug-ins can be used to change the pedestrian model, global navigation algorithm,
35 spatial query mechanism, and change how the simulation is visualized. SteerSuite in-
36 cludes an XML-based specification for designing new simulation scenario and includes
37 a suite of simple scenarios for evaluation of novel scenarios. Like ADAPT, SteerSuite
38 has included the Rekast [40] code as one of the global planning algorithms. The frame-
39 work also includes instrumentation for performance profiling. However, SteerSuite has
40 no high-level behaviors. As with the previous applications, simulation scenarios consist
41 of agents in initial positions moving toward a fixed goal. The global navigation algorithm
42 is global – all agents in the simulation must use the same mechanism to move toward their
43 goal. Finally, the XML specification is not extensible; referencing novel components in
44 the specification would require modifications to the core application. SteerSuite provides

1 a strong base for *focused development*; specific global planning or steering algorithms
2 can be inserted into the system, relying on the other, pre-existing components. The same
3 plug-in architecture enables *efficient dissemination* and *meaningful comparisons* because
4 novel models can be distributed and built independent of the SteerSuite source code. And
5 a single simulation scenario can be run multiple times with different pedestrian models
6 by specifying the plug-in to apply. SteerSuite allows for *bespoke functionality* in steering
7 and global navigation algorithms, but research into high-level behaviors would require
8 completely new code to interface with the current architecture. Finally, the XML scenario
9 specification gives limited access to the simulation constructs, indicating limited *specifica-*
10 *tion flexibility*.

11 All of these applications are available and will serve their specific purpose. A re-
12 searcher, looking to enter the domain of pedestrian simulation, could select one and
13 benefit from using it as a starting point. They each have a unique strength borne of an
14 apparent targeted, intent. JuPedSim has modeled spatial awareness and knowledge prop-
15 agation. ADAPT has high fidelity crowd visualization and state-based motion synthesis.
16 SteerSuite has a suite of scenarios and comes ready with multiple pedestrian models im-
17 plemented. OpenSteer explores a number of behavioral scenarios (e.g., capture the flag,
18 etc.) The fact that these features are spread out across multiple applications is precisely
19 the functional fracturing that occurs in the absence of a common framework. This is re-
20 grettable because these features are not mutually exclusive; they could happily co-exist in
21 a single framework.

22 Menge's initial implementation includes a subset of these features. However, the under-
23 lying architecture and design of Menge makes it possible to incorporate *all* of these fea-
24 tures in future releases. Furthermore, Menge's extensibility is different from the other ap-
25 plications which have implemented plug-in architecture. Where other systems use plug-in
26 framework to introduce *replacements* for the built-in functionality, Menge's architecture
27 uses the plug-ins to *extend* its functionality; multiple independent element implementa-
28 tions co-exist in the system and can be used in a common simulation scenario (as shown
29 in Section 4.1).

30 6. Final Remarks

31 6.1. Conclusion

32 We have presented the design of a novel, modular framework for the simulation of crowd
33 movement. Through the combination of various modular constructs, called *elements*,
34 novel crowd simulators can be dynamically constructed to simulate a wide range of sce-
35 narios and behaviors. Furthermore, because of its plug-in architecture, particular imple-
36 mentations of Menge elements can be released as code or binary objects, enabling users
37 of the framework to share their own advances and benefit from the contributions of others.
38 We have discussed the validity of Menge's paradigm in the context of representative sam-
39 ples from crowd simulation literature and shown, through specific examples, the strengths
40 and properties of this framework.

1 Menge provides a platform for crowd research that facilitates straightforward combi-
2 nations of algorithmic techniques that were previously infeasible. Rather than producing
3 algorithms which target a particular subproblem in crowd simulation without considera-
4 tion of how those algorithms fit into a larger context, Menge encourages researchers to
5 produce algorithms which are inter-operable and provides algorithmic implementations
6 which are themselves inter-operable. Researchers have the opportunity to build on com-
7 mon work in a way not previously available to them. Simulators can be constructed in
8 Menge that take advantage of a number of models which would not otherwise be compat-
9 ible without such a common core framework upon which to build.

10 Menge is open-source, cross-platform, and publicly available ⁹. Ultimately, we hope
11 that the adoption of a framework such as Menge, would foster tighter integration among
12 the crowd simulation community. New researchers would enter the domain able to exploit
13 the current state of the art and directly apply their efforts to novel algorithms. Published
14 work could be closely supported by the releases of supporting code or binaries for the
15 community's benefit and future comparisons.

16 6.2. Future Work

17 Menge is a work in progress and has definite limitations. First, it currently only allows one
18 mechanism for generating high level behaviors – the BFSM. Behavior trees are a common
19 structure in game AI [49]. Both approaches essentially encode agent behavior in graph
20 nodes, but they largely differ in how the network of nodes is traversed. Currently, this
21 traversal is not an exposed part of the Menge interface, rendering behavior trees unusable.

22 Second, planning and personality are tightly coupled in the BFSM. A single FSM-
23 state specifies both *what* the agent seeks to accomplish and *how* (i.e., its personality and
24 mood). While this does not actually limit Menge's ability to model complex scenarios,
25 it can make the task more difficult, requiring redundancies in the specification where two
26 FSM-states share the same objective but possess different behavioral profiles.

27 Menge has implicitly excluded the subproblem of motion synthesis, but Menge's ar-
28 chitecture does not prevent a `Pedestrian Model` implementation from considering
29 biomechanical factors in adapting preferred velocity. Menge would certainly benefit from
30 the inclusion of a system for synthesizing motion in a modular manner similar to the other
31 elements.

32 Additionally, Menge is an agent-based crowd simulation framework. Some recent
33 work, including [55] and [24], uses motion-patches to create populated scenes of pedestri-
34 ans. These methods create agents as needed to fill motion scripts and do not contain agents
35 exploring shared spaces and planning/interacting as they accomplish disparate goals. Al-
36 though a `Pedestrian Model` and `Velocity Component` could be implemented
37 that compute paths for agents with respect to a predefined set of motion-patches, this
38 would be a substantial undertaking.

39 Menge's implementation is in its infancy. As such, there are some short-term imple-
40 mentation issues which limit its utility. As previously noted, it uses a primitive paral-

⁹<http://gamma.cs.unc.com/Menge/>

1 elism mechanism which causes its scalability to suffer. In addition, Menge simulations
2 use a fixed population; there is no mechanism in place for removing or introducing agents
3 during the course of the simulation. Menge's core element implementations have been
4 written with this eventual functionality in mind, so that it can be introduced in the future
5 without losing backwards compatibility with prior implementations.

6 In the future, Menge will seek to address these limitations and others as the community
7 explores spaces as yet unconsidered. We invite others to explore the Menge framework
8 and produce novel implementations of the many elements. We hope that Menge's future
9 growth will be fueled by groups around the world expanding its feature set according to
10 their varied needs. We invite those eager to contribute; contact information can be found
11 on Menge's website.¹⁰

12 References

- 13 [1] Bandini, S., Federici, M., Manzoni, S., Vizzari, G.: Towards a methodology for
14 situated cellular agent based crowd simulations. *Engineering societies in the agents
15 world VI* pp. 203–220 (2006)
- 16 [2] van Basten, B.J.H., Stuvell, S.A., Egges, A.: A hybrid interpolation scheme for
17 footprint-driven walking synthesis. *Graphics Interface* pp. 9–16 (2011)
- 18 [3] van den Berg, J., Guy, S.J., Lin, M., Manocha, D.: Reciprocal n-body collision
19 avoidance. In: *Inter. Symp. on Robotics Research*, pp. 3–19 (2011)
- 20 [4] Best, A., Curtis, S., Kasik, D., Senesac, C., Sikora, T., Manocha, D.: Ped-air: a
21 simulator for loading, unloading, and evacuating aircraft. In: *7th International Con-
22 ference on Pedestrian and Evacuation Dynamics*. Delft, Netherlands (2014)
- 23 [5] Best, A., Narang, S., Curtis, S., Manocha, D.: Densesense: Interactive crowd sim-
24 ulation using density-dependent filters. In: *Symposium on Computer Animation
25 (2014)*
- 26 [6] Bruderlin, A., Calvert, T.W.: Goal-directed, dynamic animation of human walking.
27 In: *Proc. of SIGGRAPH '89*, pp. 233–242 (1989)
- 28 [7] Burghardt, S., Klingsch, W., Seyfried, A.: Analysis of flow-influencing factors in
29 mouths of grandstands. In: *Pedestrian and Evacuation Dynamics* (2012)
- 30 [8] Chraibi, M., Seyfried, A., Schadschneider, A.: Generalized centrifugal-force model
31 for pedestrian dynamics. *Phys. Rev. E* **82**(4), 046,111 (2010)
- 32 [9] Cohen-Or, D., Chrysanthou, Y., Silva, C., Durand, F.: A survey of visibility
33 for walkthrough applications. *IEEE Transactions on Visualization and Computer
34 Graphics* **9**(3), 412–431 (2003). doi:10.1109/TVCG.2003.1207447

¹⁰<http://gamma.cs.unc.edu/Menge>

- 1 [10] Curtis, S., Snape, J., Manocha, D.: Way portals: Efficient multi-agent navigation
2 with line-segment goals. *Proc. of Symposium on Interactive 3D Graphics and Games*
3 pp. 15–22 (2012)
- 4 [11] Duives, D.C., Daamen, W., Hoogendoorn, S.P.: State-of-the-art crowd motion simu-
5 lation models. *Transportation Research Part C: Emerging Technologies* **37**(0), 193–
6 209 (2013). doi:10.1016/j.trc.2013.02.005
- 7 [12] Forschungszentrum, J.: Open ped sim. [http://sourceforge.net/
8 projects/openpedsim/](http://sourceforge.net/projects/openpedsim/) (2014)
- 9 [13] Funge, J., Tu, X., Terzopoulos, D.: Cognitive modeling: knowledge, reasoning and
10 planning for intelligent characters. In: *Proc. of SIGGRAPH*, pp. 29–38 (1999)
- 11 [14] Geraerts, R., Kamphuis, A., Karamouzas, I., Overmars, M.: Using the corridor map
12 method for path planning for a large number of characters. In: *Motion in Games*,
13 pp. 11–22. Springer, Heidelberg (2008)
- 14 [15] Gloor, C.: Pedsim: A microscopic pedestrian crowd simulation system. [http:
15 //pedsim.silmaril.org/](http://pedsim.silmaril.org/) (2003)
- 16 [16] Golas, A., Narain, R., Lin, M.: Hybrid long-range collision avoidance for crowd
17 simulation. *Proc. of Symposium on Interactive 3D Graphics and Games (I3D)* pp.
18 29–36 (2013)
- 19 [17] Gu, Q., Deng, Z.: Generating freestyle group formations in agent-based crowd sim-
20 ulations. *IEEE Computer Graphics and Applications* **33**(1), 20–31 (2013). URL
21 <http://doi.ieeecomputersociety.org/10.1109/MCG.2011.87>
- 22 [18] Guy, S.J., Kim, S., Lin, M.C., Manocha, D.: Simulating heterogeneous crowd be-
23 haviors using personality trait theory. In: *SCA '11*, pp. 43–52 (2011)
- 24 [19] He, L., van den Berg, J.: Meso-scale planning for multi-agent navigation. In: *Proc.*
25 *IEEE Int. Conf. on Robotics and Automation - ICRA* (2013)
- 26 [20] Heck, R., Gleicher, M.: Parametric motion graphs. In: *Proceedings of Symposium*
27 *on Interactive 3D Graphics and Games*, pp. 129–136 (2007)
- 28 [21] Helbing, D., Farkas, I., Vicsek, T.: Simulating dynamical features of escape panic.
29 *Nature* **407**, 487–490 (2000)
- 30 [22] Helbing, D., Molnár, P.: Social force model for pedestrian dynamics. *Physical Re-*
31 *view E* **51**(5), 4282–4286 (1995)
- 32 [23] Hutton, A.: London bridge station, the role of ped modelling: Pedestrian modelling
33 and design development. In: *6th International Conference on Pedestrian and Evac-*
34 *uation Dynamics*. Zurich, Switzerland (2012)

- 1 [24] Hyun, K., Kim, M., Hwang, Y., Lee, J.: Tiling motion patches. *IEEE*
2 *Trans. Vis. Comput. Graph* **19**(11), 1923–1934 (2013). URL [http://doi.](http://doi.ieeecomputersociety.org/10.1109/TVCG.2013.80)
3 [ieeecomputersociety.org/10.1109/TVCG.2013.80](http://doi.ieeecomputersociety.org/10.1109/TVCG.2013.80)
- 4 [25] Jaillet, L., Simeon, T.: A PRM-based motion planning for dynamically changing
5 environments. In: *Proc. IEEE RSJ Int. Conf. Intell. Robot. Syst.*, vol. 2, pp. 1606–
6 1611 (2004)
- 7 [26] Ju, E., Choi, M.G., Park, M., Lee, J., Lee, K.H., Takahashi, S.: Morphable crowds.
8 *ACM Trans. Graph* **29**(6), 140 (2010). URL [http://dx.doi.org/10.1145/](http://dx.doi.org/10.1145/1882261.1866162)
9 [1882261.1866162](http://dx.doi.org/10.1145/1882261.1866162)
- 10 [27] Kallman, M., Mataric, M.: Motion planning using dynamic roadmaps. In: *Proc.*
11 *IEEE Int. Conf. Robot. Autom.*, vol. 5, pp. 4399–4404 (2004)
- 12 [28] Kapadia, M., Marshak, N., Shoulson, A., Badler, N.I.: ADAPT:
13 The agent development and prototyping testbed. *IEEE Transactions*
14 *on Visualization and Computer Graphics* **20**(7), 1035–1047 (2014).
15 doi:<http://doi.ieeecomputersociety.org/10.1109/TVCG.2013.251>
- 16 [29] Kapadia, M., Singh, S., Hewlett, W., Faloutsos, P.: Egocentric affordance fields in
17 pedestrian steering. In: *Proceedings of the Symposium on Interactive 3D Graphics*
18 *and Games*, pp. 215–223 (2009)
- 19 [30] Karamouzas, I., Heil, P., van Beek, P., Overmars, M.H.: A predictive collision avoid-
20 ance model for pedestrian simulation. In: *Motion in Games, Lecture Notes in Com-*
21 *puter Science*, vol. 5884, pp. 41–52. Springer (2009)
- 22 [31] Khatib, O.: Real-time obstacle avoidance for manipulators and mobile robots. *Int.*
23 *J. Robot. Res.* **5**, 90–98 (1986)
- 24 [32] Kim, S., Guy, S.J., Manocha, D., Lin, M.C.: Interactive simulation of dynamic
25 crowd behaviors using general adaptation syndrome theory. *ACM Symposium on*
26 *Interactive 3D Graphics and Games* pp. 55–62 (2012)
- 27 [33] Klupfel, H.: <http://sourceforge.net/projects/openpedsim/>
28 (2014)
- 29 [34] Kovar, L., Gleicher, M., Pighin, F.: Motion graphs. *ACM Trans. Graph.* **21**(3),
30 473–482 (2002)
- 31 [35] Kretz, T.: Pedestrian traffic: on the quickest path. *Journal of Sta-*
32 *tistical Mechanics: Theory and Experiment* **2009**(03), P03,012 (2009).
33 doi:10.1088/1742-5468/2009/03/P03012
- 34 [36] Lamarche, F., Donikian, S.: Crowd of virtual humans: a new approach for real
35 time navigation in complex and structured environments. *Computer Graphics Fo-*
36 *rum* **23**(3), 509–518 (2004). doi:10.1111/j.1467-8659.2004.00782.x.
37 URL <http://dx.doi.org/10.1111/j.1467-8659.2004.00782.x>

- 1 [37] Latombe, J.C.: Robot Motion Planning. Springer, Heidelberg (1991)
- 2 [38] Lee, K.H., Choi, M.G., Hong, Q., Lee, J.: Group behavior from video: a data-driven
3 approach to crowd simulation. In: Symposium on Computer Animation, pp. 109–
4 118 (2007)
- 5 [39] Marcus, M.P., Marcinkiewicz, M.A., Santorini, B.: Building a large annotated cor-
6 pus of english: The penn treebank. *Comput. Linguist.* **19**(2), 313–330 (1993). URL
7 <http://dl.acm.org/citation.cfm?id=972470.972475>
- 8 [40] Mononen, M.: Recast: navigation-mesh construction toolset for games. [http:
9 //code.google.com/p/recastnavigation/](http://code.google.com/p/recastnavigation/) (2009)
- 10 [41] Multon, F., France, L., Cani-Gascuel, M.P., Debunne, G.: Computer animation of
11 human walking: a survey, vol. 10, pp. 39–54 (1999)
- 12 [42] Narain, R., Golas, A., Curtis, S., Lin, M.C.: Aggregate dynamics for dense crowd
13 simulation. *ACM Trans. Graph.* **28**, 122:1–122:8 (2009)
- 14 [43] Newell, A.: Unified theories of cognition. Harvard University Press, Cambridge,
15 MA, USA (1990)
- 16 [44] Oliva, R., Pelechano, N.: Automatic generation of suboptimal navmeshes. In: J.M.
17 Allbeck, P. Faloutsos (eds.) *MIG, Lecture Notes in Computer Science*, vol. 7060,
18 pp. 328–339. Springer (2011). URL [http://dx.doi.org/10.1007/978-
19 3-642-25090-3](http://dx.doi.org/10.1007/978-3-642-25090-3)
- 20 [45] Ondřej, J., Pettré, J., Olivier, A.H., Donikian, S.: A synthetic-vision based steering
21 approach for crowd simulation. In: Proc. SIGGRAPH, pp. 123:1–123:9 (2010)
- 22 [46] Paris, S., Donikian, S.: Activity-driven populace: A cognitive approach to crowd
23 simulation. *Computer Graphics and Applications*, IEEE **29**(4), 34–43 (2009)
- 24 [47] Pelechano, N., Allbeck, J., Badler, N.: Controlling individual agents in high-density
25 crowd simulation. In: Symposium on Computer Animation, pp. 99–108 (2007)
- 26 [48] Pettré, J., Ondřej, J., Olivier, A.H., Cretual, A., Donikian, S.: Experiment-based
27 modeling, simulation and validation of interactions between virtual walkers. In:
28 Symposium on Computer Animation, pp. 189–198 (2009)
- 29 [49] Rabin, S.: *AI Game Programming Wisdom 4 (AI Game Programming Wisdom
30 (W/CD))*, 1 edn. Charles River Media (2008)
- 31 [50] Reynolds, C.: Flocks, herds and schools: A distributed behavioral model. In: Proc.
32 of SIGGRAPH (1987)
- 33 [51] Reynolds, C.W.: Steering behaviors for autonomous characters. *Game Developers
34 Conference* (1999)

- 1 [52] Samet, H.: Foundations of MultiDimensional and Metric Data Structures. Morgan
2 Kaufmann (2006)
- 3 [53] Schadschneider, A.: Cellular automaton approach to pedestrian dynamics - theory.
4 Pedestrian and Evacuation Dynamics pp. 75–86 (2002)
- 5 [54] Shao, W., Terzopoulos, D.: Autonomous pedestrians. In: Symposium on Computer
6 Animation, pp. 19–28 (2005)
- 7 [55] Shum, H.P.H., Komura, T., Shiraishi, M., Yamazaki, S.: Interaction patches for
8 multi-character animation. *ACM Trans. Graph* **27**(5), 114 (2008). URL <http://doi.acm.org/10.1145/1457515.1409067>
- 9
- 10 [56] Singh, S., Kapadia, M., Faloutsos, P., Reinman, G.: An open framework for de-
11 veloping, evaluating, and sharing steering algorithms. In: Proceedings of the 2nd
12 International Workshop on Motion in Games, pp. 158–169 (2009)
- 13 [57] Singh, S., Kapadia, M., Faloutsos, P., Reinman, G.: Steerbench: a benchmark suite
14 for evaluating steering behaviors. *Computer Animation and Virtual Worlds* **20**(5-6),
15 533–548 (2009). doi:10.1002/cav.277
- 16 [58] Snook, G.: Simplified 3D movement and pathfinding using navigation meshes. In:
17 Game Programming Gems, chap. 3, pp. 288–304. Charles River, Hingham, Mass.
18 (2000)
- 19 [59] Sun, H.C., Metaxas, D.N.: Automating gait generation. In: Proc. of ACM SIG-
20 GRAPH, pp. 261–270 (2001)
- 21 [60] Treuille, A., Cooper, S., Popović, Z.: Continuum crowds. In: Proc. of ACM SIG-
22 GRAPH, pp. 1160–1168 (2006)
- 23 [61] Ulicny, B., Thalmann, D.: Towards interactive real-time crowd behavior simulation.
24 *Computer Graphics Forum* **21**(4), 767–775 (2002)
- 25 [62] Weidmann, U.: Transporttechnik der fussgaenger. Tech. Rep. 90 (1993)
- 26 [63] Yang, Y., Brock, O.: Elastic roadmaps: globally task-consistent motion for au-
27 tonomous mobile manipulation. In: Proc. Robot. Sci. Syst., pp. 279–286 (2007)
- 28 [64] Zhang, P., Liu, H., Ding, Y.h.: Crowd simulation based on constrained
29 and controlled group formation. *The Visual Computer* pp. 1–14 (2013).
30 doi:10.1007/s00371-013-0900-7. URL [http://dx.doi.org/10.](http://dx.doi.org/10.1007/s00371-013-0900-7)
31 [1007/s00371-013-0900-7](http://dx.doi.org/10.1007/s00371-013-0900-7)
- 32 [65] Zheng, X., Zhong, T., Liu, M.: Modeling crowd evacuation of a building based
33 on seven methodological approaches. *Building and Environment* **44**(3), 437–445
34 (2009). doi:10.1016/j.buildenv.2008.04.002

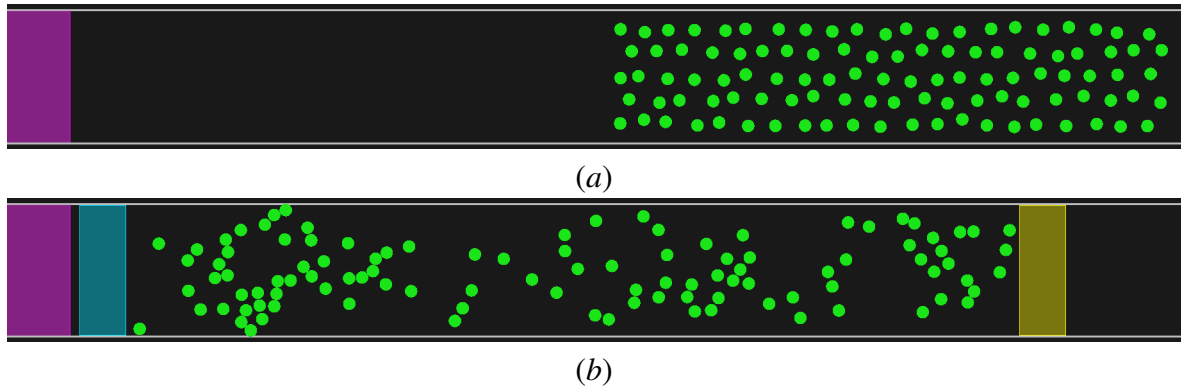


Figure 11 The scenario described in the appendix. Agents move from right to left down a corridor. The *effect* of periodic boundaries is realized with a `teleport Action`. As agents move toward the left-hand purple goal region, they enter the cyan box immediately preceding it. Upon entering this region, the agents are teleported back to the yellow region on the right. This motion allows agents to walk down the corridor indefinitely – well approximating periodic boundaries.

1 A. Menge Simulation Specification Example

2 Here we give an example of Menge’s simulation specification language. We present and
 3 discuss the complete description of a simple scenario: uni-directional flow down a corri-
 4 dor with periodic boundaries. Figure 11 shows the initial condition and some later point
 5 in the simulation. See below for a detailed description of the figures.

6 A.1. Scene Specification

7 Listing 1 provides the complete scene specification. It is responsible for defining the agent
 8 population and initial state.

9 **Line 1** The root element of the specification XML.

10 **Line 2** The declaration of the `SpatialQuery` type – in this case, a kd-tree.

11 **Lines 4-6** The specification of the global `Pedestrian Model` parameters, including
 12 those shared by all pedestrian models (`Common`) and those particular to the simple
 13 social force model (`Helbing`) and the predictive social force model (`Karamouzas`)¹¹.

14 **Lines 8-15** The definition of an “agent profile”, defining the space of values of agent
 15 b-state parameters. The profile is named, `group1`, for reference purposes.

16 **Lines 9-14** The per-agent `Pedestrian Model` parameters, including the shared
 17 parameters (`Common`), and for three particular implementations (`Helbing`,
 18 `Karamouzas`, `ORCA`).

¹¹The ORCA model does not have any global parameters.

1 **Line 10** The property allows for definitions of b-state parameters using a nu-
2 merical distribution. In this case, the preferred speed is defined as a nor-
3 mal distribution with a mean value and standard deviation of 1.3 m/s and
4 0.15 m/s, respectively.

5 **Lines 17-21** The instantiation of a group of agents. The number and position of each
6 agent is defined by the `Generator`, assigned b-state parameter values by its `ProfileSelector`,
7 and assigned an initial FSM-state by its `StateSelector`.

8 **Line 18** The `ProfileSelector` uses a `const` type. Which means that all
9 agents will be assigned the `group1` agent profile. In contrast, distribu-
10 tion-style `ProfileSelector` could assign a profile from a set of specified pro-
11 files.

12 **Line 19** The `StateSelector`, like the `ProfileSelector`, is of `const` type
13 and assigns all agents to the same initial FSM-state.

14 **Line 20** The `AgentGenerator` instantiates a hexagonal lattice of agent posi-
15 tions. The arguments specify the geometry of the lattice, average density, and
16 the approximate count of agents. In addition, it provides a displacement dis-
17 tribution to perturb the initial positions from the perfect lattice positions. The
18 noisy lattice can be seen in Figure 11(a).

19 **Lines 23-30** These define the obstacles in the environment. In this case, the type of the
20 `ObstacleSet` is `explicit`; each obstacle is explicitly defined in the specifica-
21 tion file (in contrast to being read from an external file).

22 **Lines 24-29** The definition of a single obstacle. The obstacle is a closed polygon,
23 defined by a two-dimensional vertex list. The order of vertices defines the
24 “inside” and “outside” of the obstacle.

25 **A.2. Behavior Specification**

26 The behavior specification includes the explicit instantiation of a particular BFSM, as
27 well as supporting data structures. Listing 2 contains the full BFSM specification for the
28 example scenario. The key feature to this BFSM is the teleport `Action` element on line
29 13. This is what creates the effect of periodic boundary conditions.

30 **Line 1** The root element of the behavior specification XML.

31 **Lines 2-4** The definition of a set of goals. A behavior specification can contain any num-
32 ber of such sets. Each goal set contains one or more `Goals`. Each goal set must
33 possess a unique, numerical id for referencing by other entities.

34 **Line 3** The single `Goal` defined in this scenario. In this case, the `Goal` is an AABB
35 (axis-aligned bounding box). Agents will always move to the closest point in
36 the goal region. The box is shown as the purpose region in Figure 11.

1 **Lines 6-9** The definition of the “walking” FSM-state. Uniquely identified by the name
2 Walk. The state also indicates that it is a non-final state – the simulation will not
3 end if there are agents in this FSM-state.

4 **Line 7** The `GoalSelector` for this FSM-state. When agents enter the state, they
5 are assigned a `Goal`. In this case, every agent explicitly is assigned a specific
6 `Goal` from a specific goal set (`Goal 0` from goal set 0).

7 **Line 8** The `VelocityComponent` which causes agents to move directly toward
8 their `Goal`. In this simple scenario, no more sophisticated mechanism is nec-
9 essary beyond simply walking straight to the goal.

10 **Lines 10-14** The definition of the “goal reached” FSM-state. This state serves a single
11 purpose, to discontinuously move (teleport) agents to a target region. Its various
12 components will reflect this purpose.

13 **Line 11** This FSM-state’s `GoalSelector` is of type `identity`. This means
14 that each agent’s `Goal` is the point at which the agent is when it enters the
15 state. This is useful for causing agents to hold position.

16 **Line 12** This FSM-state’s `VelocityComponent` is the zero type. Every agent
17 in this state will have the zero preferred velocity.

18 **Line 13** The `teleport Action` is assigned to this FSM-state. When agents en-
19 ter this FSM-state, the action is applied and the agents are moved to a random
20 point inside the box implied by the `min_x`, `max_x`, `min_y`, and `max_y` pa-
21 rameters (shown in yellow in Figure 11(b)).

22 **Lines 16-18** The definition of the transition from the `Walk` to `GoalReached` FSM-
23 states. This makes use of the implied transition `Target` element.

24 **Line 17** The transition `Condition` which causes an agent to move FSM-states.
25 This transition is taken when the agent enters an AABB. The region is shown
26 as the cyan box on the left in Figure 11(b). Because of this transition, no
27 agent will ever actually reach its goal, but will, instead, be teleported back to
28 the yellow region.

29 **Lines 19-21** The definition of the transition from the `GoalReached` back to `Walk`
30 FSM-states.

31 **Line 17** This transition `Condition` is an `auto` condition; it is the tautology. It
32 implies that any agent entering the `GoalReached` FSM-state will automat-
33 ically be transitioned to the `Walk` FSM-state. This type of automatic transi-
34 tions allows FSM-states to be introduced which have a one-time effect. This
35 transition is also the reason why the `GoalSelector` and `VelocityComponent`
36 in the `GoalReached` state are immaterial; they will never really be used.

1 **A.3. Additional Documentation**

2 The examples provided with Menge illustrate the various methods of creating and run-
3 ning scenes. Complete documentation of the Menge codebase is available at the project
4 website, <http://gamma.cs.unc.edu/Menge/>.

- 5 • An installation guide and Getting Started is available at <http://gamma.cs.unc.edu/Menge/learn/gettingStarted.html>
- 7 • Documentation on the Namespaces in Menge can be found at <http://gamma.cs.unc.edu/Menge/docs/code/menge/html/namespaces.html>.
- 9 • A complete class reference can be found at <http://gamma.cs.unc.edu/Menge/docs/code/menge/html/classes.html>.
- 11 • Documentation on the plugins included with Menge can be found at <http://gamma.cs.unc.edu/Menge/docs/code/PedPlugins/html/index.html>

1 A.4. Specification XML Files

Listing 1 Scene specification for a periodic hallway

```

2 1 <Experiment version="2.0">
3 2 <SpatialQuery type="kd-tree" test_visibility="false" />
4 3
5 4 <Common time_step="0.1" />
6 5 <Helbing agent_scale="2000" obstacle_scale="4000" reaction_time="
7 0.5" body_force="1200" friction="2400" force_distance
8 ="0.015" />
9 6 <Karamouzas orient_weight="0.8" fov="200" reaction_time="0.4"
10 wall_steepness="2" wall_distance="2"
11 colliding_count="5" d_min="1" d_mid="8" d_max="10" agent_force
12 ="4" />
13 7
14 8 <AgentProfile name="group1" >
15 9 <Common max_angle_vel="360" max_neighbors="10" obstacleSet="1"
16 neighbor_dist="5" r="0.19" class="2" pref_speed="1.04"
17 max_speed="2" max_accel="5" priority="0.0">
18 10 <Property name="pref_speed" dist="n" mean="1.3" stddev="0.15"
19 />
20 11 </Common>
21 12 <Helbing mass="80" />
22 13 <Karamouzas personal_space="0.69" anticipation="8" />
23 14 <ORCA tau="3.0" tauObst="0.15" />
24 15 </AgentProfile>
25 16
26 17 <AgentGroup>
27 18 <ProfileSelector type="const" name="group1" />
28 19 <StateSelector type="const" name="Walk" />
29 20 <Generator type="hex_lattice" anchor_x="1.5" anchor_y="0.0"
30 alignment="center" row_direction="y" density="1.8" width="
31 4.0" population="100" rotation="-90" displace_dist="n"
32 displace_mean="0.1" displace_stddev="0.03" />
33 21 </AgentGroup>
34 22
35 23 <ObstacleSet type="explicit" class="1">
36 24 <Obstacle closed="1" >
37 25 <Vertex p_x="-20" p_y="2.0" />
38 26 <Vertex p_x="20" p_y="2.0" />
39 27 <Vertex p_x="20" p_y="-2" />
40 28 <Vertex p_x="-20" p_y="-2" />
41 29 </Obstacle>
42 30 </ObstacleSet>
43 31 </Experiment>

```

Listing 2 Behavior specification for a periodic hallway

```

44 1 <BFSM>
45 2 <GoalSet id="0">
46 3 <Goal type="AABB" id="0" min_x="-20" max_x="-15" min_y="-2.0"
47 max_y="2" />

```

```
1 4 </GoalSet>
2 5
3 6 <State name="Walk" final="0" >
4 7   <GoalSelector type="explicit" goal_set="0" goal="0" />
5 8   <VelComponent type="goal" />
6 9 </State>
7 10 <State name="GoalReached" final="0">
8 11   <GoalSelector type="identity" />
9 12   <VelComponent type="zero" />
10 13   <Action type="teleport" dist="u" min_x="13.5" max_x="14" min_y=
11     "-1.5" max_y="1.5" />
12 14 </State>
13 15
14 16 <Transition from="Walk" to="GoalReached" >
15 17   <Condition type="AABB" min_x="-40" max_x="-13.5" min_y="-2.0"
16     max_y="2.0" inside="1" />
17 18 </Transition>
18 19 <Transition from="GoalReached" to="Walk" >
19 20   <Condition type="auto" />
20 21 </Transition>
21 22 </BFSM>
```