

A Comparative Study of Navigation Meshes

Wouter van Toll*
Utrecht University

Roy Triesscheijn
Utrecht University

Marcelo Kallmann
University of California, Merced

Ramon Oliva
Universitat Politecnica de Catalunya

Nuria Pelechano
Universitat Politecnica de Catalunya

Julien Pettré
INRIA

Roland Geraerts
Utrecht University




Figure 1: Navigation meshes computed for the Military environment. Regions are shown in different colors. From left to right: the Local Clearance Triangulation, the Explicit Corridor Map, the Clearance Disk Graph, Recast, NEOGEN, and a grid.

Abstract

A navigation mesh is a representation of a 2D or 3D virtual environment that enables path planning and crowd simulation for walking characters. Various state-of-the-art navigation meshes exist, but there is no standardized way of evaluating or comparing them. Each implementation is in a different state of maturity, has been tested on different hardware, uses different example environments, and may have been designed with a different application in mind.

In this paper, we conduct the first comparative study of navigation meshes. First, we give general definitions of 2D and 3D environments and navigation meshes. Second, we propose theoretical properties by which navigation meshes can be classified. Third, we introduce metrics by which the quality of a navigation mesh implementation can be measured objectively. Finally, we use these metrics to compare various state-of-the-art navigation meshes in a range of 2D and 3D environments.

We expect that this work will set a new standard for the evaluation of navigation meshes, that it will help developers choose an appropriate navigation mesh for their application, and that it will steer future research on navigation meshes in interesting directions.

Keywords: navigation meshes, path planning, comparative study

Concepts: •Computing methodologies → Mesh geometry models; Motion path planning; •General and reference → Metrics; Evaluation;

*Corresponding author. E-mail: W.G.vanToll@uu.nl.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org. © 2016 ACM. MiG '16, October 10-12, 2016, Burlingame, CA, USA ISBN: 978-1-4503-4592-7/16/10...\$15.00 DOI: <http://dx.doi.org/10.1145/2994258.2994262>

1 Introduction

Path planning for moving characters in virtual environments is a fundamental task in simulations and games. Modern applications feature increasingly large crowds of characters; each character needs to autonomously compute and follow a path while avoiding collisions with obstacles and other characters. This leads to many queries related to e.g. path planning, path following, point location, and collision avoidance [van Toll et al. 2015]. The simulation is expected to run in real-time despite all these demands. This stresses the need for high-quality data structures and algorithms.

Path planning for characters is different from robot motion planning, in which the high-dimensional *configuration space* of a robot [Lozano-Perez 1983] is often represented as a sampling-based graph (e.g. [Kavraki et al. 1996; LaValle 2006]). In our domain, the environments are typically three-dimensional, but characters are constrained to surfaces that are sufficiently flat to walk on. The behavior of characters may also include crawling, running, and other surface-based movement, but we will speak of *walking* and *walkable surfaces* for simplicity. The walkable surfaces of an environment form the *free space* \mathcal{E}_{free} , which is usually less complex than the environment itself.

A *navigation mesh* is a representation of \mathcal{E}_{free} as a set of (usually polygonal) regions, along with a graph that describes how these regions are connected. For path planning, characters first find an optimal path in the graph and then compute a suitable geometric route through the corresponding regions. A research topic of increasing importance is the *automatic* construction of a navigation mesh for any input environment. Current construction algorithms can roughly be placed into one of two categories: *voxel-based* algorithms that approximate the walkable surfaces from raw 3D geometry, or *exact* algorithms that require pre-processed input (e.g. a set of 2D layers) to compute a navigation mesh with a provable worst-case complexity. This difference makes the two categories difficult to compare. Furthermore, each method uses its own set of test environments to show its (dis)advantages. To steer subsequent research into interesting directions, an objective comparison between navigation meshes is required.

Goals and Contributions. In this paper, we conduct a comparative study of navigation mesh implementations by using the same hardware, quality metrics, and input environments for all methods. Our goal is to propose a way to objectively measure how suitable par-

ticular navigation meshes are for particular types of environments. Because navigation meshes have many applications with different requirements, it is difficult to propose a single criterion that can identify ‘the best’ navigation mesh. Instead, we present a collection of criteria, each of which is relevant for particular applications.

The main contributions of this paper are the following:

- We propose properties by which the data structures and algorithms of navigation meshes can be classified (Section 4).
- We present quantitative metrics to measure the quality and performance of a navigation mesh implementation for a given input environment (Section 5). In particular, we address the concept of coverage in 3D.
- We combine these metrics into a benchmark tool, and we use it to compare state-of-the-art navigation mesh implementations in a range of 2D and 3D environments (Section 6).

We emphasize that our goal is not to expose which navigation mesh implementation works best for particular input environments. Instead, by comparing navigation meshes using a common test platform and settings, we intend to set a standard for the analysis of navigation meshes and to expose interesting areas for future research.

2 Related Work

2.1 Navigation Meshes

Snook [2000] and Tozour [2002] were among the first to use the term ‘navigation mesh’ for a subdivision of the walkable space into polygonal regions. Because constructing a navigation mesh by hand is time-consuming and subject to human error, there has been increasing interest in automatically computing a navigation mesh from an input environment.

Voxel-based methods [Deusdado et al. 2008; Mononen 2014; Oliva and Pelechano 2013b; Pettré et al. 2005] usually take an unprocessed 3D environment as their input. To construct a navigation mesh, they discretize the environment into a 3D grid of *voxels* using GPU techniques, extract the voxels that correspond to walkable regions, and summarize this information in a navigation mesh that *approximates* the geometry of \mathcal{E}_{free} . This reconstruction is based on the assumption that the environment has a single direction of gravity \vec{g} , and that characters are cylinders with a fixed height and (sometimes) a fixed radius. Voxel-based methods can handle arbitrary 3D geometry; the approximation automatically resolves issues caused by e.g. intersecting obstacles. However, the precision and efficiency of these methods depends to a certain degree on the grid resolution. The quality of the navigation mesh depends on how well the free space is extracted from the 3D geometry.

Exact methods [Geraerts 2010; Hale et al. 2008; Kallmann 2014; Oliva and Pelechano 2011; van Toll et al. 2011; Tozour 2002] require that the exact geometry of \mathcal{E}_{free} is already known, and that this free space has been pre-processed into one or more planar layers. In exchange, they represent their input *precisely*, and they often have provable worst-case construction times and storage sizes, which implies better scalability to large environments. However, extracting \mathcal{E}_{free} from a 3D environment without using approximations is still a topic of ongoing research [Polak 2016].

Researchers have also investigated navigation meshes for other types of geometry or movement. An environment can be subdivided into 3D volumes to encode height differences and variable vertical clearance [Hale and Youngblood 2009; Lamarche 2009]. Alternatively, one could perform crowd simulations on arbitrary surfaces with no consistent direction of gravity [Ricks and Egbert

2014; Berseth et al. 2015]. Other methods allow characters to jump between surfaces by either checking for jumping possibilities on the fly [Lopez et al. 2012] or annotating a navigation mesh with jump links beforehand [Budde 2013]. However, these extensions are outside the scope of our comparative study.

Navigation meshes are useful for simulating crowds of characters with individual properties and goals. Crowd simulation is a large research field with many components including path planning, collision avoidance between characters, animation, and the evaluation of realism. Several books exist that give good overviews of this field [Ali et al. 2013; Kapadia et al. 2015; Pelechano et al. 2016; Thalmann and Musse 2013]. Also, there are multiple crowd simulation frameworks in which navigation meshes play a central role [Curtis et al. 2014; van Toll et al. 2015]. In this paper, we focus on the fundamental properties of navigation meshes, so we will not treat the field of crowd simulation in more detail. However, crowd simulation is an important motivation for many of the properties and metrics that we will propose.

2.2 Comparative Studies

Our comparative study of navigation meshes is inspired by comparisons for *other* aspects of path planning and crowd simulation.

Sturtevant [2012] has developed a test set of environments for 2D *grid-based* path planning. This set includes mazes of various complexities and levels from computer games, and it is often used to analyze variants of the A* search algorithm [Hart et al. 1968]. Although we study general navigation meshes rather than grids, we will also include grid-based environments in our experiments.

SteerBench [Singh et al. 2009] focuses on local behavior such as collision avoidance. It presents a comprehensive set of scenarios that local methods are expected to solve, such as two characters crossing paths, or characters switching places in a narrow corridor. Given the output of a crowd simulation for such a scenario, SteerBench can compute metrics such as the distance that all characters traverse and the amount of energy that they spend. However, the results need to be put in perspective because steering methods typically have many parameters and implementation choices.

In this paper, we compare navigation meshes in a similar way based on metrics, input environments, and a single test platform. We will see that parameter settings are influential in our study as well.

3 Definitions

In this section, we give definitions of environments and navigation meshes. This is useful because all existing papers and algorithms use slightly different terminology; the content of Sections 4 and 5 requires unified definitions.

3.1 2D Environment

A *2D environment* is a finite subset of the 2D plane with polygonal holes; we refer to these holes as obstacles. We will not consider point or line segment obstacles in this paper. The obstacle space \mathcal{E}_{obs} is the union of all obstacles. Its complement is the free space \mathcal{E}_{free} . Let n be the number of vertices required to define \mathcal{E}_{obs} or \mathcal{E}_{free} using simple polygons. We call n the *complexity* of \mathcal{E} .

In our experiments, we want to treat 2D and 3D environments similarly. We will therefore embed our 2D environments in \mathbb{R}^3 by assigning a height component of zero to each vertex.




Figure 2: Different representations of an environment, and an example of its navigation mesh. (a) A 3D environment consists of unprocessed 3D geometry. (b) A walkable environment (WE) contains only walkable surfaces. It is the free space \mathcal{E}_{free} of a 3D environment. (c) A multi-layered environment (MLE) is a WE subdivided into layers. (d) A navigation mesh is a description of a WE for path planning purposes.

3.2 3D Environment

In this paper, a 3D environment (3DE) is a raw collection of polygons in \mathbb{R}^3 . These polygons may include floors, ceilings, walls, or any other type of geometry. Figure 2a shows an example. To define the free space \mathcal{E}_{free} of a 3DE, we need to specify various parameters that describe on which surfaces a character may walk. Examples of such parameters are the maximum slope with respect to the direction of gravity, the maximum height difference between nearby polygons (e.g. the maximum step height of a staircase), and the required vertical distance between a floor and a ceiling.

Characters are typically approximated by cylinders. Some navigation meshes use a predefined character radius to determine \mathcal{E}_{free} . In this paper, we will use a radius of zero to enable an objective comparison to other navigation meshes.

3.3 Walkable Environment

A walkable environment (WE) is a set of interior-disjoint polygons in \mathbb{R}^3 on which characters can stand and walk. Thus, a WE is a clean representation of the free space \mathcal{E}_{free} of a 3DE, based on the filtering parameters and character properties mentioned earlier. Any two polygons are directly connected if and only if characters can walk directly between them. Figure 2b shows an example. In our experiments, all environments will be WEs, so we know beforehand which area should be covered by a navigation mesh.

All polygons in the WE have a maximum slope with respect to the ground plane P , which is the plane perpendicular to the gravity direction \vec{g} . It is common for a navigation mesh to project the length of a path onto P as well, i.e. to ignore height differences along a path during planning. Therefore, in this paper, we will not judge a navigation mesh by its preservation of height differences.

The complexity of a WE is the total number of polygon vertices. The free space \mathcal{E}_{free} is simply the set of polygons itself. The obstacle space \mathcal{E}_{obs} can be thought of as ‘anything beyond the boundary of \mathcal{E}_{free} ’, but (unlike in 2D) it is difficult to represent or visualize because it does not necessarily consist of polygons on a plane.

It is important to see that a WE can be self-overlapping when projected onto the ground plane P , i.e. it is not guaranteed that all surfaces are visible from a single top view. This strongly influences the construction of navigation meshes: an algorithm for 2D environments cannot easily be applied to WEs in general.

3.4 Multi-Layered Environment

Some navigation meshes require the WE to be subdivided into 2D components. A multi-layered environment (MLE) [Pettré et al. 2005; van Toll et al. 2011] is a subdivision of a WE into layers

such that the walkable polygons of each individual layer are non-overlapping when projected onto P . The layers are connected by line segments. An example of an MLE is shown in Figure 2c.

The complexity of an MLE is given by the number of layers l , the number of connections k , and the number of boundary vertices n in all layers combined. Converting a WE to an MLE with a minimal number of connections is NP-hard, but good results can be obtained using heuristics [Hillebrand 2012]. In our experiments, we will subdivide all WEs into layers to facilitate the construction of navigation meshes.

Finally, note that a 2D environment is a special case of an MLE with only one layer (or, equivalently, a WE that can be projected onto P without overlap). Section 5 will define quality metrics for WEs in general, so that 2D and 3D input can be treated equally.

3.5 Navigation Mesh

Now that we have a definition of the free space \mathcal{E}_{free} , we can define a navigation mesh as a tuple $\mathcal{M} = (\mathcal{R}, \mathcal{G})$:

- $\mathcal{R} = \{R_0, R_1, \dots\}$ is a collection of geometric regions in \mathbb{R}^3 that represents \mathcal{E}_{free} . Each region R_i is P -simple, by which we mean that a region cannot intersect itself when projected onto the ground plane P .
- $\mathcal{G} = (V, E)$ is an undirected graph that describes how characters can navigate between the regions in \mathcal{R} .

Figure 2d shows an abstract example of a navigation mesh. For many navigation meshes, \mathcal{R} consists of non-overlapping simple polygons, and \mathcal{G} is the dual graph of \mathcal{R} , with one vertex per region and one edge per pair of adjacent region sides. However, other possibilities exist. In the Clearance Disk Graph [Pettré et al. 2005], \mathcal{R} consists of overlapping disks, and \mathcal{G} contains an edge wherever two disks overlap. The Explicit Corridor Map [van Toll et al. 2011] is explicitly defined as a graph, and the mesh regions can be derived from its annotations. Still, all meshes have in common that \mathcal{R} and \mathcal{G} can be obtained from their representation in some way.

4 Properties of Navigation Meshes

In this section, we propose a set of properties that describe a navigation mesh’s data structure, algorithms, and limitations. These properties do not depend on a specific implementation or environment. They can serve as a ‘checklist’ to simplify choosing an appropriate mesh for a particular application.

Region type The type of regions in \mathcal{R} , e.g. triangles or disks.

Graph type A description of the path planning graph \mathcal{G} , e.g. ‘the dual graph of \mathcal{R} ’ or ‘the medial axis of \mathcal{E}_{free} ’.

Overlap Whether or not the regions in \mathcal{R} can overlap by definition. Having overlapping regions is generally discouraged because geometric algorithms that assume non-overlapping regions may not work properly. Also, a query point (or an agent) can be in multiple regions at the same time in case of overlap, which may complicate path planning and crowd simulation.

Pipeline The conversion pipeline performed by the construction algorithm, e.g. ‘from a 2D environment to a navigation mesh’ or ‘from a 3DE via an MLE to a navigation mesh’. We also indicate whether this pipeline is voxel-based or exact.

Parameters The parameters that the user needs to set for the construction algorithm of the navigation mesh. Having fewer parameters implies a more automated process for computing the mesh. These parameters are often related to the filtering process that extracts the walkable surfaces from the 3D geometry.

Computational complexity The asymptotic construction time of the navigation mesh. This is usually expressed in terms of the environment complexity or a grid resolution.

Storage complexity The asymptotic size of the navigation mesh data structure.

Clearance Whether or not the navigation mesh supports the computation of paths with an arbitrary clearance from obstacles, i.e. paths for disks with an arbitrary radius.

Dynamic updates Whether or not the navigation mesh supports dynamic insertions and deletions of obstacles.

Due to space constraints, we cannot include our full theoretical comparison of the navigation meshes that will be used in our experiments. Instead, Section 6.1 will describe each method briefly, and Table 1 classifies each method by the properties described above. We will include the full comparison in an extended publication.

5 Quality Metrics for Navigation Meshes

For a navigation mesh $\mathcal{M} = (\mathcal{R}, \mathcal{G})$ that has been constructed for an environment using a particular implementation, we want to objectively measure the quality. Many possible evaluation criteria exist, and each application area may have its own view of what is good or desirable. In this paper, we choose to focus on the navigation mesh itself and on the performance of its construction algorithm. We will present metrics that answer the following questions:

1. (Coverage) How accurately do the regions of \mathcal{R} cover the geometry of \mathcal{E}_{free} ? If parts of the free space are not covered, characters might not find a path in \mathcal{G} even though a path exists in \mathcal{E}_{free} . If parts outside the free space are covered, characters might accidentally find paths through obstacles.
2. (Connectivity) How accurately does the graph \mathcal{G} represent the connectivity of \mathcal{E}_{free} ? This question is related to coverage because it determines whether or not appropriate paths can be found; however, it concerns topology rather than geometry.
3. (Complexity) How efficiently does \mathcal{M} represent \mathcal{E}_{free} , i.e. how ‘compact’ is the mesh? This can refer to the size of the graph \mathcal{G} (a smaller graph allows faster path planning queries) or to the complexity of each individual region in \mathcal{R} (simpler regions allow faster basic operations such as point location). It depends on the application which property is more desirable.
4. (Performance) How efficiently is \mathcal{M} computed in terms of time and memory? An efficient algorithm allows the construction of navigation meshes in interactive applications such as

level editors. Even if the navigation mesh is precomputed in an off-line stage, performance is still desirable.

Of course, many other questions are interesting, e.g. questions related to the performance of path planning queries, or to the quality or realism of paths. We will discuss a number of options in Section 7 as suggestions for future work.

Analyzing coverage and connectivity is only useful for voxel-based navigation meshes that attempt to ‘discover’ \mathcal{E}_{free} themselves; exact methods are known to yield perfect coverage. Also, some properties can only be analyzed if the ground truth (the structure of \mathcal{E}_{free}) is known. Therefore, each input environment in our experiments will be a ‘clean’ walkable environment, i.e. a manifold that contains only walkable polygons. While this implies that voxel-based methods will not fully use their advantage of handling raw (non-clean) 3D geometry, we believe that using the same input for all methods yields a more objective comparison.

Since the outcome of each metric depends on implementation details, the results should always be judged in combination with the theoretical properties of Section 4.

5.1 Coverage

The first set of metrics describes how well the free space is covered. Coverage is a complicated property to evaluate due to the 3D structure of \mathcal{R} and \mathcal{E}_{free} . We need to introduce a number of concepts before we can define actual metrics. These concepts are based on the assumption that the environment has a consistent direction of gravity. Coverage is the only category of metrics in which this assumption comes into play.

5.1.1 Mapping the Navigation Mesh onto the Free Space

Comparing the geometry of \mathcal{R} to the geometry of \mathcal{E}_{free} requires us to vertically map these two structures onto each other. This is straightforward if the environment consists of a single layer because everything can then be projected onto the ground plane P . However, for general WEs in \mathbb{R}^3 , mapping \mathcal{R} onto \mathcal{E}_{free} is ambiguous. In an abstract sense, there should be a function m such that, for any point p in a navigation mesh, $m(p)$ vertically maps p to an appropriate point in \mathcal{E}_{free} if possible (and if not, p is assumed to lie in \mathcal{E}_{obs}). Several choices can be made here, such as the maximum allowed height difference between p and $m(p)$. We will describe our own implementation of m in Section 6.

Using the function m , we define a *mapped region* R_i^* as a region R_i that has been mapped onto \mathcal{E}_{free} wherever possible, i.e. $R_i^* = \{m(p) \mid p \in R_i \text{ and } m(p) \text{ exists}\}$. Because each mapped region is a subset of \mathcal{E}_{free} , we can use the mapped regions to define unions, coverage, and overlap. Let the *mapped region set* \mathcal{R}^* be a version of \mathcal{R} in which all regions have been mapped onto \mathcal{E}_{free} , i.e. $\mathcal{R}^* = \{R_i^* \mid R_i \in \mathcal{R}\}$. The regions in \mathcal{R}^* may overlap: in that case, some points of \mathcal{E}_{free} are represented more than once.

5.1.2 Computing the Projected Area

Because we ignore height differences in our problem domain, our coverage metrics are based on projected areas onto the ground plane P . We define the *projected area* of a shape S as follows:

- If S does not overlap itself when projected onto P (i.e. if S is a P -simple shape as defined in Section 3.5), the projected area $\|S\|$ is the signed area of the projection of S onto P .
- Otherwise, let $\{S_0, \dots, S_{s-1}\}$ be any subdivision of S into P -simple shapes. The projected area of S is the sum of projected areas of these components, i.e. $\|S\| = \sum_i \|S_i\|$.

We assume that \mathcal{E}_{free} is given as a subdivision into P -simple shapes, such that $||\mathcal{E}_{free}||$ can be computed.

5.1.3 Coverage Metrics

We introduce three coverage metrics. Each metric has a regular version M with range $\mathbb{R}_{\geq 0}$ and a normalized version M' with range $[0, 1]$, as described below.

Free space covered The area of \mathcal{E}_{free} that is correctly covered by at least one navigation mesh region. Because the regions in \mathcal{R}^* may overlap and we do not want to count overlapping regions twice, we first compute the union of \mathcal{R}^* in \mathbb{R}^3 . High coverage is desirable: it allows characters to use more of \mathcal{E}_{free} .

$$Cov = ||\bigcup_i R_i^*|| \text{ and } Cov' = \frac{Cov}{||\mathcal{E}_{free}||}$$

Incorrect area The area of the mesh that ‘overshoots’ \mathcal{E}_{free} and lies in the obstacle space. Intuitively, this is the difference between \mathcal{R} and the part of \mathcal{R} that can be mapped onto \mathcal{E}_{free} . Ideally, the incorrect area should be zero because areas outside \mathcal{E}_{free} should not be accessible to characters.

$$A_{inc} = \sum_i (||R_i|| - ||R_i^*||) \text{ and } A'_{inc} = \frac{A_{inc}}{\sum_i ||R_i||}$$

Note: while it may seem more intuitive to express this metric as ‘the area of \mathcal{E}_{obs} that is covered’, this would be impossible because (for WEs in 3D) \mathcal{E}_{obs} does not have an area.

Overlap The amount of overlap among the regions in the navigation mesh. Intuitively, overlap is the sum of all region areas minus the area that is covered at least once. Because coverage is only defined properly inside \mathcal{E}_{free} , overlap is also based on the mapped region set \mathcal{R}^* . The normalized version indicates which fraction of \mathcal{R}^* is redundant.

$$Ov = \sum_i ||R_i^*|| - ||\bigcup_i R_i^*|| \text{ and } Ov' = \frac{Ov}{\sum_i ||R_i^*||}$$

If a navigation mesh is deliberately based on overlapping regions (e.g. [Pettré et al. 2005]), then this metric simply indicates how much space is covered more than once. Otherwise, overlap may indicate an implementation bug, which is not the focus of our comparative study.

5.2 Connectivity

The second set of metrics analyzes how well the graph $\mathcal{G} = (V, E)$ represents the dual graph of \mathcal{E}_{free} .

Connected components The number of connected components in \mathcal{G} . Ideally, this value is equal to the number of connected components in \mathcal{E}_{free} . Having more components implies that not all adjacencies in \mathcal{E}_{free} are captured. Having fewer components implies that regions have been made adjacent when there are actually obstacles in-between.

Boundaries The number of environment boundaries perceived by the navigation mesh. Ideally, this value is equal to the actual number of boundaries of \mathcal{E}_{free} . It can be computed by traversing the graph \mathcal{G} , checking the corresponding regions in \mathcal{R} , and collecting the region edges that are not shared by multiple regions. The number of boundaries is the number of closed loops that are traced. Note: if the number of boundaries is perfect, the geometry of \mathcal{R} is not necessarily correct.

5.3 Complexity

The third set of metrics measures how efficiently the navigation mesh represents \mathcal{E}_{free} . The size of \mathcal{G} , the number of regions, and the complexity of these regions may have implications for the efficiency of path planning and crowd simulation.

Vertices, # Edges The number of vertices and the number of edges in \mathcal{G} , i.e. $|V|$ and $|E|$. A larger graph implies that path planning queries (and other algorithms that browse the graph) typically take more time to answer. Therefore, lower numbers imply faster path planning.

Regions The number of regions in the navigation mesh: $|\mathcal{R}|$. This indicates how efficiently the free space is represented by elementary parts. It also suggests how often a character in the simulation may move from one region to another. Moving to another region typically triggers computational overhead in the simulation. Hence, having fewer regions may cause some aspects of the simulation to run more efficiently. Note that $|\mathcal{R}| = |V|$ if \mathcal{G} is simply the dual graph of \mathcal{R} .

Region complexity The number of floating-point numbers required to describe the regions in \mathcal{R} . Since we treat regions as shapes in \mathbb{R}^3 , we will say that a polygonal region with p vertices has complexity $3p$. A disk has a complexity of 4 because it can be defined by a center point in \mathbb{R}^3 and a radius. Naturally, other choices are possible as well. Some navigation meshes have extra annotations, such as the maximum allowed radius of a character for an edge traversal [Kallmann 2014]. We will not include such annotations in this metric.

We measure three variants: the average complexity among all regions, the standard deviation, and the total complexity of all regions combined. A low region complexity implies that geometric operations within these regions are computationally cheap. If a mesh has a small number of regions, a low region complexity, and high coverage, then it is a very efficient description of \mathcal{E}_{free} .

5.4 Performance

The final set of metrics concerns the practical performance of the navigation mesh implementations. One issue to take into account is that voxel-based methods perform more steps than exact methods. Another issue is that different implementations are in drastically different states: some are a ‘proof of concept’ for research purposes, while others are highly optimized for the gaming and simulation industry. Still, these metrics can indicate if an implementation corresponds to the asymptotic complexity of a navigation mesh, and how well a navigation mesh scales to large or complex environments.

Construction time The time (in milliseconds) spent on computing the navigation mesh. Naturally, fast construction is encouraged because it makes the algorithm suitable for interactive applications.

Memory usage The maximum amount of memory (in MB) used during the execution of the program. A small value implies that the mesh can be computed in situations with limited resources, e.g. on a game console with little working memory.

To obtain more reliable results, we will run each navigation mesh program 10 times and report the average values and standard deviations. This is not needed for the other categories of metrics because the output of each program is deterministic.

6 Experimental Comparison

In this section, we use our metrics to experimentally compare various navigation meshes in a range of environments. All experiments were run on a Windows 7 PC with a 3.20 GHz Intel i7-3930K CPU, an NVIDIA GeForce GTX 680 GPU, and 16 GB of RAM.

6.1 Navigation Meshes

We compare five state-of-the-art navigation mesh implementations and one extra baseline method. The first two navigation meshes are exact; the others are voxel-based and cover the full 3D pipeline. We currently include only the navigation meshes that are designed specifically for the environments described in Section 3, and for which we could obtain robust source code from their respective authors. Naturally, we encourage others to join this comparison.

Table 1 summarizes each navigation mesh based on the properties from Section 4. Most of these navigation meshes depend on various parameters. The parameter settings that we use in our experiments are listed in Appendix A in the supplementary file. Examples of the output for each method are shown in Figure 1.

Local Clearance Triangulation (LCT). The LCT [Kallmann 2014] subdivides a 2D environment of complexity n into $\mathcal{O}(n)$ triangles by first computing a constrained Delaunay triangulation and then adapting it in $\mathcal{O}(n^2)$ worst-case time. The tested implementation runs in $\mathcal{O}(n\sqrt{n})$ expected time by using a special point-location method. These triangles are the regions of \mathcal{R} , and \mathcal{G} is their dual graph. Triangle edges are annotated with clearance values to allow path planning for characters of an arbitrary radius. The LCT also supports dynamic updates. An extension to MLEs has not been developed, but an approach similar to the Explicit Corridor Map (described next) should be possible.

The LCT uses line segments as input and output, so our benchmark program needs to translate between walkable regions and boundary representations. These steps will not be included in our time measurements.

Explicit Corridor Map (ECM). The ECM [van Toll et al. 2011] is an exact navigation mesh. Its graph $\mathcal{G} = (V, E)$ is the medial axis of \mathcal{E}_{free} , where V contains the medial axis vertices of degree 1, 3, or higher, and each edge E is a sequence of medial axis arcs. Certain medial axis points are annotated with their two nearest obstacle points, which induces a subdivision of \mathcal{E}_{free} into polygonal regions.

The ECM enables path planning for characters of any radius, and the nearest obstacle in any region can be found in constant time. The ECM also supports dynamic updates. For a 2D environment of complexity n , the ECM has size $\mathcal{O}(n)$ and is computed in $\mathcal{O}(n \log n)$ time. For an MLE with k connections, it has size $\mathcal{O}(kn)$ and can be computed in $\mathcal{O}(kn \log n)$ time by iteratively opening the connections. In exchange for its advantages, the ECM is mathematically more complex than e.g. a triangulation.

Clearance Disk Graph (CDG). Pettré et al. [2005] presented the first *voxel-based* navigation mesh for 3D environments. In this paper, we refer to it as the Clearance Disk Graph (CDG). The CDG uses voxelization to approximate the areas where characters can stand. Next, the voxels are extracted for which the clearance is locally largest. These form an approximation of the medial axis of \mathcal{E}_{free} , and each cell represents an obstacle-free disk. A subset of these disks is chosen as the set of regions \mathcal{R} , and the graph \mathcal{G} describes the disks and their overlap. Extra disks (that do not lie on the medial axis) can be added to improve coverage.

The asymptotic construction time of the CDG is difficult to assess because the algorithm relies on rendering techniques. Also, the

number of disks cannot be expressed in terms of the environment complexity, but it is at least limited by the number of voxels S .

Recast. The Recast Navigation toolkit [Mononen 2014] is a popular choice for game development that is also used in the Unity3D game engine [2016]. Like the CDG, it uses voxelization to approximate \mathcal{E}_{free} . However, Recast converts the walkable voxels to non-overlapping convex polygonal regions. This conversion involves many parameters that the user needs to tweak to get the best results. One parameter is the character radius, which is subtracted from the navigation mesh during its construction. As mentioned, we will use a radius of zero to allow a fair comparison to other methods.

Recast computes two versions of the navigation mesh: a coarse mesh used for path planning, and a detailed mesh with more accurate height differences. We will use the coarse mesh to determine \mathcal{R} and \mathcal{G} , and the detailed mesh to measure coverage.

NEOGEN. The NEOGEN method [Oliva and Pelechano 2013b] also starts with voxelization, but it groups walkable voxels into 2D layers. Next, the method obtains a more precise floorplan for each layer in a way that does not depend on the voxel size. Compared to Recast, the overall precision of NEOGEN is therefore less dependent on the grid resolution. Based on these floorplans, an exact 2D algorithm [Oliva and Pelechano 2011] is used to compute the final navigation mesh. This 2D algorithm subdivides the layer into convex polygons in $\mathcal{O}(nr)$ time, where $r < n$ is the number of convex polygon vertices in the input. In our experiments, for simplicity, we will use the voxel-based method in both MLEs and 2D environments.

A contribution of NEOGEN is the *convexity relaxation* parameter that can be used to allow slightly non-convex regions. This decreases the total number of regions in exchange for having more complex region shapes. Clearance information can also be added to the navigation mesh if desired [Oliva and Pelechano 2013a].

Grid. As a baseline for our comparison, we have implemented a simple grid method. It voxelizes the environment similarly to Recast and NEOGEN, but it uses the walkable voxels directly as navigation mesh regions. Therefore, each region in \mathcal{R} is a square.

We include this method because grids are still frequently used for path planning. They are easy to implement in 2D environments and WEs [Sturtevant 2011], and they are a common choice for games that are grid-based by design [Sturtevant 2012]. Another advantage is that algorithms such as A* search can be optimized for grids [García et al. 2014; Harabor and Grastien 2011; Lee and Lawrence 2013; Sturtevant and Rabin 2016]. Many variants of A* are either designed with grids in mind [Botea et al. 2004] or explained in terms of grids [Koenig and Likhachev 2002; Koenig et al. 2004; Likhachev et al. 2005]. However, grids are typically more dense than other navigation meshes, which makes them less appropriate for planning many paths in real-time.

6.2 Implementation

We have converted each navigation mesh program to a stand-alone executable that reads an input file, computes a navigation mesh, and returns the result. We have written a benchmark tool that communicates with these programs, converts environments between different file formats, and calculates all metrics. Also, the CDG requires the walkable surfaces to be visible from all sides; to ensure this, we extrude all input polygons downwards by a small amount.

An important detail is our choice of the mapping function m that is used to compute coverage. For a point p in the navigation mesh, we define $m(p)$ as the *nearest* point in \mathcal{E}_{free} above or below p up to a threshold distance T . The threshold distance is required to prevent

Navigation mesh	Region type	Graph type	Overlap	Pipeline	Parameters	Computational complexity	Storage complexity	Dynamic updates	Arbitrary clearance
LCT	Triangles	Dual of \mathcal{R}	No	2D $\rightarrow \mathcal{M}$ (exact)	None	2D: $\mathcal{O}(n\sqrt{n})$ (expected)	$\mathcal{O}(n)$	+	+
ECM	Polygons	Medial axis	No	2D/MLE $\rightarrow \mathcal{M}$ (exact)	None	2D: $\mathcal{O}(n \log n)$ MLE: $\mathcal{O}(kn \log n)$	$\mathcal{O}(n)$ $\mathcal{O}(kn)$	+	+
CDG	Disks	Dual of \mathcal{R}	Yes	3D $\rightarrow \mathcal{M}$ (voxel-based)	3D filtering Voxel precision Min character radius Min/max disk size	?	$\mathcal{O}(S)$	+/-	+
Recast	Convex polygons	Dual of \mathcal{R}	No	3D $\rightarrow \mathcal{M}$ (voxel-based)	3D filtering Voxel precision Region refinement Character radius	?	?	+/-	-
NEOGEN	(Convex) polygons	Dual of \mathcal{R}	No	3D \rightarrow MLE $\rightarrow \mathcal{M}$ (voxel-based + exact)	3D filtering Voxel precision Convexity relaxation	2D: $\mathcal{O}(n^2)$ MLE: $\mathcal{O}(n^2)$ 3D: ?	$\mathcal{O}(n)$ $\mathcal{O}(n)$	+/-	+/-
Grid	Squares	Dual of \mathcal{R}	No	3D $\rightarrow \mathcal{M}$ (voxel-based)	3D filtering Voxel precision	?	$\mathcal{O}(S)$	+/-	-

Table 1: Overview of the navigation meshes compared in this paper. For the rightmost columns, ‘+’ means that a property is supported in the current implementation, ‘+/-’ means that a property could be added in theory, and ‘-’ means that a property is not supported by definition.

erroneous points of \mathcal{R} from getting mapped onto surfaces that are too far away. We choose a value of $T = 1\text{m}$ because the vertical clearance is at least 2 m in all our test environments. Admittedly, this choice for m requires that the height coordinates of the navigation mesh are sufficiently close to the ground truth. It may fail in environments with gradual yet large height differences that are not captured by the navigation mesh. (In 2D environments, T can be ignored because a vertical mapping is already unambiguous.)

We have implemented our coverage metrics using a CGAL-based program [CGAL 2016] that can compute the intersection of two OBJ files based on the threshold distance T . For this program, we approximated the disks of the CDG by polygons of 16 vertices. We used *inner* approximations: the approximated disks were smaller than the actual disks. This leads to slightly lower numbers for coverage, incorrect area, and overlap, but the chosen precision is sufficient for comparative purposes.

6.3 Environments

We have computed navigation meshes for the 2D and 3D input environments shown in Figures 3 and 4, which range in scale and complexity. Due to space constraints, we focus on environments that have appeared in previous publications. To test for scalability, we have also added randomly generated 2D mazes of various sizes, inspired by Sturtevant [2012]. In an extended publication, we will propose a more comprehensive set of benchmark environments. We have converted each environment to a clean representation of \mathcal{E}_{free} , subdivided into layers whenever necessary. The corresponding OBJ files are included in the supplementary file of this paper.

6.4 Results

The full set of results can be found in the attached supplementary file. Tables 2 and 3 show the results for coverage and connectivity, and Tables 4 and 5 show the results for complexity and performance. We will now highlight a number of observations.

Coverage. We have chosen our parameters to maximize coverage (see Appendix A). Still, in terms of absolute values, the voxel-based methods sometimes missed large areas or covered large incorrect parts, up to hundreds of square meters in large environments. However, the *relative* coverage was still high (typically over 90%).

The maze environments are an exception: even though their free space was perfectly aligned with grid cells of $1 \times 1\text{m}$, the CDG and Recast could not capture them accurately. However, these mazes are quite detailed relative to their size, so a finer grid resolution could improve the results. NEOGEN generally yielded better coverage due to its extra processing step per layer. It would be interesting to let methods automatically choose an appropriate resolution based on a user-specified balance between coverage and performance. A theoretically stronger alternative would be to reconstruct \mathcal{E}_{free} without relying on a grid resolution.

Connectivity. Recast and NEOGEN captured connectivity quite well for most environments, except in the mazes where each accidental gap causes parts to become disconnected. For the CDG, the graph usually contained many connected components, and gaps in the covered space led to a large number of boundaries. The grid also contained unexpected gaps at times, due to small errors in the current implementation. Still, the grid method works sufficiently well for the purpose of a comparison.

It is motivating to see that bad connectivity values corresponded to navigation meshes that were also *visually* incorrect. For example, Recast gave overlapping regions in some mazes, and it accidentally connected layers vertically in the *Tower* environment. Still, we acknowledge that metrics can never fully replace visual inspection.

Complexity. NEOGEN typically yielded the smallest graph in exchange for the highest average region complexity. This makes sense because the algorithm is deliberately designed to produce a small number of regions [Oliva and Pelechano 2011]. Its convexity relaxation parameter could enlarge this effect. The ECM often produced smaller graphs than the LCT, while the LCT usually had a lower total region complexity.

Recast appears to average between graph complexity and region complexity using our current settings. The method contains several parameters (such as the maximum number of vertices per region) with which this balance can be controlled. Recast and NEOGEN may not always capture all details of \mathcal{E}_{free} , but this does generally lead to simpler navigation meshes. This is useful for applications in which low storage requirements and fast path planning are more important than perfect coverage. In fact, exact methods could also benefit from a pre-processing step that simplifies \mathcal{E}_{free} .

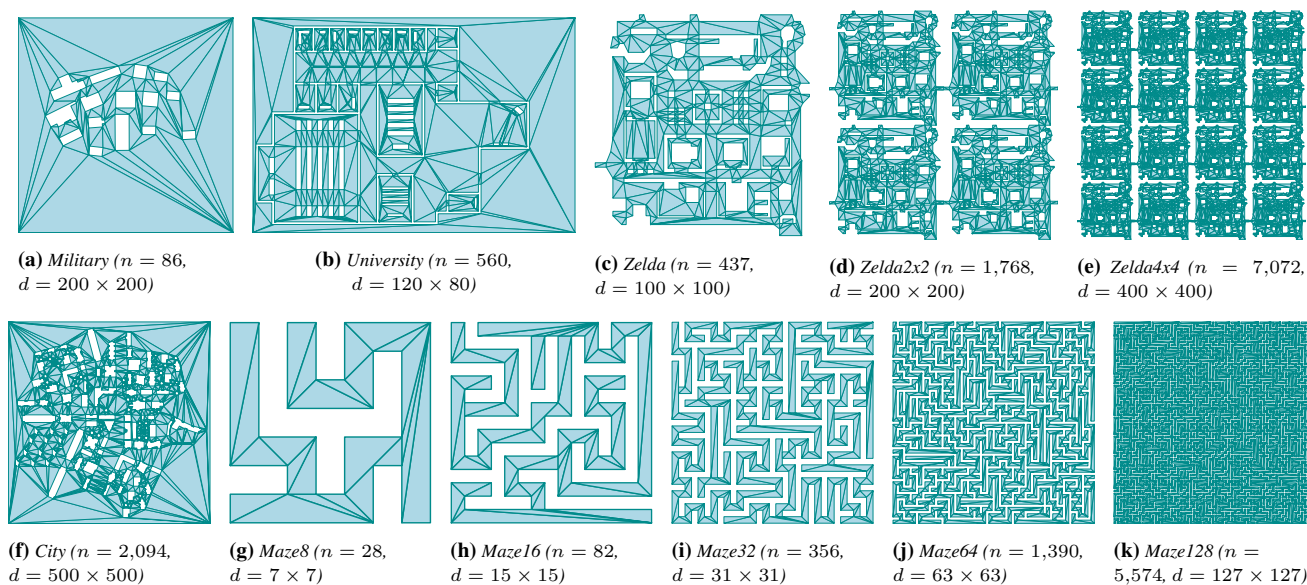


Figure 3: Top views of the 2D environments used in our experiments. The number of polygon vertices n and the physical dimensions d (in meters) are shown in brackets.

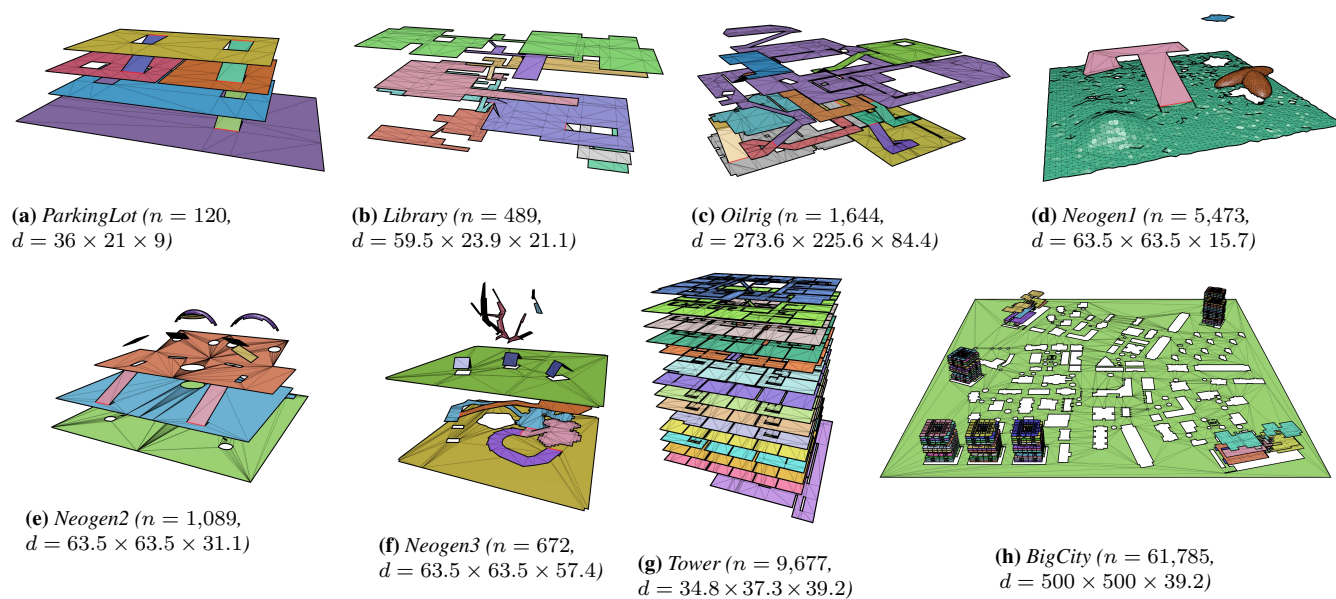


Figure 4: Renders of the multi-layered environments used in our experiments. Each layer of an environment is shown in a different color. Connections between layers are shown in red. The number of polygon vertices n and the physical dimensions d (width \times depth \times height, in meters) are shown in brackets.

As expected, our grid implementation always gave the largest graph, except in some of the mazes. This confirms that grids are usually inefficient representations, although we acknowledge their ease of use and their attractiveness for grid-aligned applications.

Performance. The LCT implementation was by far the fastest in all environments, although it required pre-processing that we have not included in our measurements. As expected, exact methods scaled better to large environments than voxel-based methods: while the LCT and ECM remained fast, the running times increased strongly for Recast and the CDG in particular. NEOGEN was usually the fastest voxel-based method. The *BigCity* environment challenged the limits of all voxel-based methods: only Recast could produce a navigation mesh using our settings. Increasing the voxel resolution caused Recast to crash as well, most likely due to memory usage. Recast can subdivide the environment into *tiles* to alleviate this, but we have excluded this option to simplify our comparison.

The differences in scalability are difficult to judge because voxel-based methods include the reconstruction of \mathcal{E}_{free} in their algorithm. Combined with the results for coverage, this indicates that obtaining \mathcal{E}_{free} without voxels is an interesting topic for future work.

7 Discussion

A limitation of our comparison lies in the current set of input environments. We have focused on examples from previous publications; these are all realistic scenarios that have been considered interesting before. Also, we have deliberately only used ‘clean’ walkable environments and not raw 3D geometry, to allow a fair comparison between exact and voxel-based methods for the same input. The goal of this paper was not to provide an exhaustive set of environments, or to expose all strengths and weaknesses of an implementation. Ultimately, it would be good to create an open database of input scenarios for researchers to use, similarly to the ones that currently exist for local character steering [Singh et al. 2009] and grid-based A* search [Sturtevant 2012].

We would also like to investigate more types of metrics. For instance, it would be useful to measure the efficiency of a navigation mesh for path planning: how much time does it take to compute paths in \mathcal{G} , and how efficiently can these be converted to geometric routes using the regions of \mathcal{R} ? Another option is to look at the *quality* of these routes: how short are they, and how well do they correspond to real-life behavior? If shortness is important, a dense grid may yield better results than a navigation mesh with a small dual graph. Ultimately, for real-world applications, we would like to quantify how well a navigation mesh captures the navigation abilities of real humans. This is a challenging direction for future work. We expect that not everything can be analyzed mechanically, and that user studies will also be required.

Finally, to simplify the comparison, we have chosen a single set of parameter settings for all methods. It would be interesting to see how different settings influence the results for each method, and how these settings can be optimized for particular metrics. For example, Oliva and Pelechano have discussed how the voxel size affects the results of Recast and NEOGEN [2013b]. We would like to combine such ideas with the quantitative metrics of our paper.

8 Conclusions and Future Work

A navigation mesh enables path planning and crowd simulation for walking characters in 2D and 3D environments. In this paper, we have performed a comparative study of multiple state-of-the-art navigation meshes. We have proposed properties by which a mesh and its construction algorithm can be classified, and metrics that

measure the quality of a mesh in practice. We have used these components to compare the Local Clearance Triangulation, the Explicit Corridor Map, the Clearance Disk Graph, NEOGEN, and a grid.

While we intend to use more environments, metrics, and settings, our results already suggest interesting properties. Voxel-based methods can be tuned to yield good coverage, but they do not always preserve connectivity, and their construction time does not seem to scale well to physically large environments. Furthermore, grids are usually not space-efficient representations of \mathcal{E}_{free} , although they may be attractive for particular applications.

The goal of this paper was not to find ‘the best’ navigation mesh, but to develop a way of comparing navigation meshes based on theoretical properties and quantitative metrics. Users may decide which properties and metrics are the most relevant for their application. We expect that this study will set a new standard for the evaluation and development of navigation meshes, and that it can help users choose an appropriate navigation mesh for particular applications.

Aside from the discussion points mentioned in Section 7, a topic for future work lies in developing *exact* algorithms that automatically extract the walkable space from arbitrary 3D input in real-time. Our experiments suggest that voxel-based approaches do not always preserve coverage and connectivity, and that they are not very scalable to large environments. However, an advantage of voxelization is that the input is automatically simplified to a certain level of precision. *Exact* filtering algorithms should yield a perfect representation of \mathcal{E}_{free} within provable time bounds, but they may be sensitive to small details or imprecisions in the input (such as gaps or overlap). In the end, it may turn out that the best solution is to combine various approaches, e.g. a filtering algorithm without voxels that is based on rounded coordinates.

Acknowledgements

We thank all research groups involved in this study for sharing their source code and environments, and for joining us in helpful discussions. Furthermore, we thank Mihai Polak and Arne Hillebrand for helping us with pre-processing 3D geometry and computing coverage using their software.

References

- ALI, S., NISHINO, K., MANOCHA, D., AND SHAH, M. 2013. *Modeling, Simulation and Visual Analysis of Crowds: A Multi-disciplinary Perspective*. Springer.
- BERSETH, G., KAPADIA, M., AND FALOUTSOS, P. 2015. AC-CLMesh: Curvature-based navigation mesh generation. In *Proc. 8th ACM SIGGRAPH Conf. on Motion in Games*, 97–102.
- BOOST, 2015. The Boost C++ library. <http://www.boost.org/>.
- BOTEA, A., MÜLLER, M., AND SCHAEFFER, J. 2004. Near optimal hierarchical path-finding. *Journal of Game Development 1*, 7–28.
- BUDDE, S. 2013. *Automatic generation of jump links in arbitrary 3D environments for navigation meshes*. Master’s thesis, Humboldt-Universität zu Berlin.
- CGAL, 2016. The Computational Geometry Algorithms Library. <http://www.cgal.org/>.
- CURTIS, S., BEST, A., AND MANOCHA, D. 2014. Menge: A modular framework for simulating crowd movement. Tech. rep., University of North Carolina at Chapel Hill.

- DEUSDADO, L., FERNANDES, A. R., AND BELO, O. 2008. Path planning for complex 3D multilevel environments. *Proc. 24th Spring Conf. on Computer Graphics*, 187–194.
- GARCÍA, F. M., KAPADIA, M., AND BADLER, N. M. 2014. GPU-based dynamic search on adaptive resolution grids. In *Proc. IEEE Int. Conf. on Robotics and Automation*, 1631–1638.
- GERAERTS, R. 2010. Planning short paths with clearance using Explicit Corridors. In *Proc. IEEE Int. Conf. on Robotics and Automation*, 1997–2004.
- HALE, D. H., AND YOUNGBLOOD, G. M. 2009. Full 3D spacial decomposition for the generation of navigation meshes. In *Proc. 5th Artificial Intelligence and Interactive Digital Entertainment Conf.*, 143–147.
- HALE, D. H., YOUNGBLOOD, G. M., AND DIXIT, P. N. 2008. Automatically-generated convex region decomposition for real-time spatial agent navigation in virtual worlds. In *Proc. 4th Artificial Intelligence and Interactive Digital Entertainment Conf.*, 173–178.
- HARABOR, D., AND GRASTIEN, A. 2011. Online graph pruning for pathfinding on grid maps. In *Proc. 52th AAAI Conf. on Artificial Intelligence*, 1114–1119.
- HART, P., NILSSON, N., AND RAPHAEL, B. 1968. A formal basis for the heuristic determination of minimum cost paths. *IEEE Transactions on Systems Science and Cybernetics* 4, 2, 100–107.
- HILLEBRAND, A. 2012. *Separating a polygonal environment into a multi-layered environment*. Master’s thesis, Utrecht University, The Netherlands.
- KALLMANN, M. 2014. Dynamic and robust Local Clearance Triangulations. *ACM Transactions on Graphics* 33, 5.
- KAPADIA, M., PELECHANO, N., ALLBECK, J., AND BADLER, N. I. 2015. *Virtual Crowds: Steps Toward Behavioral Realism*. Morgan & Claypool Publishers.
- KAVRAKI, L. E., ŠVESTKA, P., LATOMBE, J.-C., AND OVERMARS, M. H. 1996. Probabilistic roadmaps for path planning in high-dimensional configuration spaces. *IEEE Transactions on Robotics and Automation* 12, 4, 566–580.
- KOENIG, S., AND LIKHACHEV, M. 2002. D* Lite. In *Proc. AAAI Conf. of Artificial Intelligence*, 476–483.
- KOENIG, S., LIKHACHEV, M., AND FURCY, D. 2004. Lifelong Planning A*. *Artificial Intelligence* 155, 1-2, 93–146.
- LAMARCHE, F. 2009. TopoPlan: a topological path planner for real time human navigation under floor and ceiling constraints. *Computer Graphics Forum* 28, 2, 649–658.
- LAVALLE, S. M. 2006. *Planning Algorithms*. Cambridge University Press.
- LEE, W., AND LAWRENCE, R. 2013. Fast grid-based path finding for video games. In *Advances in Artificial Intelligence*, vol. 7884 of *Lecture Notes in Computer Science*. Springer, 100–111.
- LIKHACHEV, M., FERGUSON, D., GORDON, G., STENTZ, A., AND THRUN, S. 2005. Anytime Dynamic A*: An anytime, replanning algorithm. In *Proc. Int. Conf. on Automated Planning and Scheduling*, 262–271.
- LOPEZ, T., LAMARCHE, F., AND LI, T.-Y. 2012. Space-time planning in changing environments: using dynamic objects for accessibility. *Computer Animation and Virtual Worlds* 23, 87–99.
- LOZANO-PEREZ, T. 1983. Spatial planning: A configuration space approach. *IEEE Transactions on Computing* 32, 2, 108–120.
- MONONEN, M., 2014. Recast Navigation. <https://github.com/memononen/recastnavigation/>.
- OLIVA, R., AND PELECHANO, N. 2011. Automatic generation of suboptimal navmeshes. In *Proc. 4th Int. Conf. on Motion in Games*, 328–339.
- OLIVA, R., AND PELECHANO, N. 2013. A generalized exact arbitrary clearance technique for navigation meshes. In *Proc. 6th Int. Conf. on Motion in Games*, 103–110.
- OLIVA, R., AND PELECHANO, N. 2013. NEOGEN: Near optimal generator of navigation meshes for 3D multi-layered environments. *Computers & Graphics* 37, 5, 403–412.
- PELECHANO, N., ALLBECK, J. M., KAPADIA, M., AND BADLER, N. I. 2016. *Simulating Heterogeneous Crowds with Interactive Behaviors*. CRC Press.
- PETTRÉ, J., LAUMOND, J., AND THALMANN, D. 2005. A navigation graph for real-time crowd animation on multilayered and uneven terrain. In *Proc. 1st Int. Workshop on Crowd Simulation*, 81–89.
- POLAK, R. M. 2016. *Extracting walkable areas from 3D environments*. Master’s thesis, Utrecht University.
- RICKS, B. C., AND EGBERT, P. K. 2014. A whole surface approach to crowd simulation on arbitrary topologies. *IEEE Trans. Visualization and Computer Graphics* 20, 159–171.
- SINGH, S., KAPADIA, M., FALOUTSOS, P., AND REINMAN, G. 2009. An open framework for developing, evaluating, and sharing steering algorithms. In *Proc. 2nd Int. Workshop on Motion in Games*, 158–169.
- SNOOK, G. 2000. Simplified 3D movement and pathfinding using navigation meshes. In *Game Programming Gems*, M. DeLoura, Ed. Charles River Media, 288–304.
- STURTEVANT, N., AND RABIN, S. 2016. Canonical orderings on grids. In *Proc. Int. Joint Conf. on Artificial Intelligence*, 683–689.
- STURTEVANT, N. 2011. A sparse grid representation for dynamic three-dimensional worlds. In *Proc. AAAI Conf. on Artificial Intelligence and Interactive Digital Entertainment*.
- STURTEVANT, N. 2012. Benchmarks for grid-based pathfinding. *Transactions on Computational Intelligence and AI in Games* 4, 2, 144–148.
- THALMANN, D., AND MUSSE, S. R. 2013. *Crowd Simulation*, 2 ed. Springer.
- VAN TOLL, W. G., COOK IV, A. F., AND GERAERTS, R. 2011. Navigation meshes for realistic multi-layered environments. In *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, 3526–3532.
- VAN TOLL, W., JAKLIN, N., AND GERAERTS, R. 2015. Towards believable crowds: A generic multi-level framework for agent navigation. In *ASCI.OPEN*.
- TOZOUR, P. 2002. Building a near-optimal navigation mesh. In *AI Game Programming Wisdom*, S. Rabin, Ed. Charles River Media, 171–185.
- UNITY3D GAME ENGINE, 2016. <http://www.unity3d.com/>.

A Parameter Settings

Most navigation meshes are built based on various parameters. For simplicity, we use one set of parameter settings for all experiments.

Precision. For the CDG, Recast, and NEOGEN, we used voxels of 0.2 meters in all three dimensions. This is the smallest voxel size at which we could obtain results in most environments. For the CDG, we enforced a maximum resolution of 512 pixels in all dimensions. This kept the construction times manageable for environments that were too large to allow a precision of 0.2m.

For the grid baseline method, we reduced the voxel size to 1×1 m in the horizontal plane to prevent the program from taking too much time and memory. We could keep the voxel height at 0.2m.

For the CDG, we used a minimum disk radius of 0.1m (to prevent the method from generating many small disks) and a maximum disk radius of 100,000m (which is essentially infinite in our examples).

Filtering. We have created our input environments such that they are entirely walkable and no more surfaces need to be filtered out. Therefore, for all voxel-based methods, we used a maximum surface slope of 60 degrees, a character radius of 0, and a character height of 0.5m. These settings ensured that all environments could (in theory) be covered completely.

Recast offers various other parameters that need to be tweaked for each environment to get the best results. Through preliminary experiments, we obtained the following values that gave good results in most environments: Tiling: Off; Max climb: 0.5; Min region size: 0; Merged region size: 10,000; Partitioning: Watershed; Max edge length: 500; Max edge error: 1.3; Vertices per polygon: 6; Detail mesh sample distance: 6; Detail mesh max error: 1.

Other. To compute ECMs, we used the implementation based on the Boost Voronoi library [Boost 2015]. It computes a 2D Voronoi diagram in $\mathcal{O}(n \log n)$ worst-case time, and it can process multiple layers of an MLE at the same time by using parallel threads.

For NEOGEN, we used a convexity relaxation parameter of 0. Increasing this parameter leads to fewer regions (i.e. a smaller dual graph) and a higher region complexity. Thus, the ‘best’ value for this parameter depends on the application.

B Results

The values of all metrics can be found in Tables 2 to 5 on the following pages.

Environment				Navigation mesh		Coverage						Connectivity	
Name	Total area (m ²)	#CCs	#Boundaries		Free space covered		Incorrect area		Overlap		#CCs	#Boundaries	
					Absolute	Relative	Absolute	Relative	Absolute	Relative			
Military	36,876.82	1	16	CDG	36,198.31	0.98	7.90	0.00	15,613.55	0.30	3	176	
				Recast	36,629.42	0.99	7.27	0.00	0.01	0.00	1	16	
				NEOGEN	36,876.82	1.00	0.00	0.00	0.00	0.00	1	16	
				Grid	36,736.63	1.00	107.37	0.00	0.00	0.00	1	17	
University	8,370.68	1	82	CDG	7,821.35	0.93	11.42	0.00	3,387.29	0.30	274	214	
				Recast	8,106.06	0.97	51.38	0.01	0.00	0.00	1	81	
				NEOGEN	8,370.68	1.00	0.00	0.00	0.00	0.00	1	82	
				Grid	8,335.13	1.00	27.87	0.00	0.00	0.00	1	45	
Zelda	5,642.25	1	57	CDG	5,262.52	0.93	10.40	0.00	2,311.22	0.31	609	211	
				Recast	5,457.23	0.97	35.56	0.01	0.00	0.00	1	57	
				NEOGEN	5,642.24	1.00	0.00	0.00	0.00	0.00	1	57	
				Grid	5,545.39	0.98	135.61	0.02	0.00	0.00	1	59	
Zelda2x2	22,632.42	1	226	CDG	19,364.63	0.86	85.68	0.00	6,790.28	0.26	9	594	
				Recast	21,900.31	0.97	122.55	0.01	0.00	0.00	1	226	
				NEOGEN	22,632.42	1.00	0.00	0.00	0.00	0.00	1	226	
				Grid	22,178.87	0.98	484.12	0.02	0.01	0.00	1	227	
Zelda4x4	90,529.70	1	906	CDG	65,922.95	0.73	842.89	0.01	16,245.82	0.20	244	1,536	
				Recast	87,652.50	0.97	542.35	0.01	0.00	0.00	1	906	
				NEOGEN	90,529.69	1.00	0.01	0.00	0.00	0.00	1	906	
				Grid	88,679.73	0.98	1,911.23	0.02	0.03	0.00	1	921	
City	207,518.40	1	181	CDG	197,001.40	0.95	302.53	0.00	82,909.38	0.30	204	367	
				Recast	206,304.50	0.99	103.55	0.00	0.40	0.00	3	183	
				NEOGEN	207,518.40	1.00	0.04	0.00	0.00	0.00	4	185	
				Grid	206,050.10	0.99	1,524.91	0.01	0.00	0.00	2	203	
Maze8	31.00	1	1	CDG	22.10	0.71	0.35	0.01	6.37	0.22	11	4	
				Recast	23.23	0.75	0.57	0.02	0.00	0.00	1	1	
				NEOGEN	29.00	0.94	0.00	0.00	0.00	0.00	2	2	
				Grid	31.00	1.00	0.00	0.00	0.00	0.00	1	1	
Maze16	127.00	1	1	CDG	90.96	0.72	1.77	0.01	35.08	0.28	29	39	
				Recast	100.47	0.79	2.15	0.02	0.00	0.00	1	1	
				NEOGEN	125.00	0.98	0.00	0.00	0.00	0.00	3	3	
				Grid	126.00	0.99	0.00	0.00	0.00	0.00	2	2	
Maze32	511.00	1	1	CDG	352.45	0.69	6.36	0.01	74.51	0.17	92	135	
				Recast	418.44	0.82	10.61	0.03	0.00	0.00	1	1	
				NEOGEN	507.00	0.99	0.00	0.00	0.00	0.00	11	11	
				Grid	511.00	1.00	0.00	0.00	0.00	0.00	1	1	
Maze64	2,047.00	1	1	CDG	1,445.65	0.71	25.04	0.01	503.57	0.26	618	374	
				Recast	1,415.99	0.69	28.25	0.02	0.00	0.00	30	11	
				NEOGEN	2,003.00	0.98	0.00	0.00	0.00	0.00	45	45	
				Grid	2,045.00	1.00	0.00	0.00	0.00	0.00	4	1	
Maze128	8,191.00	1	1	CDG	5,066.94	0.62	109.58	0.02	896.21	0.15	1,188	3,548	
				Recast	7,464.19	0.91	6,428.45	0.52	4,807.87	0.39	123	89	
				NEOGEN	7,951.00	0.97	0.00	0.00	0.00	0.00	164	164	
				Grid	8,184.00	1.00	0.00	0.00	0.00	0.00	8	8	

Table 2: Results for the coverage and connectivity metrics in the 2D environments. We report these only for CDG, Recast, and the grid because the other methods do not attempt to discover $\mathcal{E}_{\text{free}}$ themselves. ‘#CCs’ is an abbreviation of ‘number of connected components’. The descriptions of all metrics can be found in Section 5.

Environment				Navigation mesh	Coverage						Connectivity		
Name	Total area (m ²)	#CCs	#Boundaries		Free space covered		Incorrect area		Overlap		#CCs	#Boundaries	
					Absolute	Relative	Absolute	Relative	Absolute	Relative			
ParkingLot	1,921.50	1	9	CDG	1,819.68	0.95	0.07	0.00	675.84	0.27	178	81	
				Recast	1,849.52	0.96	1.24	0.00	0.00	0.00	2	10	
				NEOGEN	1,920.88	1.00	0.63	0.00	0.00	0.00	1	9	
				Grid	1,861.50	0.97	4.51	0.00	0.00	0.00	4	11	
Library	3,154.06	1	4	CDG	2,951.06	0.94	9.62	0.00	1,108.37	0.27	256	137	
				Recast	3,046.73	0.97	9.14	0.00	0.41	0.00	1	4	
				NEOGEN	3,132.65	0.99	21.41	0.01	0.00	0.00	1	4	
				Grid	2,947.75	0.93	243.26	0.08	0.00	0.00	1	4	
Oilrig	75,746.52	1	26	CDG	70,551.81	0.93	3,688.22	0.04	27,946.73	0.28	578	468	
				Recast	74,933.66	0.99	124.29	0.00	3.30	0.00	1	29	
				NEOGEN	73,475.38	0.97	2,271.13	0.03	0.00	0.00	1	26	
				Grid	74,183.09	0.98	1,714.92	0.02	0.00	0.00	1	30	
Neogen1	4,748.44	3	12	CDG	4,341.39	0.91	550.79	0.09	1,731.02	0.29	43	176	
				Recast	4,659.20	0.98	16.38	0.00	7.65	0.00	3	17	
				NEOGEN	4,519.90	0.95	203.78	0.05	0.00	0.00	3	15	
				Grid	4,449.82	0.94	70.04	0.02	0.00	0.00	6	46	
Neogen2	9,600.56	11	33	CDG	9,237.81	0.96	214.07	0.02	4,120.91	0.31	175	159	
				Recast	9,431.68	0.98	14.29	0.00	5.07	0.00	10	27	
				NEOGEN	9,371.87	0.98	20.71	0.00	0.00	0.00	6	27	
				Grid	9,334.44	0.97	237.17	0.03	0.00	0.00	10	36	
Neogen3	9,642.51	10	20	CDG	9,349.64	0.97	35.66	0.00	3,884.67	0.29	58	347	
				Recast	9,440.83	0.98	8.82	0.00	0.57	0.00	15	22	
				NEOGEN	9,527.26	0.99	15.06	0.00	0.00	0.00	10	18	
				Grid	9,297.01	0.96	133.04	0.01	0.00	0.00	28	27	
Tower	12,093.88	1	19	CDG	10,615.44	0.88	141.53	0.01	3,746.06	0.26	2,052	1,105	
				Recast	11,672.40	0.97	425.70	0.04	5.26	0.00	1	256	
				NEOGEN	-	-	-	-	-	-	-	-	-
				Grid	11,121.91	0.92	945.09	0.09	0.00	0.00	1	225	
BigCity	280,897.00	1	301	CDG	-	-	-	-	-	-	-	-	
				Recast	277,326.20	0.99	2,820.01	0.01	15.14	0.00	8	1,796	
				NEOGEN	-	-	-	-	-	-	-	-	-
				Grid	-	-	-	-	-	-	-	-	-

Table 3: Results for the coverage and connectivity metrics in the multi-layered environments. An empty row indicates that the navigation mesh could not be computed for the corresponding algorithm and environment.

Environment	Navigation mesh	Complexity			Performance						
		$ V $	$ E $	$ \mathcal{R} $	Region complexity			Construction time (ms)		Memory usage (MB)	
					Average	SD	Total	Average	SD	Average	SD
Military	LCT	120	134	120	9.00	0.00	1,080	0.80	0.40	0.88	0.00
	ECM	58	72	214	14.83	2.16	3,174	7.23	0.01	29.99	0.02
	CDG	1,168	2,078	1,168	4.00	0.00	4,672	44,256.03	85.58	79.74	0.01
	Recast	101	115	101	11.55	2.88	1,167	1,055.51	9.84	24.76	0.00
	NEOGEN	52	66	52	15.40	3.58	801	15.90	0.70	66.53	0.32
	Grid	36,844	72,755	36,844	12.00	0.00	442,128	205.39	2.24	77.88	0.38
University	LCT	732	812	732	9.00	0.00	6,588	5.90	0.30	1.16	0.01
	ECM	329	409	1,134	15.04	2.29	17,055	31.14	0.06	33.02	0.02
	CDG	3,309	4,369	3,309	4.00	0.00	13,236	27,866.70	17.94	68.72	0.00
	Recast	402	481	402	11.88	2.78	4,776	209.72	1.79	9.08	0.00
	NEOGEN	261	341	261	16.80	8.13	4,386	94.90	2.51	86.46	3.07
	Grid	8,363	15,460	8,363	12.00	0.00	100,356	106.69	1.23	54.05	0.31
Zelda	LCT	554	608	554	9.00	0.00	4,986	4.60	0.49	1.07	0.00
	ECM	289	344	895	14.93	2.30	13,359	23.53	0.07	32.34	0.01
	CDG	3,579	4,233	3,579	4.00	0.00	14,316	28,802.21	26.55	68.47	0.00
	Recast	321	376	321	12.12	2.77	3,891	159.42	0.80	8.14	0.00
	NEOGEN	205	260	205	16.04	6.31	3,288	59.10	2.39	80.67	0.41
	Grid	5,681	10,193	5,681	12.00	0.00	68,172	92.35	0.85	51.69	0.28
Zelda2x2	LCT	2,248	2,472	2,248	9.00	0.00	20,232	19.00	0.77	1.79	0.00
	ECM	1,148	1,372	3,602	14.92	2.30	53,754	107.45	0.47	41.89	0.03
	CDG	5,636	8,850	5,636	4.00	0.00	22,544	31,915.01	57.39	75.47	0.03
	Recast	1,281	1,505	1,281	12.16	2.84	15,573	736.57	3.04	24.00	0.17
	NEOGEN	820	1,044	820	16.15	6.35	13,245	246.60	8.49	125.54	1.80
	Grid	22,663	40,658	22,663	12.00	0.00	271,956	201.52	1.97	73.82	0.24
Zelda4x4	LCT	9,007	9,911	9,007	9.00	0.00	81,063	96.21	1.08	4.73	0.00
	ECM	4,580	5,484	14,436	14.92	2.30	215,424	438.66	2.81	78.67	0.31
	CDG	11,996	16,564	11,996	4.00	0.00	47,984	39,629.89	40.84	105.19	0.01
	Recast	5,105	6,009	5,105	12.17	2.84	62,148	3,275.23	8.29	89.36	0.68
	NEOGEN	3,289	4,193	3,289	16.15	6.36	53,103	1,040.90	19.45	253.26	2.49
	Grid	90,591	162,537	90,591	12.00	0.00	1,087,092	796.62	13.17	153.34	0.30
City	LCT	2,553	2,732	2,553	9.00	0.00	22,977	33.40	0.49	2.11	0.00
	ECM	1,442	1,621	4,679	14.42	2.16	67,491	162.69	0.41	45.54	0.02
	CDG	3,451	5,278	3,451	4.00	0.00	13,804	44,140.58	380.99	83.88	0.02
	Recast	1,527	1,706	1,527	11.90	3.08	18,168	11,833.70	9.93	135.15	0.00
	NEOGEN	1,164	1,343	1,164	13.61	5.17	15,846	330.70	6.18	126.86	0.79
	Grid	207,575	408,383	207,575	12.00	0.00	2,490,900	1,431.77	10.35	200.95	3.17
Maze8	LCT	26	25	26	9.00	0.00	234	0.20	0.40	0.84	0.00
	ECM	30	29	51	14.71	2.32	750	1.54	0.01	29.40	0.00
	CDG	68	59	68	4.00	0.00	272	1,105.43	9.50	32.18	0.01
	Recast	14	13	14	11.57	1.05	162	1.00	0.00	2.12	0.00
	NEOGEN	12	10	12	14.25	3.70	171	3.20	0.60	63.72	2.02
	Grid	31	30	31	12.00	0.00	372	62.39	1.44	41.17	0.17
Maze16	LCT	81	80	81	9.00	0.00	729	0.70	0.46	0.88	0.00
	ECM	84	83	156	14.85	2.25	2,316	4.28	0.02	29.84	0.01
	CDG	310	319	310	4.00	0.00	1,240	2,508.26	9.16	38.90	0.01
	Recast	42	41	42	11.71	1.44	492	7.30	0.46	2.25	0.00
	NEOGEN	39	36	39	14.62	3.05	570	9.20	0.75	63.98	2.66
	Grid	126	124	126	12.00	0.00	1,512	63.77	1.32	42.09	0.22
Maze32	LCT	363	362	363	9.00	0.00	3,267	3.40	0.49	1.02	0.00
	ECM	358	357	686	14.89	2.23	10,212	18.05	0.03	31.87	0.00
	CDG	1,084	1,138	1,084	4.00	0.00	4,336	5,458.99	5.88	45.90	0.00
	Recast	184	183	184	11.77	1.34	2,166	391.64	0.66	2.83	0.00
	NEOGEN	168	157	168	14.75	2.86	2,478	38.90	0.83	77.59	1.89
	Grid	511	510	511	12.00	0.00	6,132	72.58	1.51	46.24	0.11
Maze64	LCT	1,422	1,421	1,422	9.00	0.00	12,798	12.50	0.67	1.64	0.00
	ECM	1,392	1,391	2,672	14.88	2.24	39,756	70.34	0.25	40.16	0.01
	CDG	4,805	4,488	4,805	4.00	0.00	19,220	13,448.17	19.51	55.80	0.01
	Recast	602	572	602	11.63	1.45	6,999	11,834.30	11.52	4.64	0.00
	NEOGEN	636	591	636	14.74	2.91	9,372	161.30	6.08	109.03	2.69
	Grid	2,045	2,041	2,045	12.00	0.00	24,540	94.05	0.62	54.51	0.31
Maze128	LCT	5,674	5,673	5,674	9.00	0.00	51,066	62.21	0.75	3.99	0.00
	ECM	5,567	5,566	10,710	14.88	2.24	159,336	306.97	0.63	71.75	0.23
	CDG	25,135	44,910	25,135	4.00	0.00	100,540	56,150.68	93.81	92.84	0.01
	Recast	2,590	2,480	2,590	11.73	2.01	30,393	52,746.20	19.62	11.19	0.00
	NEOGEN	2,574	2,410	2,574	14.73	2.98	37,914	654.80	12.03	199.44	1.89
	Grid	8,184	8,176	8,184	12.00	0.00	98,208	250.38	2.70	81.05	0.26

Table 4: Results for the complexity and performance metrics in the 2D environments. The descriptions of all metrics can be found in Section 5.

Environment	Navigation mesh	Complexity						Performance			
		$ V $	$ E $	$ \mathcal{R} $	Region complexity		Total	Construction time (ms)		Memory usage (MB)	
					Average	SD			Average	SD	Average
ParkingLot	ECM	61	68	108	15.08	2.33	1,629	7.16	1.53	31.45	0.20
	CDG	779	873	779	4.00	0.00	3,116	6,686.30	18.02	53.33	0.01
	Recast	60	66	60	11.35	1.98	681	31.60	0.66	3.53	0.00
	NEOGEN	24	31	24	20.75	4.58	498	53.00	0.00	63.07	0.01
	Grid	1,866	3,493	1,866	12.00	0.00	22,392	88.84	3.13	49.18	0.29
Library	ECM	216	218	377	14.51	2.35	5,469	13.07	2.27	33.48	0.38
	CDG	1,758	2,173	1,758	4.00	0.00	7,032	11,374.29	31.39	59.02	0.01
	Recast	111	113	111	12.19	2.65	1,353	62.41	0.80	4.72	0.00
	NEOGEN	74	76	74	20.64	8.85	1,527	56.10	4.35	71.39	0.05
	Grid	3,191	5,801	3,191	12.00	0.00	38,292	211.99	22.58	54.12	0.29
Oilrig	ECM	603	629	1,283	14.72	2.23	18,891	55.67	10.68	38.52	0.76
	CDG	4,858	6,316	4,858	4.00	0.00	19,432	45,585.04	61.95	100.93	0.01
	Recast	324	353	324	12.35	3.08	4,002	2,280.03	6.43	49.79	0.23
	NEOGEN	253	279	253	23.44	12.99	5,931	219.80	6.18	107.94	0.49
	Grid	75,898	147,409	75,898	12.00	0.00	910,776	2,827.53	51.75	707.46	10.02
Neogen1	ECM	438	444	1,148	14.67	2.41	16,842	158.49	25.01	45.62	0.00
	CDG	1,017	1,651	1,017	4.00	0.00	4,068	19,911.95	24.38	101.37	0.03
	Recast	103	114	103	11.80	3.27	1,215	281.58	0.92	8.89	0.00
	NEOGEN	193	202	193	23.21	28.35	4,479	1,924.90	18.25	195.75	0.02
	Grid	4,519	8,494	4,519	12.00	0.00	54,228	1,498.18	43.75	192.09	0.12
Neogen2	ECM	390	403	1,240	14.86	2.32	18,426	43.40	6.22	37.58	0.46
	CDG	2,170	2,911	2,170	4.00	0.00	8,680	32,356.20	50.54	95.81	0.01
	Recast	198	213	198	11.74	2.97	2,325	208.62	1.28	8.58	0.00
	NEOGEN	295	312	295	15.31	7.11	4,515	245.60	4.20	93.33	0.45
	Grid	9,571	18,374	9,571	12.00	0.00	114,852	498.23	23.05	80.71	6.43
Neogen3	ECM	439	439	984	14.54	2.35	14,304	24.52	2.67	36.07	0.34
	CDG	2,070	3,319	2,070	4.00	0.00	8,280	32,919.54	80.55	82.33	0.01
	Recast	275	276	275	11.66	3.08	3,207	213.52	1.63	8.40	0.00
	NEOGEN	218	218	218	15.26	8.77	3,327	100.90	3.42	76.77	0.50
	Grid	9,430	17,962	9,430	12.00	0.00	113,160	523.35	54.51	109.36	1.40
Tower	ECM	4,988	5,019	9,416	14.37	2.38	135,300	139.45	17.94	72.63	0.58
	CDG	12,643	14,281	12,643	4.00	0.00	50,572	51,898.69	129.44	169.52	0.03
	Recast	1,155	1,437	1,155	12.47	3.02	14,400	538.25	3.79	16.22	0.05
	NEOGEN	-	-	-	-	-	-	-	-	-	-
	Grid	12,067	21,755	12,067	12.00	0.00	144,804	3,699.97	60.07	192.12	1.08
BigCity	ECM	32,167	32,554	69,176	14.41	2.37	997,053	756.03	5.45	299.44	0.90
	CDG	-	-	-	-	-	-	-	-	-	-
	Recast	9,394	11,345	9,394	12.27	3.11	115,260	16,353.00	28.15	216.66	0.20
	NEOGEN	-	-	-	-	-	-	-	-	-	-
	Grid	-	-	-	-	-	-	-	-	-	-

Table 5: Results for the complexity and performance metrics in the multi-layered environments. An empty row indicates that the navigation mesh could not be computed for the corresponding algorithm and environment.