

Received January 15, 2022, accepted April 16, 2022, date of publication April 20, 2022, date of current version April 28, 2022.

Digital Object Identifier 10.1109/ACCESS.2022.3168832

Procedural Generation of 3D Maps With Snappable Meshes

RAFAEL C. E SILVA¹, NUNO FACHADA², DIOGO DE ANDRADE², AND NÉLIO CÓDICES³

¹ECATI, Lusófona University, Campo Grande, 1749-024 Lisboa, Portugal

²COPELABS, Lusófona University, Campo Grande, 1749-024 Lisboa, Portugal

³Instituto Superior Técnico, Universidade de Lisboa, 1049-001 Lisboa, Portugal

Corresponding author: Nuno Fachada (nuno.fachada@ulusofona.pt)

This work was supported by the Fundação para a Ciência e a Tecnologia under Grant UIDB/04111/2020 (COPELABS).

ABSTRACT In this paper we present a technique for procedurally generating 3D maps using a set of premade meshes which snap together based on designer-specified visual constraints. The proposed approach avoids size and layout limitations, offering the designer control over the look and feel of the generated maps, as well as immediate feedback on a given map's navigability. A prototype implementation of the method, developed in the Unity game engine, is discussed, and a number of case studies are analyzed. These include a multiplayer game where the method was used, together with a number of illustrative examples which highlight various parameterizations and piece selection methods. The technique can be used as a designer-centric map composition method and/or as a prototyping system in 3D level design, opening the door for quality map and level creation in a fraction of the time of a fully human-based approach.

INDEX TERMS 3D maps, computer games, designer-centric methods, layout, procedural content generation (PCG).

I. INTRODUCTION

In this paper we describe a method for procedurally generating 3D maps using visual constraints. The approach, termed *snappable meshes*, consists of a system of connectors with pins and colors—similar in concept to a jigsaw puzzle—which constrains how any two map pieces (i.e., meshes) can snap together. Through the visual design and specification of these connection constraints, and an easy-to-follow and fully explainable [1] generation procedure, the method is accessible to game designers and/or other non-experts in procedural content generation (PCG), artificial intelligence (AI) or programming.

While the maps are procedurally generated, the technique grants the game designer considerable influence over the final result. More specifically, the designer has complete control on the modeling of individual meshes—including the placement and configuration of associated connectors—and can achieve substantial customization on the generated maps through a small number of intuitive parameters. At the same time, the technique avoids size and layout limitations, common in grid-based approaches [2]. Maps are generated almost immediately and can be quickly validated for navigability. With a

focus on fast iteration and respect for existing development workflows [3], the snappable meshes PCG technique can be used in practical industry settings, while following academic best practices such as transparency and reproducibility [4]. Further, the technique can be easily tweaked and extended by programmers, since one of its core components—the *selection method*, discussed in detail in Section III—is fully swappable, and implementing new ones is relatively simple. Four of these *selection methods* are described in this paper and included in the provided prototype implementation.

The snappable meshes PCG technique was initially developed for a multiplayer combat game, created as a semester project at Lusófona University's Bachelor in Videogames—an industry-focused, interdisciplinary game development degree [5]. The decision of using PCG in this particular game project was made to promote its replayability, requiring players to adapt to a new map on every match, keeping the experience from turning stale once the combat loop is mastered. Given the technique's capability of creating general 3D maps (i.e., not specific to action games), its core ideas were previously presented in a conference [6]. The present paper extends that publication in several ways, making the following contributions:

- The technique is presented in a carefully formalized, fully reproducible and implementation-independent

The associate editor coordinating the review of this manuscript and approving it for publication was Songwen Pei.

fashion. Furthermore, it is comprehensively classified and contextualized within the PCG state of the art, highlighting the advantages it brings with respect to explainability [1] and potential integration with industry workflows [3].

- A new optional connection constraint is introduced in the form of a color compatibility layer, specified through a color matrix. Additionally, two navigability-focused map validation metrics are proposed and analyzed.
- A reference and standalone implementation of the technique, developed in the Unity game engine [7], is presented and examined. This implementation is made available as free and open-source software, and the source code is thoroughly annotated. Further, it includes concrete solutions for issues considered separate from the technique itself, such as mesh overlap detection (i.e., collision avoidance) and specific issues related with validating map navigability.
- Using the prototype implementation, an in-depth analysis of the technique's generation and validation capabilities is undertaken, namely on how the various parameters and *selection methods* influence the created maps, as well as the generation and validation times.
- A thorough discussion is presented on several aspects of the proposed technique, related to its designer-centric approach, applicability (i.e., in which scenarios could the method be more useful), limitations, as well as alternative uses and possible improvements.

The paper is organized as follows. In Section II, we review related work concerning the use of PCG for map design in computer games. In Section III, we describe the snappable meshes PCG technique, namely the components that make up the system, and classify the method according to a well-known PCG taxonomy. A reference implementation of the method, developed in the Unity game engine, is presented in Section IV. Several case studies are examined in Section V, including one where the method was used for generating arenas for the aforementioned multiplayer combat game. A discussion of the method, possible uses, limitations and potential improvements follows in Section VI. Section VII closes the paper, offering some conclusions.

II. BACKGROUND

The use of PCG for creating maps and levels in computer games began in the 1970's with games such as *Beneath Apple Manor* and *Rogue* [8], and is unarguably the most common use case of PCG in games [9]. Academic interest in this area has been growing, with a number of important developments occurring in the last 15 years [10], [11].

PCG techniques are often fully autonomous, in the sense that the user simply performs some initial configuration, such as defining parameters and/or output constraints, before content is generated. Recently, mixed-initiative content creation (MICC) techniques, in which a combination of human input and computer-assisted PCG are used, have gained

traction in game development in general and map design in particular [9], [12]. However, the line that separates autonomous PCG and MICC-PCG is ill-defined. For example, Yannakakis & Togelius [9] claim that “the PCG process is autonomous when the initiative of the designer is limited to algorithmic parameterizations before the generation starts”. Then, how to classify a PCG technique that works according to this statement, but produces results which mainly depend on designer-provided building blocks, such as the one presented in this paper? Although we do not provide a definitive answer, this question guides the type of related works discussed in this section. More specifically, we will focus on PCG methods and tools for creating maps and levels in which the human designer has considerable influence on the look and feel of the generated output, irrespective of whether the PCG algorithm tends to be more autonomous (i.e., the designer specifies the parameters and possibly provides the building blocks) or works in a more MICC fashion, where human and computer iteratively collaborate on the design process.

Tanagra is one such MICC tool for 2D platformers [13], allowing human designers to partially specify a level's geometry and pacing, leaving it up to the computer to fill in the gaps. Using constraint programming, Tanagra guarantees that the generated levels are playable when human-defined constraints are valid.

Occupancy-regulated extension (ORE) [14] is a 2D geometry assembly algorithm that, similar to the algorithm presented in this paper, (i) creates maps from premade pieces (or chunks), and, (ii) delegates context and chunk selection to separate subroutines. However, contrary to the snappable meshes technique, ORE utilizes the potential positions a player might occupy in a chunk during play to determine which other chunks are adequate for combination, thus requiring pieces of geometry annotated with gameplay information.

PCG can be used together with pre-established level design patterns to generate desirable levels. Two proposals have employed a MICC approach, where the contribution from the level designers is used in the optimization process of the generators' evolutionary algorithm [15], [16]. This way, their accumulated knowledge of level design is fed into the generators, refining not only level layouts but also agent and item placement on each iteration. This is accomplished by identifying and then reinforcing the level design patterns in the heuristics, leading to both playable and enjoyable experiences.

Sentient Sketchbook [17] is a MICC level design tool for strategy games, where designers sketch maps while being presented with similar, more detailed alternatives. These recommendations are presented in real time using evolutionary search and employing gameplay metrics—such as balance, exploration, resources or navigability—as fitness dimensions.

A number of 3D map and level generation approaches have also been proposed. SuSketch [18], aimed at generating

first-person shooter levels, is one such case. Like Sentient Sketchbook, the user provides the initial designs and the tool offers alternative layouts and feedback on a number of gameplay metrics in a full MICC fashion. Additionally, SuSketch also provides gameplay predictions which take into account different character classes.

FPSEvolver is a 3D MICC generation proposal in which the human contribution for level design comes from the players themselves. It consists of a multiplayer shooter featuring a novel grid-based interactive evolution approach for generating maps according to the players' preferences [19]. Players vote on a selection of evolving scenarios, with the goal of generating levels in accordance with what they consider to be a good map.

Oblige is a 3D level generator for the DOOM family of games [20]. It allows the level designer to set a number of parameters, such as level size and approximate quantities of each type of monster, power-up and level section (e.g. outdoors, caves, hallways, etc). Levels are created using shape grammars on a grid-based layout, and are limited to a single floor—an inherent limitation of the DOOM family of games.

Butler *et al.* [21] explored MICC-based PCG in the context of macro level design, where the focus is the entire player experience and not just single levels. Their solution relies on a set of authoring tools where level designers can define progression constraints that lead to the generation of a progression plan. The solution then generates and/or validates the individual level designs against the progression plan.

Regarding recent commercial applications, map and/or level PCG has been employed for level creation in a number of games, namely 2D roguelikes such as *Spelunky Classic* [22], *The Binding of Isaac: Rebirth* [23] and *Enter the Gungeon* [24]. *Spelunky* uses premade room templates to fill out a grid. Rooms with specific characteristics such as top entrances and bottom pits are considered when generating levels in order to create a valid path for the player to traverse towards the end [25], [26]. In *The Binding of Isaac: Rebirth*, maps are created by connecting several rooms together, fit into a grid [27], [28]. However, each room may take more than one grid space, and each grid space it occupies can be connected to other rooms on adjacent grid spaces. This allows for large rooms to connect to small rooms and vice-versa. *Enter the Gungeon* does not connect its rooms directly to one another, employing a more complex algorithm to obtain the desired layout. It uses nodes and connectors, placing different premade rooms as those nodes and afterwards creating corridors to join the rooms for the final map layout [29].

Generate Worlds [30] is a commercial tool for creating 3D maps from user-provided voxel tiles or blocks. It uses a very simple premise: two blocks are only adjacently placed if they have the same color in all the places where they touch, in a similar fashion to Wang tiles [31], often used for image and texture generation [32], [33]. If the user carefully designs these blocks, the method is able to generate varied and interesting content, from dungeons to landscapes. Generate Worlds builds on a previously proposed open-source

tool [34], and essentially solves the constraint satisfaction problem of correctly tiling blocks according to their colors. The Wave Function Collapse (WFC) family of algorithms, initially proposed by Maxim Gumin [35], takes these ideas further, requiring only one input example (e.g., an image) and a rule to decompose it into blocks (e.g., tiles). For the purpose of image, map and level generation, WFC works essentially on a grid, although it has been shown to handle graph representations for more specific generation tasks, such as node placement in navigation meshes [36]. Nonetheless, gridless map generation has yet to be demonstrated. WFC has been gaining momentum among game developers, and was recently analyzed and formalized by Karth and Smith [37], which also look into several of the technique's uses and extensions since its inception. A recent tool, Tessera [38], addresses several common issues with the original WFC implementation, allowing users to experiment with the technique and several of its extensions within the Unity game engine.

The map generation technique proposed in this paper uses some of the ideas present in the works discussed thus far, but follows its own distinct approach. For example, contrary to the work of Butler *et al.* [21], the proposed method does not deal with gameplay or inter-level balance, focusing instead on the layout of individual maps. Furthermore, unlike evolutionary-based designer feedback approaches [15]–[18], the presented PCG technique relies on the designer to carefully model individual map pieces, since one of the goals is to generate multiple quality maps, rather than refining one design. Like Sentient Sketchbook [17] and SuSketch [18], the snappable meshes technique uses path finding for map evaluation; contrary to these tools, however, the technique is less bound to specific game genres (at the cost of less detailed maps) and offers a fully explainable generation process. Snappable meshes follows a constraint-based constructive approach, similar to Generate Worlds [30] and WFC [35], but is not bound to a grid, allowing free-form map generation. Indeed, it is considerably more flexible than Generate Worlds, as it does not require pixel-perfect colored block interfaces, and is simpler than WFC, which several users found difficult to grasp and refactor [37]. Further, unlike many of the works discussed here, the proposed technique is not a specific tool designed for certain use cases, but a generic and easy-to-understand procedure which can be integrated in various designer tools and workflows.

III. THE SNAPPABLE MESHES PCG TECHNIQUE

The snappable meshes PCG technique is presented in this section in a game engine-independent fashion. The technique, summarized in Fig. 1, requires a set of premade *map pieces* to generate maps. These pieces can be manually created by the game designer and should contain one or more *connectors*. Connectors, which have a *color* and one or more *pins*, are placed on the mesh in locations where pieces can snap together. Map pieces, connectors and pins/colors are detailed in Subsection III-A. The set of human-designed map

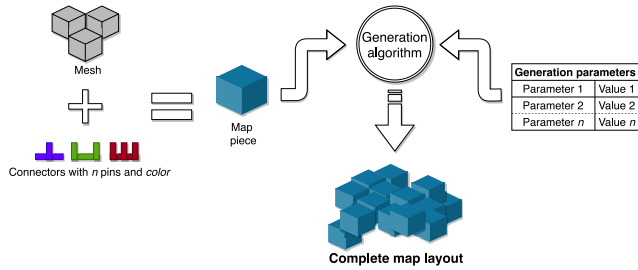


FIGURE 1. Summary of the snappable meshes map generation process.




FIGURE 2. Map pieces and connectors. (a) An example of a map piece with four visible connectors; (b) an isolated 4-pin connector with the straight line representing its heading; (c) two map pieces snapped together at their compatible connectors.

pieces, together with a number of generation parameters—also defined by the game designer—are fed to the *generation algorithm*. The algorithm is then able to create a playable map, as described in Subsection III-B. Certain aspects of the algorithm depend on the chosen *selection method*. A number of selection methods, and their influence on the generation algorithm, are presented in Subsection III-C. Two metrics for validating the generated maps are discussed in Subsection III-D. Finally, in Subsection III-E, the snappable meshes PCG technique is classified according to the PCG taxonomy proposed by Togelius *et al.* [11].

A. MAP PIECES AND CONNECTORS

Map pieces are sets of geometry that include one or more connectors. An example of a map piece with visible connectors is shown in Fig. 2a. For a map piece to be usable by the generation algorithm it needs to have at least one associated connector, since these determine where two pieces will be joined, i.e., snapped together. Map pieces are in effect content parameters for the generation algorithm. As shown in Table 1, the generation algorithm accepts a list of pieces for generating the map and an optional list of pieces which can be used as the *starting piece*, i.e., the first piece to be placed on the map.

Compatibility between connectors from different pieces is defined by two parameters set by the designer: *pin count* and *color*. Connectors have a heading, represented in Fig. 2b by a straight line. Pieces with matching connectors can be snapped together, as shown in Fig. 2c. When two connectors are matched, a copy of the *tentative piece*—the piece being evaluated for placement on the map—is positioned so its connector and the connector of the *guide piece*—a piece previously added to the map—are facing each other. The connection is made such that the connector pair completely fits/overlaps, unless a predefined connector distance is set. As will be discussed next, this connection can be optionally invalidated if the tentative piece’s geometry overlaps with

TABLE 1. Content parameters for the generation algorithm.

Parameter	Description
piecesList	List of map pieces to be used by the algorithm to create the map.
useStarter	Boolean. If true, the starting piece will be drawn from the starterList instead of the piecesList.
starterList	List of map pieces to be used as the starting piece if useStarter is true.

TABLE 2. Connection parameters for the map pieces.

Parameter	Description
matchingRules	Selection. Selects whether to use <i>none</i> , <i>pins</i> , <i>colors</i> or <i>both</i> as connector matching rules.
pinTolerance	Integer. The maximum allowed difference between pin counts in two connectors to allow them to be paired up.
colorMatrix	Boolean matrix. Valid color combinations, one-way or two-way (i.e., symmetric).

existing map geometry. Matching rules between connectors are defined as follows:

- 1) Connectors may be compatible if their pin count is equal, or if their difference is within a tolerance level defined in the generation parameters.
- 2) Connectors may be compatible if their colors match according to a color matrix given as a generation parameter. The use of a non-symmetric color matrix allows the designer to specify one- or two-way compatibility between the connectors of the guide and tentative pieces.

These matching rules are optional. The designer can activate both rules, only one of them, or even none. If both rules are active, connector compatibility is established only if both rules are verified. If none of the rules is set, all connectors in different pieces are compatible and pieces can snap together on their connectors without restriction. Table 2 summarizes how the matching rules are given as generation parameters.

B. GENERATION ALGORITHM

The pseudo-code of the generation algorithm is presented in Algorithm 1, and the general algorithm parameters are described in Table 3. The algorithm begins by selecting the starting piece (line 1) and placing it on the map (line 2). The exact process of selecting the starting piece depends on the chosen selection method, and is detailed in the next subsection. By default, the selection method should choose the starting piece from the *piecesList*, i.e., from the list of pieces used during the generation process. However, if the *useStarter* option is enabled, the selection method will instead pick the starting piece from a separate *starterList*, as described in Table 1. In any case, when a piece is selected for placement, it is not used directly. Instead, a copy is made and it is the copy that is placed on the map. Thus, pieces in these lists act as blueprints for the pieces actually deployed during the generation process.

Before entering the main loop of the algorithm, the starting piece is selected as the guide piece (line 3), since it is the only piece currently placed on the map. When the main loop

Algorithm 1 Map Generation

```

1 startingPiece ← selMethod.SelectStartingPiece(useStarter ?
  starterList : piecesList)
2 map.InitializeWith(startingPiece)
3 guidePiece ← startingPiece
4 do
5   failCount ← 0
6   connection ← none
7   do
8     tentativePiece ← piecesList.GetRandomItem()
9     connList ←
      guidePiece.GetConnectionsWith(tentativePiece)
10    while connList is not empty and connection is none do
11      connection ← connList.GetRandomItem()
12      if checkOverlaps and map.IsOverlap (connection)
      then
13        connList.Remove (connection)
14        connection ← none
15    if connection is not none then
16      guidePiece.SnapWith(tentativePiece, connection)
17    else
18      failCount ← failCount + 1
19  while connection is none and failCount < maxFails
20  guidePiece ← selMethod.SelectGuidePiece(map)
21 while guidePiece is not none

```

begins (line 4), the `failCount` and `connection` variables are initialized to zero and `none`, respectively (lines 5–6). The former counts how many failed connection attempts have occurred between the current guide piece and tentative pieces. The latter represents a valid connection pair between the current guide and tentative pieces.

The algorithm then enters the tentative piece selection and placement loop (lines 7–19). A tentative piece is randomly selected from `piecesList` (line 8), and all possible connector pairings between the guide piece and the tentative piece are evaluated. Valid pairings are stored in a temporary list (line 9). A connector pairing is considered valid if, and only if, the following conditions are met:

- 1) Both connectors are unused, thus available for pairing.
- 2) Both connectors fulfill the criteria described in the previous subsection, in accordance with the `matchingRules` parameter (see Table 2).

The next steps—i.e., the while loop in lines 10–14—depend on whether overlap verification (i.e., the `checkOverlaps` parameter, see Table 3) is enabled or not. If this verification is disabled, the while loop finishes right after a random pairing is drawn from the temporary list and placed in the `connection` variable (line 11). However, if overlap verification is enabled, the algorithm will check if the tentative piece, when connected to the guide piece by the randomly selected pairing at a distance of `pieceDistance` (see Table 3), overlaps with existing geometry (line 12). If so, the pairing is removed from the temporary list (line 13) and the `connection` variable is set to `none` (line 14). In such case, the while loop (lines 10–14) continues until a non-overlapping

TABLE 3. General parameters for the generation algorithm.

Parameter	Description
<code>maxFails</code>	Integer. Number of tentative pieces that the algorithm tries to connect with a guide piece.
<code>pieceDistance</code>	Real value. The spacing between two joined connectors.
<code>checkOverlaps</code>	Boolean. If true, connections will be invalidated if there is an overlap between tentative piece and map geometry.

solution is found or there are no more pairings left in the temporary list.

If the previous step yielded a valid pairing (line 15), the guide and tentative pieces are finally snapped together at that location (line 16) and with a distance defined by the `pieceDistance` parameter (see Table 3). However, if no valid pairing was found, the algorithm will keep the same guide piece and randomly select another tentative piece from `piecesList`. This process is repeated until a valid pairing is found or a limit of failed attempts (defined in the `maxFails` parameter, see Table 3) is reached for the current guide piece (line 19).

Whether or not a new piece was placed on the map during the tentative piece selection and placement loop (lines 7–19), a new guide piece will be chosen by the selection method (line 20). As will be discussed in the next subsection, selection methods decide the guide piece based on the pieces currently placed on the map. Consequently, if the tentative piece selection and placement loop was unsuccessful, the selection method will determine the next guide piece based on exactly the same scenario. Thus, if the selection method is deterministic, the same guide piece will be picked again. This leads to an infinite loop if no other piece can be connected with the guide piece in question. Therefore, the algorithm needs to detect if the same guide piece is chosen two times in a row and the number of free connectors in it remains the same—which most likely means the generation process is unable to go any further. In the proposed algorithm, this detection mechanism is assumed to be encapsulated in the `SelectGuidePiece()` function (line 20), and works by returning `none` when such a scenario is detected. The function may also return `none` according to the internal logic of the selection method in place, as discussed in the next subsection. In any case, when no guide piece is returned, the map generation ends (line 21).

C. SELECTION METHODS

The different selection methods produce distinct map layouts by determining how the generation algorithm selects the starting piece (line 1 of Algorithm 1) and the guide pieces (line 20 of Algorithm 1). Four methods are proposed and discussed here, namely *arena*, *corridor*, *star* and *branch*. Fig. 3 shows examples of maps created by each of these methods.

Different methods have specific parameter sets, as detailed in Table 4. The first parameter, `starterConTol`, impacts all methods, determining the connector tolerance for selecting the starting piece. Selection methods choose the starting piece based on the amount of connectors it has, and `starterConTol` provides a tolerance in this selection.




FIGURE 3. Example outputs of the proposed selection methods, namely (a) *arena* selection method, (b) *corridor* selection method, (c) *star* selection method, and, (d) *branch* selection method. In these examples, the piece color is lighter the later it is placed on the map. Thus, the starting piece is shown with the darkest color, while the last snapped piece appears with the lightest color. Pieces are very simple and placed in grid-like fashion in order to highlight how the different selection methods work.

Subsections III-C1–III-C4 describe each of the four methods and their parameter sets with additional detail.

1) THE ARENA SELECTION METHOD

This method (example in Fig.3a) aims to create maps that sprawl in all directions from the starting piece, covering the surrounding area with geometry. It is a generic and useful approach, having been used for the Trinity game with satisfactory results, as discussed in Section V-C.

The piece with most connectors is selected as the starting piece. If there are multiple pieces with the same highest number of connectors and/or if more pieces are considered for selection due to the `starterConTol` parameter, one of them is picked at random. The goal is to have a central piece with many connections so that geometry can spread in various directions.

For the selection of the next guide piece, the *arena* method checks if the current guide piece has unused connectors, and if so, keeps it as the guide piece in the next tentative piece selection and placement loop. Otherwise, the piece placed immediately after the current guide piece is selected as the next guide piece. This way, by performing a breadth-first traversal of the search space, the method preferentially fills center locations near the original starting piece. The *arena* method stops when the number of pieces placed on the map is higher than `maxPieces`, as described in Table 4.

2) THE CORRIDOR SELECTION METHOD

The goal of the *corridor* selection method is to generate narrow maps where the geometry seemingly follows a line, as shown in Fig. 3b.

The piece with least connectors is selected as the starting piece. If there are multiple pieces with the same lowest number of connectors and/or if more pieces are considered for

selection due to the `starterConTol` parameter, one of them is selected at random. The selected guide piece is always the most recently placed piece, such that pieces will have at most two connections: one with the previous placed piece and the other with next compatible tentative piece. Thus, the method performs a depth-first traversal of the search space.

The main motivation for implementing the *corridor* method is its potential suitability for “Capture the Flag”-style matches, in which two players or teams must capture the enemy’s flag from their end of the map and bring it back.

3) THE STAR SELECTION METHOD

This method is a mix of the *arena* and *corridor* selection methods, creating lanes sprawling from the starting piece and ending when that piece has no more available connectors. Thus, the number of arms in the “star” is equal to the number of available connectors in the starting piece, which acts as a central hub for the map layout.

The starting piece is selected the same way as in the case of the *arena* method, i.e., based on the highest amount of connectors. As with the other methods, the designer can force a specific piece or pieces to be used as the central hub by manipulating the `useStarter` and `starterList` parameters (see Table 1).

Each arm of the star, i , will have a uniformly random length (i.e., number of pieces) of $l_i = \text{armLength} \pm \text{armLengthVar}$ (see Table 4), which, nonetheless, cannot be smaller than one. For this purpose, the selected guide piece will be the last placed piece while the length of the current arm i is lower than l_i . When the length of the current arm reaches l_i , the starting piece (acting as the central hub) is returned again as the guide piece, allowing for the creation of a new arm. If this piece has no connectors left, the method returns *none*, ending the generation process.

The *star* selection method will try its best to generate a map with (a) a number of arms equal to the amount of connectors in the starting piece, and, (b) arms with length within the interval $\text{armLength} \pm \text{armLengthVar}$. However, this might not be possible for two reasons. First, the generation algorithm may fail to find a tentative piece compatible with the guide piece within `maxFails` attempts (see Table 3). Second, if there are pieces with a single connector, these would effectively work as premature arm terminals, in which case a new arm should be opened. Fig. 3c shows a map generated with the *star* selection method with four arms and arm length between 10 and 13 pieces.

A possible use case for the *star* selection method are “King of the Hill”-type game modes, in which players converge onto the middle of the map, fighting for control.

4) THE BRANCH SELECTION METHOD

This method (example in Fig. 3d) creates branches in a similar manner to the *star* method, except that it does not return the starting piece as the guide piece when a branch is finished. Instead, the *branch* selection method selects one of the previously placed pieces to start a new branch, repeating

TABLE 4. Parameters for the selection methods.

Method	Parameter	Description
All	starterConTol	Integer. The starting piece is selected among the set of pieces with most (<i>arena, star</i>) or fewest (<i>branch, corridor</i>) connectors, n_{\max} or n_{\min} , respectively. This parameter represents the tolerance, in number of connectors, from the piece(s) with most (or fewest) connectors, allowing pieces with as few as $n_{\max} - \text{starterConTol}$ (or as much as $n_{\min} + \text{starterConTol}$) connectors to be selected as the starting piece.
<i>Arena, Corridor</i> <i>Star</i>	maxPieces	Integer. The maximum number of pieces the generator will place on the map.
	armLength	Integer. The average amount of pieces an arm of the “star” will have.
	armLengthVar	Integer. The maximum variation from armLength in each arm.
<i>Branch</i>	branchCount	Integer. The number of branches to be created.
	branchLength	Integer. The average amount of pieces a branch will have.
	branchLengthVar	Integer. The maximum variation from branchLength in each branch.

this until the specified number of branches (`branchCount`, Table 4) is reached. Furthermore, in contrast to the *star* method and similarly to the *corridor* method, pieces with fewer connectors are preferentially selected as the starting piece.

The exact process of selecting the guide piece is as follows. The first branch originates at the starting piece. The branch grows by using the last placed piece as the guide piece until its maximum size (uniformly randomly drawn from the interval $\text{branchLength} \pm \text{branchLengthVar}$) is reached. At this time, it is necessary to select another map piece to be the root piece of a new branch. This selection is done by considering a list of pieces already placed on the map, from the starting piece to the last successfully placed tentative piece. The root piece is then selected by “jumping” from the starting piece in this list (index 0) to one of the other pieces. The base jump value is given by:

$$j_{\text{base}} = \max \left\{ 1, \left\lfloor \frac{\text{branchCount}}{\text{branchLength}} \right\rfloor \right\}$$

For each new branch i , where $i \in \{0, \dots, \text{branchCount} - 1\}$, the effective jump from the starting piece is given by:

$$j_i = i \cdot j_{\text{base}}$$

i.e., new branches will be based off pieces that were deployed after the root piece of the previous branch. If the new guide/root piece does not have available connectors, and to avoid premature termination of the generation process, the neighborhood $[j_i - j_{\text{base}}, j_i + j_{\text{base}}]$ is searched back and forth from its center, until a piece with available connectors is found and returned as the new guide/branch root piece. If such piece is not found in the neighborhood, the method returns *none* and the generation process terminates.

The *branch* method offers a way of generating non-linear maps without a central hub region, increasing exploration possibilities for players.

D. AUTOMATIC VALIDATION OF MAP NAVIGABILITY

Depending on the generation parameters, or simply due to “unlucky” seeds, it is possible that the procedurally generated maps are mostly untraversable or contain several unconnected regions. Although the snappable meshes technique produces fully explainable outputs, designers may be more

interested if a generated map is actually playable [1]. Therefore, some sort of automatic validation of map navigability becomes an important, if not crucial aspect in this type of algorithm [39]. We propose two metrics for this purpose, which can be obtained by deploying a predetermined amount n of randomly distributed navigation points on the map and verifying their connectivity using standard path finding. These metrics are:

- 1) The average percentage of valid connections between navigation points, or \bar{c} .
- 2) The relative area of the largest fully-connected (i.e., fully-navigable) region, or A_r^{\max} .

The first metric, \bar{c} , is given by $\bar{c} = c_t/c_{\text{all}}$, where c_t is the number of traversable connections between all n navigation points, and c_{all} is the number of all connections, traversable or not. Note that $c_{\text{all}} = n(n - 1)/2$, i.e., c_{all} is equal to the maximum number of connections between nodes in an undirected graph. Consequently, this approach has $\mathcal{O}(n^2)$ complexity, which limits the amount of deployable navigation points, n .

The second metric, A_r^{\max} , requires determining the various fully-navigable—but separate—regions in the generated map. This can be done by analyzing each pair of connected navigation points as follows: (1) if they are both isolated, create a new cluster containing them; (2) if one is isolated and the other is not, add the former to the latter’s cluster; or, (3) if they belong to different clusters, merge the clusters. Then, the approximate relative area represented by each cluster or region can be obtained by dividing the number of navigation points it contains by n , with A_r^{\max} representing the largest of these areas. Again, this procedure has $\mathcal{O}(n^2)$ complexity and the same caveat.

The first metric offers a general view of map navigability. The second metric is arguably more useful, as it can be used, for example, to select a playable area for deploying agents and/or in determining if the largest region represents a large enough area of the map to be usable or above a predefined threshold.

While the algorithmic complexity of these procedures is $\mathcal{O}(n^2)$, in practice—and as will be discussed in Subsection V-A2—the upper bound for the number of navigation points, n , is well above what is required for performing fast and accurate computations of the two metrics.

E. CLASSIFICATION

The snappable meshes technique is a constraint solver with no backtracking, focused on speed and simplicity. The method is able to quickly create maps—as will be shown in Section V—and is thus appropriate for runtime generation. Simplicity is a consequence of the visual and easy-to-understand constraints, mainly in the form of mesh connector rules.

According to the PCG taxonomy described by Togelius *et al.* [11], the proposed technique is *offline*, *necessary*, *generic* and *stochastic*. It is *offline* since maps are generated during game development or before the start of a game session. The method is *necessary* because it provides the main structure of the levels (i.e., the 3D maps). Content is generated without taking the player's previous behavior into account, hence snappable meshes is a *generic* (or *experience-agnostic* [9]) PCG technique, as opposed to an *adaptive* (or *experience-driven* [9]) one. Finally, it is a *stochastic* technique, as it offers considerable map variability given the same set of input parameters.¹

There are three other categories in this taxonomy where the snappable meshes technique does not fall under a specific classification. The first concerns autonomy, which differentiates between *autonomous* versus *MICC* PCG approaches. As stated in Section II, this separation is not clear-cut. On one hand, the technique works autonomously after the designer defines its input (parameters and map pieces). On the other, it is the designer who creates the fundamental building blocks of the generated maps, hence affecting their look and feel. Thus, the snappable meshes technique is essentially autonomous, though stating that it does not employ some form of “mixed-authorship” seems inaccurate. The second category, *degree and dimensions of control* (or *controllability* [9]), specifies the ways in which content generation can be controlled, for example using random seeds or content/parameter vectors. The proposed technique generates content using both approaches simultaneously. More specifically, an (optionally seeded) pseudo-random number generator can drive content creation based on the content and parameters specified by the designer. The third category, *constructive versus generate-and-test*, defines whether content is generated in one pass (*constructive*) or oscillates between content generation and testing in a loop, until a suitable output is found. The snappable meshes methodology allows for both approaches. It has been used as a *constructive* method in the Trinity game, as will be discussed in Subsection V-C, but supports a *generate-and-test* loop using validation metrics such as the ones presented in the previous subsection.

IV. A PROTOTYPE IMPLEMENTATION IN UNITY

A standalone demonstration prototype of the snappable meshes PCG technique was implemented in the Unity game engine, leveraging its editor tools to handle the input of the human designer. This is a simple reference implementation

¹The method is technically *deterministic*, in the sense that it will generate the same output if also given the same seed.




FIGURE 4. Screenshot of the first-person demo mode on a map generated in the *Benchmark* scene.

based on the code originally developed for the *Trinity* game, discussed in the next section, and is provided as a Unity project—i.e., it must be experimented in the Unity editor. The aim of this prototype is to allow designers and researchers to explore the proposed technique, particularly in how the chosen set of map pieces and the different algorithm parameters influence map generation. The prototype is bundled with two predefined scenes, allowing interested users to get started quickly. These scenes, denoted *Benchmark* and *Artistic*, are configured to use contrasting sets of map pieces, and a number of example maps, further discussed in Section V, are respectively shown in Fig. 7 and Fig. 11. The code is fully documented and available at <https://github.com/VideojogosLusofona/snappable-meshes-pcg> under the open-source Apache 2.0 license, meaning it can be freely adapted and used in commercial contexts.

In Subsection IV-A, the reference implementation's designer workflow and its editor-based user interface are described, while relevant implementation details are highlighted in Subsection IV-B.

A. DESIGNER WORKFLOW AND USER INTERFACE

The designer workflow is divided into roughly three parts:

- 1) Setup and configuration of the generation process. At this stage, the designer imports and/or selects the map pieces to use for the generation process, and defines the generation parameters.
- 2) Map generation and validation assessment. The designer generates the map, which appears in the scene view; the validation metrics—described in Subsection III-D—appear in the console, allowing the designer to quickly assess the map's navigability.
- 3) Demo of an NPC traversing the map. The two previous steps occur in editor mode, i.e., when the “game” is not running. To get a first-person feeling for the generated map, the designer can enter play mode, in which case a demo of an NPC traversing the map starts, as exemplified in Fig. 4.

Naturally, this workflow can be repeated and iterated upon until the designer is satisfied with the results. The next subsections further detail these steps, as well as the user interfaces which allow for this interaction.

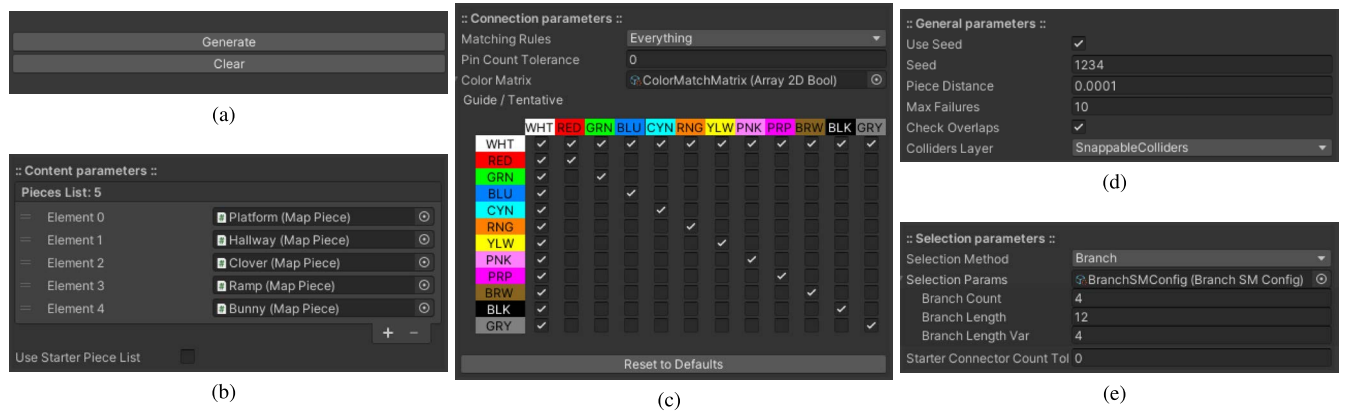


FIGURE 5. The map generation parameters and control buttons, part of the GenerationManager component, as shown in Unity's inspector. The different generation parameter blocks, (b)–(e), essentially correspond to the parameters respectively presented in Tables 1–4. More specifically, subfigures show: (a) the control buttons; (b) the content parameters; (c) the connection parameters; (d) the general parameters; and, (e) the selection method parameters (here exemplified for the case of the *branch* selection method).

1) SETUP AND CONFIGURATION OF THE GENERATION PROCESS

The designer sets up the generation process by interacting with the GenerationManager component, displayed in Unity's inspector—see Fig. 5—when the game object of the same name is selected in the scene's object hierarchy. Here, the designer can generate a new map or clear an existing one (Fig. 5a), select the pieces used to create the map (Fig. 5b, Table 1), define the connection criteria (Fig. 5c, Table 2), specify general parameters (Fig. 5d, Table 3), and choose and configure a selection method (Fig. 5e, Table 4).

Configuring map validation is also performed at this stage. This is done in a separate inspector panel, which appears when the NavController game object is selected in the object hierarchy. Among other aspects, the designer can define the dimension of NPCs and/or players navigating the map, as well as the number of navigation points used for determining the navigation metrics, as explained in Subsection III-D.

2) MAP GENERATION AND VALIDATION ASSESSMENT

After configuring the generation process, the designer can create a new map by clicking the “Generate” button (Fig. 5a). Besides the produced map, the process also outputs a detailed generation log in Unity's console, listing each selected guide and tentative pieces, as well as successful and unsuccessful snaps, allowing the designer to follow all the steps of the generation process.

After the map is generated, the validation step is triggered, and a predefined number of navigation points are randomly placed on the map's surface in order to obtain the validation metrics previously described, namely, \bar{c} , the map's average relative connectivity, and, A_r^{\max} , the relative area of the largest fully-connected region. This information is also provided via a log in Unity's console (separate from the generation log).

These metrics can be used to request a regeneration of the level if the average connectivity is below a certain threshold and/or the largest region does not represent a large enough relative area of the generated map. Additionally, A_r^{\max} can

be used to define the playable area of the map, where new players or NPCs are spawned, for example.

3) DEMO OF AN NPC TRAVERSING THE MAP

Having successfully generated a map, the designer can enter Unity's play mode, where an NPC traverses the map using the navigation points (initially created for map validation) as randomly selected waypoints. The NPC is placed in the largest fully-navigable area of the map, also determined during the validation procedure, thus avoiding getting stuck in poorly connected regions. While this demo is not useful in a real-world automatic generation process, it constitutes a visual aid for the designer to get a first-person awareness of the generated map. An example of this demo mode is shown in Fig. 4.

B. IMPLEMENTATION DETAILS

In this section we highlight a number of implementation details in the Unity prototype which might be relevant for helping researchers understand the code and/or for developers implementing their own versions of snappable meshes.

1) PIECE DESIGN AND DEPLOYMENT

When the designer creates a map piece, it is necessary to specify its mesh, individual connectors (each with a color and a number of pins), and one or more colliders² when piece overlap is to be avoided (as discussed in Subsection IV-B2). In this context, human-designed pieces are created as *prefabs*, Unity's implementation of the Prototype design pattern [40]. Therefore, pieces in the *piecesList* and *starterList* (Table 1) are prototypes, and the pieces actually placed on the map are effectively copies of the original designs.

In this implementation, and for the purpose of snapping two meshes together, tentative pieces can only be rotated about their vertical axis, so that the involved connectors face each other. Therefore, meshes will never tilt or flip.

²A collider is Unity's terminology for a bounding volume.

2) OVERLAP DETECTION

An important part of the proposed technique is the ability to generate maps without intersecting geometry. To guarantee this, when a tentative piece is selected, an optional verification ensures that the piece does not overlap with existing geometry for each of its possible connections (line 12 of Algorithm 1). The `checkOverlaps` parameter (Fig. 5d, Table 3), determines whether this verification is performed or not.

For overlap verification to work, pieces must contain one or more box colliders, i.e., rectangular cuboid-shaped bounding volumes. These should approximately mirror the piece's shape, allowing Unity to detect if the tentative piece overlaps with existing map geometry in a quick and relatively accurate fashion. We opted for box colliders since in Unity, due to optimization concerns, general convex mesh colliders are limited to 255 triangles and often display inaccurate behavior. These box colliders are tagged in a separate object layer, so that the application is able to find them while creating the map and safely delete them when the generation process is finalized. By default, this layer is named *SnappableColliders* (Fig. 5d). Nonetheless, the user can specify another name.

3) SELECTION METHODS

The Strategy design pattern [40] was used to decouple the selection methods from the generation algorithm itself. What this means is that the different methods are placed in distinct classes derived from a common base class. Any existing selection method classes are then “discovered” using C#'s reflection, and matched with an appropriate configuration object. This object is then used by the `GenerationManager` component to present the user the available parameters for the current selection method, chosen from a drop-down list populated during the “discovery” process. Fig. 5e shows what is presented to the user when the *branch* method is selected.

This approach simplifies the creation of new selection methods, requiring only two classes to be implemented: one for the selection method itself, and the other for configuring it.

4) MAP VALIDATION AND DEMO

Map validation occurs after the generation process, and is performed using Unity's built-in navigation mesh—or *navmesh*—system. A navmesh is a mesh of convex polygons that define traversable areas on a map. These polygons can be considered nodes in a graph, with adjacent polygons forming valid paths, or in graph terminology, a connection between nodes. Thus, a path finding algorithm such as A^* can be used to determine if a path exists between any two nodes [41].

Unity's navmesh system allows the runtime creation of navmeshes on existing geometry, and is used in our reference implementation for this purpose. After a navmesh is created for a generated map, a predetermined number of navigation points is deployed in the navmesh. At this stage,

Unity's path finding system is used for determining if there are valid paths between each pair of navigation points. With this information, the metrics discussed in Subsections III-D and IV-A2 can be easily computed. Finding traversable paths for the first-person demo (Subsection IV-A3) is similarly straightforward.

To aid in the visual inspection of the generated maps, navigation points placed in the largest fully-traversable region of the map are rendered in green, while points deployed in other regions are shown in red (see Fig. 9).

V. CASE STUDIES

In this section several case studies are investigated with the purpose of highlighting the potential of the snappable meshes PCG technique, as well as its limitations. In Subsection V-A we start by analyzing how the different parameters and selection methods influence the maps generated with the *Benchmark* scene, focusing on eight illustrative examples. These same examples are then dissected from a benchmarking perspective, both in terms of generation and validation time, as well as concerning the quality of the proposed validation metrics. In Subsection V-B, with the goal of further exploring the generative capabilities of the snappable meshes technique, we evaluate four maps generated with the *Artistic* scene, since it provides a set of building blocks which is completely different than those available in the *Benchmark* scene. Finally, in Subsection V-C, we examine how the proposed method was used to generate the maps for the *Trinity* third-person shooter game, enabling the desired gameplay style and replayability.

A. ANALYSIS AND VALIDATION OF SNAPPABLE MESHES USING THE *Benchmark* SCENE

1) EXPERIMENTING WITH THE DIFFERENT PARAMETERS AND SELECTION METHODS

A number of experiments were performed in order to provide a better understanding of the capabilities and limitations of the proposed technique. These experiments were carried out using the *Benchmark* scene, which contains the map pieces shown in Fig. 6. Many different maps were generated during these experiments, and several illustrative cases are shown in Fig. 7, with their respective parameters given in Table 5.

Figs. 7a–7d display experimental maps created with the different selection methods. A typical *arena* method-generated layout (`maxPieces=13`) is shown in Fig. 7a, where the map's tendency to grow in all directions is clear. In turn, Fig. 7b shows a map generated with the *corridor* method (`maxPieces=20`). The corridor characteristics are not immediately obvious, since, in this particular example, each piece is being placed in a way that changes the direction established by the previous piece. A larger map, created with the *star* selection method, is displayed in Fig. 7c, where the starting piece is a “ramp” (Fig. 6d), which has five connectors. Therefore, and as expected, the *star* has five arms stretching out from the initial, center piece. Note that one of these arms—the one expanding to the right in the figure—starts




FIGURE 6. Map pieces included with the *Benchmark* scene of the Unity prototype implementation. Piece names and connector configurations are as follows (each box corresponds to a connector with a specific color and pin count): (a) “Platform”:

1	1	4
---	---	---

; (b) “Hallway”:

1	1	4	4
---	---	---	---

; (c) “Clover”:

4	4	4	4
---	---	---	---

; (d) “Ramp”:

1	1	4	4	4
---	---	---	---	---

; (e) “Bunny”:

2	2	4
---	---	---

. Note that white is configured as a wildcard color in the *Benchmark* scene.

TABLE 5. Parameters used to generate the maps shown in Fig. 7.

Map	Piece count	Selection method	Selection method parameters	Matching rules	Check overlaps	Max. fails	Piece distance	Seed
(a)	13	<i>arena</i>	maxPieces=13	Pins + Colors	✓	10	0.0001	−267402550
(b)	20	<i>corridor</i>	maxPieces=20	Pins + Colors	✓	10	0.0001	−2095385667
(c)	41	<i>star</i>	armLength=8, armLengthVar=2	Pins + Colors	✓	10	0.0001	277759099
(d)	56	<i>branch</i>	branchCount=4, branchLength=12, branchLengthVar=4	Pins + Colors	✓	10	0.0001	1388449552
(e)	20	<i>arena</i>	maxPieces=21	Colors ^a	✓	10	6.0000	811974397
(f)	21	<i>arena</i>	maxPieces=21	Pins	✓	10	0.0001	−359152709
(g)	52	<i>star</i>	armLength=10, armLengthVar=4	Pins + Colors	✗	10	0.0001	1242840355
(h)	58	<i>corridor</i>	maxPieces=120	Pins	✓	50	0.0001	−1444708658

Notes: The ‘Piece count’ column denotes the number of pieces effectively placed on the map. The *piecesList* parameter includes the five pieces shown in Fig. 6, with *useStarter* set to false. When applicable, *pinTolerance* is set to 0 and *colorMatrix* is as shown in Fig. 5c. The *starterConToL* selection method parameter is set to 0.

^aConnector colors are made visible in this example. Note that white connectors are configured as wildcards (see Fig. 5c).

unfurling at a higher altitude than the remaining arms. Finally, the map displayed in Fig. 7d depicts the intended behavior of the *branch* selection method, with new “branches” created at specific points in the previously placed pieces.

In Fig. 7e, the *pieceDistance* parameter was set to 6, creating a map with several “islands”. The connectors are shown to help visualize the connections. The matching rules in this experiment were set to “colors”, i.e., pieces snap together based only on connector color, ignoring their pin count. The color matrix was set to its default values (Fig. 5c), where connectors snap with connectors of the same color, and white is set as a wildcard color, i.e., white connectors are able to snap with connectors of any color. Changing the *pieceDistance* parameter in this way generates mostly untraversable maps,³ but allows for visual inspection or debugging, for example to verify if connections are occurring according to the specified matching rules.

³Assuming characters without the ability to jump or fly between areas.

Contrary to the previous experiments, a pins-only matching rule was used to generate the map displayed in Fig. 7f. As can be observed, the map has the prototypical *arena* layout, but the type of connections—without the color matching constraint—are considerably different from those shown in Fig. 7a.

An interesting experiment is shown in Fig. 7g, where the overlap checking is disabled. Since pieces are snapped together without consideration for collisions between them, the resulting map loses the flow and clean appearance observed in the remaining experiments. Disabling this option will, in most use cases, likely create map-wide geometry and texture misalignments. Nonetheless, this may be desirable in some situations.

The goal in the last of these experiments was to create a long corridor-like map. The *corridor* selection method, with the *maxPieces* parameter set to 120, was used for this purpose. To avoid premature termination of the algorithm, a pins-only matching rule was defined—thus eliminating the color matching constraint—and the *maxFails*

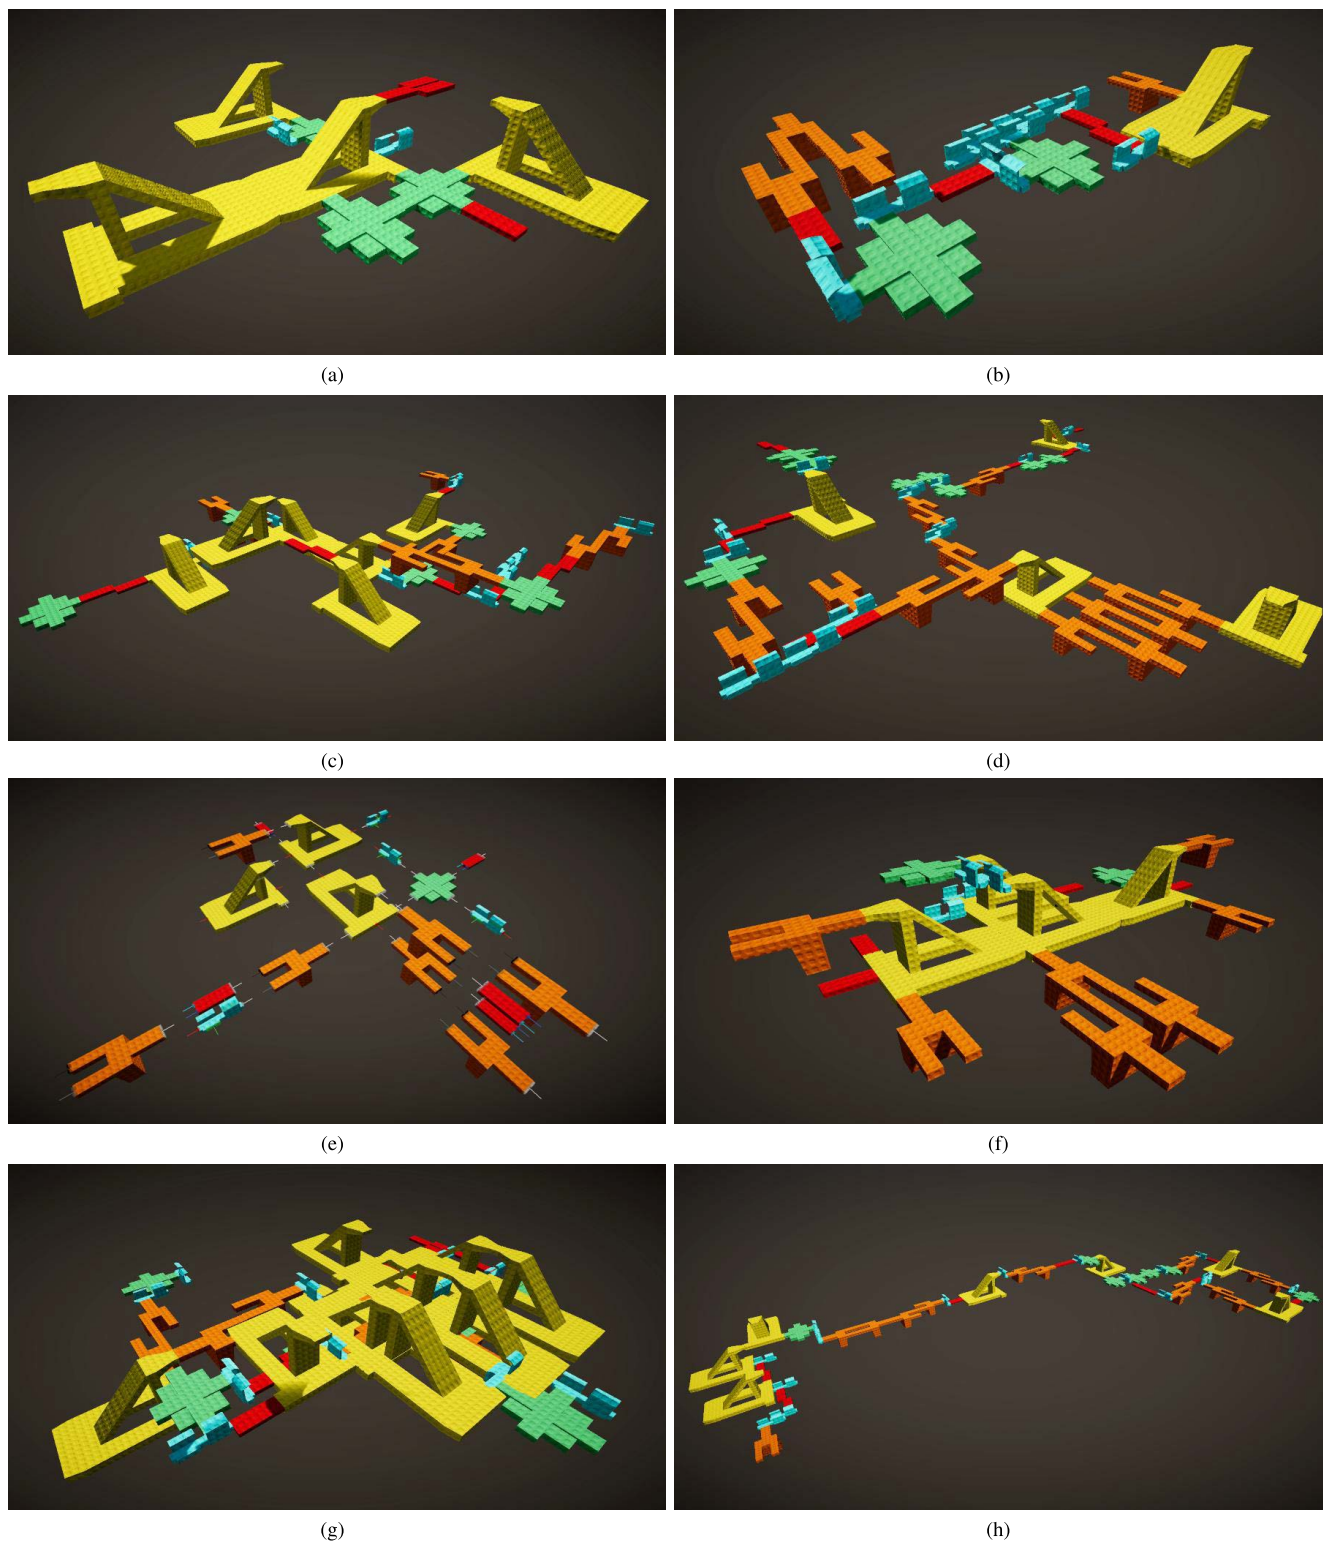


FIGURE 7. Several illustrative maps generated with snappable meshes using the *Benchmark* scene's map pieces presented in Fig. 6. The parameters used to generate each of these maps are given in Table 5.

parameter was set to 50 (a value of 10 was used in the previous experiments). The resulting map, displayed in Fig. 7h, is indeed a long corridor; however, the total number of pieces composing it is less than half of maxPieces . This highlights the fact that, depending on parameter

and geometry constraints, the generated map may end up being smaller than envisaged by the designer. Thus, it may be important to verify if the generated map attained some minimum threshold concerning the number of placed pieces.




FIGURE 8. Mean map generation times in milliseconds, obtained by generating each map (a) to (h) 30 times. Error bars denote the sample standard deviation.

2) BENCHMARKING THE EXAMPLE MAPS

The illustrative maps generated with the *Benchmark* scene, displayed in Fig. 7, were benchmarked and analyzed from three perspectives: (1) the duration of the generation step; (2) the time it takes to accurately execute the validation step; and, (3) the overall navigability of the generated maps. The first two viewpoints clarify if the snappable meshes technique, as well as the proposed validation step, allow for levels to be generated at runtime—especially given the quadratic nature of the validation metrics. The third viewpoint demonstrates whether the snappable meshes technique is overall able to generate valid, navigable maps. All experiments were performed with an AMD Ryzen 7 5800X CPU, Ubuntu 20.04.3 LTS and Unity 2020.3.25f1 LTS.

The mean generation time was obtained by generating each map 30 times. It was observed to be as low as 8.6 ms for map (a) (Fig. 7a), and up to 42.4 ms for map (d) (Fig. 7d). Generation times for all the maps are displayed in Fig. 8. Even for maps of considerable dimensions, with over 50 pieces, the proposed technique seems to be sufficiently fast to be used in a runtime level generation scenario.

The validation step was performed with the number of navigation points, n , set to 50, 500 and 5000. An example of how these points are placed on a map is shown Fig. 9, for the case of map (b). The following data was collected over 10 runs for each combination of map and n : \bar{c} (average percentage of valid connections between navigation points), A_r^{\max} (relative area of the largest fully-navigable region), number of regions which are not connected to any other, and the duration of the validation step. Results are presented in Fig. 10, while the raw data and respective analysis are available online [42].

The quadratic nature of the validation step, explained in Subsection III-D, is made obvious from the results shown in Fig. 10d: a $10\times$ increase in the number of navigation points leads to an approximately $100\times$ longer validation step. While the duration of the validation step for 50 navigation points—in the order of a few milliseconds—is certainly acceptable for runtime map generation (and likely feasible with 500 points, with validation times of a few seconds at most), that is clearly

not the case for 5000 points, where validation can take up to several minutes. The question here is whether there is a significant difference in validation accuracy when using considerably more navigation points. As Fig. 10a and Fig. 10b show, validation results are quite similar when deploying 50, 500 or 5000 points, with no clear tendency for an increase or decrease in percentage for both the \bar{c} and A_r^{\max} metrics. Thus, it seems possible to conclude that a relatively low amount of navigation points produces sufficiently accurate validation metrics, while being fast enough for runtime map generation. In turn, the number of detected low-navigability areas increases with n (Fig. 10c), which is to be expected, since a higher point coverage boosts the chances of finding these small, isolated regions. However, this is not relevant for runtime map generation, since the main concern there is in finding the largest navigable region, and this is done successfully even with very few points, as highlighted in Fig. 10b.

With respect to the overall navigability of the generated maps, results hint that the proposed method is robust, yielding highly-traversable maps for a variety of parameterizations. We have performed a large number of additional experiments, with various parameters and seeds, and have observed that \bar{c} is rarely below 90%, while A_r^{\max} is typically above 95%. This is of course assuming sensible parameterizations. For example, map (e), shown in Fig. 7e, has very poor navigability due to the large `pieceDistance`. An interesting aspect in these validation experiments, which can be observed in the red navigation points in Fig. 9, is that the low-navigability areas in the generated maps are essentially limited to the rooftops of the “hallway” piece (Fig. 6b). Therefore, it is clear that piece design has considerable influence on validation metrics, and that the results presented here, although showing promise for general use cases, should be mainly considered in the context of the utilized pieces.

B. EXPLORING THE GENERATIVE CAPABILITIES OF SNAPPABLE MESHES

In this second case study, we evaluate four maps generated with the *Artistic* scene, also bundled with the Unity prototype implementation. This scene is configured with a set of pieces notably distinct from those used in the previously discussed *Benchmark* scene, featuring cleaner and seamless connection interfaces. The scene includes 6 “hub”-like pieces and 6 “corridor”-like pieces, some of which have stairs or ramps, allowing for multistory maps. All 12 pieces are specified in `piecesList`. However, the `starterList` contains only the “hub” pieces (thus, `useStarter` is set to `true`). The “hub” pieces have 4 or 5 connectors with 2 pins each, while the “corridor” pieces have 2 or 3 connectors with 2 or 3 pins each. All connectors are white, thus connections are made only by pin count, with `pinTolerance` set to zero. Therefore, “hubs” can snap with other “hubs” without restriction (other than having free connectors), while “corridors” enforce some constraints on what connections are possible. In particular, some “corridor” connectors can only snap to compatible




FIGURE 9. Another perspective of the map in Fig. 7b, with the navigation points shown. Green navigation points belong to the largest fully-traversable region of the map, while red navigation points belong to other regions. (a) Using 50 navigation points; (b) using 500 navigation points; (c) using 5000 navigation points.




FIGURE 10. Validation metrics obtained with 50, 500 and 5000 navigation points, for the maps shown in Fig. 7, with parameters given in Table 5. Bars display mean values obtained by validating each map 10 times, while error bars denote the sample standard deviation. Experiments were performed on the same hardware and software setup used for obtaining the generation times. The following metrics are presented: (a) average percentage of valid connections between navigation points, \bar{z} ; (b) relative area of the largest fully-navigable map region, A_r^{\max} ; (c) number of isolated regions, i.e., of regions which are not connected to any another; and, (d) duration, in seconds, of the validation step—note the logarithmic scale.

connectors on other “corridor” pieces, while some other “corridor” connectors can also connect to “hubs”. The motivation for this design was in creating maps with larger areas connected with “corridor”-like pieces.

The generated maps, shown in Fig. 11, were created with the *arena*, *corridor*, *star*, and *branch* selection methods, respectively. In all instances, `pieceDistance` and `checkOverlaps`—relevant parameters for this discussion—

were set to 0.0001 and `true`, respectively. The complete configuration for these examples is predefined in one of the *Artistic* scene’s panels, making them easily reproducible. The four maps are 100% navigable—essentially due to piece design—and generation/validation times are very similar to what was observed in the case of the *Benchmark* scene. Therefore, we will focus this analysis on the maps’ aesthetic properties and their gameplay characteristics, highlighting

the diverse generative possibilities put forth by the snappable meshes algorithm.

The first map generated with the *Artistic* scene, shown in Fig. 11a, was created with the *arena* selection method, with `maxPieces` set to 18, although only 5 pieces—4 “hubs” and 1 “corridor”—are actually placed on the map before the algorithm terminates. The map is essentially a single concentrated area, typical of the *arena* selection method, and offers a number of open areas, cover positions and vantage points. Due to piece design, the map presents a more organic and playable style than the maps shown in Fig. 7, while stairs in two “hub” pieces promote the mentioned vantage points.

The second example, presented in Fig. 11b, consists of a map created with the *corridor* selection method. This is a linear map, with clear extremities. The stairs in two of the “hub” pieces allow the map to have varying altitudes along its linear path. The `maxPieces` parameter was set to 12, and this is the number of building blocks effectively placed on the map. The starting piece is the “hub” to the right edge of the figure, somewhat counter to how the *corridor* method works, since it chooses the block with least connectors as the starting piece. Note, however, that in the *Artistic* scene, the `starterList` only contains “hub” pieces, emphasizing the flexibility of the snappable meshes technique.

Fig. 11c demonstrates the prototypical *star*-generated layout, showing a basic center “hub” piece with four arms opening up in all directions. Naturally, the *star* selection method was used to create this map, with `armLength` and `armLengthVar` set to 5 and 2, respectively. The actual length of the four arms is, from the perspective of the center piece, 3 (left), 4 (top), 6 (right), and 5 (bottom), which, together with the center piece, yield a map with a total of 19 building blocks. Two of the arms are essentially “corridors”, while the other two also contain “hub” pieces.

The *branch* selection method was used to create the fourth and last example, shown in Fig. 11d. The map displays several larger areas composed of “hubs”, connected by “corridors”, as intended in the piece design. The branching effect provides a non-linear level experience with multiple areas—appropriate for exploration, for example. The map has a total of 70 pieces, and the selection method parameters, `branchCount`, `branchLength`, and `branchLengthVar`, were set to 5, 12, and 2, respectively.

These examples offer a broader picture of the proposed technique’s generative capabilities. An important aspect that stands out in these four maps is their free-flow nature, which would be difficult to achieve with grid-based map PCG methods. Additionally, and even though the algorithm does not explicitly consider multistory levels, it is still possible to add a sense of verticality with appropriate piece design, further underlining the technique’s versatility, as well as its reliance on properly constructed building blocks.

C. TRINITY—A THIRD-PERSON MULTIPLAYER SHOOTER

Trinity is a competitive, split-screen multiplayer game (Fig. 12), developed as a semester project at Lusófona

University’s Bachelor in Videogames [5]. It is a third-person shooter in which players navigate the environment trying to eliminate their opponents using weapons that shoot different types of ammunition with various effects and counter-effects.

The snappable meshes algorithm was initially developed to create maps for this game, generating a new procedural playfield at the start of each match. The method allowed designers to quickly experiment with many different map types during the development stage, and, together with testers, to swiftly home in on a set of map pieces and algorithm parameters deemed suitable for the goals set forth for *Trinity*. In the final game, a single set of map pieces and parameters were selected for generating maps. The *arena* selection method was chosen since it created satisfactory layouts—typically large areas with both open spaces and plenty of options for platforming and cover. The *corridor* method was also considered. It was better suited for “capture the flag”-style matches due to its tendency to create layouts that flowed along one direction. However, this game mode was not implemented, thus *corridor* generation was left out of the final version of the game.

Although limited to one parameter set in the final game, the technique enhanced the game’s replayability. More specifically, the procedurally generated levels changed the conventional play style of shooter games, since in *Trinity*, players have to quickly adapt to the playfield, instead of memorizing and using prior knowledge of the levels.

VI. DISCUSSION

In this section we start by contextualizing snappable meshes as a designer-centric approach, framing our reasoning with recent literature (Subsection VI-A). A discussion of the technique’s use cases is undertaken in Subsection VI-B, where its potential as a prototyping system is highlighted. Subsection VI-C points out a number of limitations in this study, namely at the level of the proposed method and the provided prototype implementation, as well as their evaluation. Finally, a number of alternative uses and possible improvements are explored in Subsection VI-D.

A. SNAPPABLE MESHES AS A DESIGNER-CENTRIC APPROACH

Togelius et al. [11] put forward five desirable properties of a PCG solution, namely *speed*, *reliability*, *controllability*, *expressivity/diversity*, and *creativity/believability*. As shown in the previous section, the snappable meshes algorithm is relatively *fast*, with generation times in the order of milliseconds, and allows the general look and feel of the generated maps to be *controlled*, although not their exact layout. The method is also able to provide *reliable* maps—from a traversability perspective, at least—if the generated maps are validated with the criteria defined in Subsection III-D. Since the technique requires map pieces to be created and provided by the designer, it has a high *expressivity* and *diversity* potential, being able to generate distinct maps given different building blocks. However, for this same reason, the snappable




FIGURE 11. Example maps created under the *Artistic* scene in the Unity prototype implementation. (a) Using the *arena* selection method; (b) using the *corridor* selection method; (c) using the *star* selection method; (d) using the *branch* selection method.



FIGURE 12. Screenshot of *Trinity* during a match.

meshes technique cannot, by itself, guarantee *creative* and *believable* maps. If given poorly designed pieces, the method will likely produce maps that are neither.

According to Zhu *et al.* [1], a common issue with AI techniques in general and PCG methods in particular is that of increasing algorithmic complexity, which hinders the designer’s understanding and trust about what the algorithm is doing. Consequently, designers are likely to avoid using such techniques to their full potential or not use them at all. A possible solution or mitigation for this problem is to develop, from the ground up, designer-oriented and fully explainable PCG methods and techniques [1]. The proposed snappable meshes technique was outlined with these considerations in mind. It has a straightforward and designer-centric workflow, and is fully explainable, as shown by the generated logs in the Unity prototype implementation. As suggested by Zhu *et al.* [1], the generative process is narrated in the form

of a sequential textual description, in which the algorithm’s decisions are explained.

Snappable meshes imposes few restrictions. Issues can generally be solved via editing map pieces, changing the algorithm parameters, or, if necessary, by creating new selection methods. Further, it provides immediate results, avoiding lengthy or computationally costly procedures common in search-based [10] and interactive evolution approaches, which may result in user fatigue [12], as well as in difficulties in fully understanding how to control the generation process [1], [43], [44]. In edge cases, new selection methods could implement or incorporate search- or machine learning-based strategies in a compositional fashion [45]—though this would undermine the simplicity and speed offered out of the box by the proposed method.

In addition, the snappable meshes technique aims to respect the designer according to the three pillars set forth by Lai *et al.* [3], namely (1) *respect designer control*, (2) *respect the creative process*, and, (3) *respect existing work processes*. In regard to pillar 1, the proposed technique respects designer control since, as already discussed, it provides “enough control to bring out the designer’s vision” [3]. Pillar 2, respecting the creative process, “concerns itself with having a feedback loop that is short enough that the creative process is not disturbed” [3]. This is guaranteed by the technique’s short generation and validation time. Finally, the snappable meshes approach respects existing work processes, the third pillar, since the algorithm can be integrated in existing workflows (e.g., game engines), and make use of existing assets—although these need to be “decorated” with connectors.

The validation metrics, discussed in Section III-D, offer yet another perspective of respect for the designer, as they cater for designers mainly interested on why a map is “deemed unplayable by the AI agent” [1].

Considering the increasing rift between academia and industry related to communication and used methodologies [3], we believe the simplicity, explainability, extensibility and respect for the designer embedded in the proposed technique grants it the potential to reduce this gap.

B. USE CASES

The most obvious use case for the snappable meshes technique is the one it was initially developed for: generating maps for 3D multiplayer shooters, possibly in a generate-and-test loop to guarantee adequately navigable levels. This is, however, a reductionist view of the technique’s potential, since there is nothing limiting its use in other game genres or scenarios. On the contrary: since the fundamental blocks of a map are created by the human designer, and given the possibility of adding new selection methods, snappable meshes can be considered more of a meta-PCG approach rather than a concrete algorithm or tool for specific use cases. However, framing the technique on such general terms is also not helpful, as one might be tempted to simply state that it can eventually create anything. Therefore, we will highlight an important use case, not related to any particular game type or genre: using snappable meshes as a prototyping method or visual map design approach [44], appropriate to designers with little to no programming experience, boosting their design space [46], allowing for fast iteration and speeding up development.

Particularizing on the use of the proposed technique as a visual prototyping method, connectors can be used as measuring aids for spacing in the game world. A designer can easily define a passage as n pins wide or tall, keeping consistency in the design of the layout of the individual pieces being made separately. Not only can a snappable meshes implementation be used as a level creation tool, but the generated maps can also be used to kickstart discussions between collaborating developers and to create basic rules for the construction of pieces, even if the algorithm ends up not being used in the final designs. Further, by using the color matching rules, the pieces developed by one designer can be grouped in the final outputs, allowing for focused design and prototyping of pieces belonging to specific areas or sections that can be seamlessly combined together while keeping a mixed-authorship approach (both human-human and human-computer) throughout the whole process.

C. LIMITATIONS

While the proposed approach frees the designer of grid and space restrictions when creating map pieces, care must be taken in their design in order to maintain cohesion and predictability of the generated output. In other words, and highlighting what was stated in Subsection VI-A, the algorithm will most likely produce poor maps given substandard

building blocks. Piece design becomes even more important considering that, to promote simplicity and explainability, the algorithm works in a greedy fashion—i.e., the first random piece that “fits” is selected—and does not perform backtracking. Since the search space is not thoroughly explored in a single algorithm run, this may lead to weak configurations and/or early dead ends, an issue potentially aggravated by poor piece design. Nonetheless, this limitation is offset by the technique’s short generation times, which allow it to work in a generate-and-test loop with validation metrics such as the ones discussed in Section III-D. Thus, the algorithm can be executed multiple times, exploring the search space until a map with the desired qualities is found.

Again, due to the focus on simplicity and explainability, there is no planning or any kind of spatial analysis when placing new pieces on the map, which would be necessary for creating loops and faultlessly avoiding the dead ends mentioned in the previous paragraph. Disabling overlap detection solves both problems at the cost of elegant map designs, but this is far from an ideal solution.

Another issue with the snappable meshes technique concerns the navigability validation presented in Subsection III-D, where we proposed deploying a predetermined amount of randomly distributed navigation points on the map, and then verify their connectivity. In the prototype Unity implementation these points are randomly placed in the runtime-generated navmesh. As shown in Fig. 9, some of these points may be placed on rooftops or other areas not intended for navigation; thus, if this approach is followed as-is for indoor maps with hollow map pieces, it may lead to invalid paths outside the intended play area. Consequently, individual map pieces may require additional metadata specifying valid movement zones.

Another limitation in the provided Unity prototype, also related with navigation, is the fact that it currently does not support jumps. While this is essentially an implementation detail—nothing stops an improved implementation of providing this functionality—it can hinder a more thorough experimentation by interested readers.

The generality of snappable meshes as a meta-PCG technique, as opposed to more objective PCG tools, creates some difficulties pertaining to runtime validation and comprehensive method evaluation. With respect to runtime validation, it is not feasible to exhaustively test or validate generated maps without knowing the specific context in which they will be deployed. As such, no validation metrics are presented in this paper other than navigability. This is the same reason why a comprehensive method evaluation is not performed in this paper. Such evaluations offer quality assurances on the generated content, and typically involve a top-down statistical analysis of the technique’s generation space [46]. An analysis of this kind is difficult to perform on an open-ended method such as snappable meshes, which, as already stated, is essentially a meta-generator, heavily dependent on the designer-provided blocks and lacking a predefined goal on the type of maps to generate. Thus, the difficulties in obtaining general

statistical assurances on the generated content are in effect a limitation of the snappable meshes technique.

D. ALTERNATIVE USES AND POSSIBLE IMPROVEMENTS

Given the open-endedness of the proposed technique, alternative or unanticipated uses are possible within the presented framework simply with smart and/or creative map piece design, as well as by implementing new generation methods.

An interesting possibility would be to use connector constraints (e.g., connector color constraints) to combine not only full map pieces, but also props, obstacles, and even simple cosmetic changes to those same pieces—possibly even player characters or NPCs. The possibility of creating more “final” maps also opens the door for validation metrics beyond navigability, such as cover ratios, target visibility, or detection of dangerous hotspots. Level difficulty could be assessed by determining path costs taking enemies and obstacles in consideration, as done by Togelius *et al.* [45], for example.

An innovative avenue of research would be to combine the input rearrangement and look-ahead capabilities of WFC with the free-form, gridless approach taken by snappable meshes. This would allow the creation of elegant and elaborate 3D levels given a single input example, such as levels carefully crafted by human designers or levels from existing games. Further, WFC’s more general constraint solving ability could potentially lead to more plausible environments with stricter constraints from a gameplay perspective.

Looking outside of the presented framework, the algorithm may be extended by performing multiple generation passes. This could potentially produce maps with several floors, adding explicit verticality to the designs, and foster more complex layouts by allowing the designer to specify different parameters for each pass. Additional passes could also be used to connect large separated clusters, for example.

VII. CONCLUSION

In this paper we presented the snappable meshes PCG technique for creating gridless 3D maps based on designer-modeled meshes with visual connection constraints. The approach was thoroughly described from an algorithmic perspective, and a Unity prototype implementation was introduced as a practical way of experimenting, testing, and studying the method. The case studies discussed in Section V showed that the technique can be used for generating levels in a concrete game scenario, and is able to generate a variety of map layouts, even with a limited set of building blocks. Further, both map generation and validation were shown to be fast procedures, opening the door for using snappable meshes in runtime generate-and-test loops. We argued that the technique respects the designer, offering a degree of control on the look and feel of the generated maps, while being adaptable to existing workflows. We also highlighted the potential of snappable meshes as a collaborative prototyping methodology, while discussing its limitations, alternative use cases and possible algorithmic improvements. In sum, the proposed method was shown to be a viable map creation

solution, allowing fast and/or collaborative level design and prototyping.

ACKNOWLEDGMENT

The authors would like to thank Phil Lopes for his valuable suggestions.

REFERENCES

- [1] J. Zhu, A. Liapis, S. Risi, R. Bidarra, and G. M. Youngblood, “Explainable AI for designers: A human-centered perspective on mixed-initiative co-creation,” in *Proc. IEEE Conf. Comput. Intell. Games (CIG)*, Aug. 2018, pp. 1–8.
- [2] R. M. Smelik, T. Tuteneel, R. Bidarra, and B. Benes, “A survey on procedural modelling for virtual worlds,” *Comput. Graph. Forum*, vol. 33, no. 6, pp. 31–50, Sep. 2014.
- [3] G. Lai, W. Latham, and F. F. Leymarie, “Towards friendly mixed initiative procedural content generation: Three pillars of industry,” in *Proc. Int. Conf. Found. Digit. Games*, Sep. 2020, pp. 1–4.
- [4] R. D. Peng, “Reproducible research in computational science,” *Science*, vol. 334, no. 6060, pp. 1226–1227, Dec. 2011.
- [5] N. Fachada and N. Códices, “Top-down design of a CS curriculum for a computer games BA,” in *Proc. ACM Conf. Innov. Technol. Comput. Sci. Educ.*, Jun. 2020, pp. 300–306, doi: [10.1145/3341525.3387378](https://doi.org/10.1145/3341525.3387378).
- [6] R. C. e Silva, N. Fachada, N. Códices, and D. de Andrade, “Procedural game level generation by joining geometry with hand-placed connectors,” in *Proc. Videojogos 12th Int. Conf. Videogames Sci. Arts*, I. Barbedo, B. Barroso, B. Legerén, L. Roque, and J. P. Sousa, Eds. Mirandela, Portugal: SPCV, 2020, pp. 80–93.
- [7] Unity Technologies. (2021). *Unity*. [Online]. Available: <https://unity.com/>
- [8] J. Aycock, “Procedural content generation,” in *Retrogame Archeology: Exploring Old Computer Games*. Cham, Switzerland: Springer, 2016, ch. 6, pp. 109–143.
- [9] G. N. Yannakakis and J. Togelius, *Artificial Intelligence and Games*. Cham, Switzerland: Springer, 2018, [Online]. Available: <http://gameaibook.org>
- [10] J. Togelius, G. N. Yannakakis, K. O. Stanley, and C. Browne, “Search-based procedural content generation: A taxonomy and survey,” *IEEE Trans. Comput. Intell. AI in Games*, vol. 3, no. 3, pp. 172–186, Sep. 2011.
- [11] J. Togelius, N. Shaker, and M. J. Nelson, “Introduction,” in *Procedural Content Generation in Games: A Textbook and an Overview of Current Research*, N. Shaker, J. Togelius, and M. J. Nelson, Eds. Cham, Switzerland: Springer, 2016, ch. 1, pp. 1–15.
- [12] A. Liapis, G. Smith, and N. Shaker, “Mixed-initiative content creation,” in *Procedural Content Generation in Games: A Textbook and an Overview of Current Research*, N. Shaker, J. Togelius, and M. J. Nelson, Eds. Cham, Switzerland: Springer, 2016, ch. 11, pp. 195–214.
- [13] G. Smith, J. Whitehead, and M. Mateas, “Tanagra: A mixed-initiative level design tool,” in *Proc. 5th Int. Conf. Found. Digit. Games (FDG)*, 2010, pp. 209–216.
- [14] P. Mawhorter and M. Mateas, “Procedural level generation using occupancy-regulated extension,” in *Proc. IEEE Conf. Comput. Intell. Games*, Aug. 2010, pp. 351–358.
- [15] A. Baldwin, S. Dahlskog, J. M. Font, and J. Holmberg, “Mixed-initiative procedural generation of dungeons using game design patterns,” in *Proc. IEEE Conf. Comput. Intell. Games (CIG)*, Aug. 2017, pp. 25–32.
- [16] S. Walton, A. Rahat, and J. Stovold, “Evaluating mixed-initiative procedural level design tools using a triple-blind mixed-method user study,” *IEEE Trans. Games*, early access, Jun. 3, 2021, doi: [10.1109/TG.2021.3086215](https://doi.org/10.1109/TG.2021.3086215).
- [17] A. Liapis, G. N. Yannakakis, and J. Togelius, “Sentient sketchbook: Computer-assisted game level authoring,” in *Proc. 8th Int. Conf. Found. Digit. Games*. New York, NY, USA: ACM, 2013, pp. 1–8.
- [18] P. Miglotzidis and A. Liapis, “SuSketch: Surrogate models of gameplay as a design assistant,” *IEEE Trans. Games*, early access, Mar. 23, 2021, doi: [10.1109/TG.2021.3068360](https://doi.org/10.1109/TG.2021.3068360).
- [19] P. T. Olsted, B. Ma, and S. Risi, “Interactive evolution of levels for a competitive multiplayer FPS,” in *Proc. IEEE Congr. Evol. Comput. (CEC)*, May 2015, pp. 1527–1534.
- [20] A. Apted. (Oct. 2017). *Oblige Level Maker 7.70*. Sourceforge. Accessed: Jul. 17, 2021. [Online]. Available: <http://oblige.sourceforge.net/>
- [21] E. Butler, A. M. Smith, Y.-E. Liu, and Z. Popovic, “A mixed-initiative tool for designing level progressions in games,” in *Proc. 26th Annu. ACM Symp. User Interface Softw. Technol.*, Oct. 2013, pp. 377–386, doi: [10.1145/2501988.2502011](https://doi.org/10.1145/2501988.2502011).

- [22] D. Yu, *Spelunky Classic*. San Francisco, CA, USA: Mossmouth, LLC, Dec. 2009. [Online]. Available: <https://www.spelunkyworld.com/original.html>
- [23] E. McMillen, *The Binding of Isaac: Rebirth*. Santa Ana, CA, USA: Nicalis, Nov. 2014. [Online]. Available: <https://www.nicalis.com/>
- [24] D. Roll, *Enter the Gungeon*. Austin, TX, USA: Devolver Digital, Apr. 2016. [Online]. Available: <https://www.dodgeroll.com/gungeon/>
- [25] D. Yu, *Spelunky: Boss Fight Books# 11*. Los Angeles, CA, USA: Boss Fight Books, Mar. 2016.
- [26] D. Kazemi, *Spelunky Generator Lessons*. Portland, Oregon: Tiny Subversions, Oct. 2013. Accessed: Jul. 13, 2021. [Online]. Available: <http://tinysubversions.com/spelunkyGen>
- [27] F. Himsl. (Apr. 2020). *Binding of ISAAC: Room Generation Explained*. Gamesquid—YouTube. Accessed: Jun. 19, 2021. [Online]. Available: <https://www.youtube.com/watch?v=1-HIA6-LBJc>
- [28] Trashboyjet. (Jul. 2020). *The Binding of ISAAC: Rebirth (Windows)*. Cutting Room Wiki. Accessed: Jul. 1, 2021. [Online]. Available: [https://trcf.net/The_Binding_of_Isaac:_Rebirth_\(Windows\)](https://trcf.net/The_Binding_of_Isaac:_Rebirth_(Windows))
- [29] Boris. (Jul. 2019). *Dungeon Generation in Enter the Gungeon*. Boris the Brave Personal Blog. Accessed: Jun. 29, 2021. [Online]. Available: <https://www.boristhebrave.com/2019/07/28/dungeon-generation-in-enter-the-gungeon/>
- [30] I. Dykeman. (May 2019). *Introducing Generate Worlds for Procedural Generation Without Code*. Personal Blog. Accessed: May 15, 2021. [Online]. Available: https://ijdykeman.github.io/procedural_generation/2019/05/17/introducing-generate-worlds.html
- [31] H. Wang, “Proving theorems by pattern recognition—II,” *Bell Syst. Tech. J.*, vol. 40, no. 1, pp. 1–41, Jan. 1961.
- [32] M. F. Cohen, J. Shade, S. Hiller, and O. Deussen, “Wang tiles for image and texture generation,” *ACM Trans. Graph.*, vol. 22, no. 3, pp. 287–294, Jul. 2003.
- [33] A. Lagae, *Wang Tiles in Computer Graphics* (Synthesis Lectures on Computer Graphics and Animation). San Rafael, CA, USA: Morgan & Claypool Publishers, 2009.
- [34] I. Dykeman. (Oct. 2017). *Procedural Worlds From Simple Tiles*. Personal Blog. Accessed: May 15, 2021. [Online]. Available: <https://ijdykeman.github.io/ml/2017/10/12/wang-tile-procedural-generation.html>
- [35] M. Gumin. (2016). *Wavefunctioncollapse: Bitmap & Tilemap Generation From a Single Example With the Help of Ideas From Quantum Mechanics*. GitHub. [Online]. Available: <https://github.com/mxgmn/WaveFunctionCollapse>
- [36] H. Kim, S. Lee, H. Lee, T. Hahn, and S. Kang, “Automatic generation of game content using a graph-based wave function collapse algorithm,” in *Proc. IEEE Conf. Games (CoG)*, Aug. 2019, pp. 1–4.
- [37] I. Karth and A. M. Smith, “WaveFunctionCollapse: Content generation via constraint solving and machine learning,” *IEEE Trans. Games*, early access, May 3, 2021, doi: [10.1109/TG.2021.3076368](https://doi.org/10.1109/TG.2021.3076368).
- [38] A. Newgas, “Tessera: A practical system for extended WaveFunctionCollapse,” in *Proc. 16th Int. Conf. Found. Digit. Games (FDG)*, Aug. 2021, pp. 1–7.
- [39] L. H. S. Leis, W. M. P. Reis, and Y. A. Gal, “Procedural generation of game maps with human-in-the-loop algorithms,” *IEEE Trans. Games*, vol. 10, no. 3, pp. 271–280, Sep. 2018.
- [40] R. Nystrom. (2014). *Game Programming Patterns*. Geneva Benning. [Online]. Available: <https://gameprogrammingpatterns.com/>
- [41] I. Millington, *AI for Games*, 3rd ed. Boca Raton, FL, USA: CRC Press, Mar. 2019, doi: [10.1201/9781351053303](https://doi.org/10.1201/9781351053303).
- [42] N. Fachada, *Snappable Meshes Performance Dataset*, document Zenodo, (Version 1.0.0) [Data Set], Jan. 2022.
- [43] A. Baldwin, S. Dahlskog, J. M. Font, and J. Holmberg, “Towards pattern-based mixed-initiative dungeon generation,” in *Proc. 12th Int. Conf. Found. Digit. Games*, Aug. 2017, pp. 1–10.
- [44] M. S. O. Almeida and F. S. C. da Silva, “A systematic review of game design methods and tools,” in *Proc. Int. Conf. Entertainment Comput. (ICEC)* in Lecture Notes in Computer Science, vol. 8215, J. C. Anacleto, E. W. G. Clua, F. S. C. da Silva, S. Fels, and H. S. Yang, Eds. Berlin, Germany: Springer, 2013, pp. 17–29.
- [45] J. Togelius, T. Justinussen, and A. Hartzen, “Compositional procedural content generation,” in *Proc. 3rd Workshop Procedural Content Gener. Games (PCG)*, 2012, pp. 1–4.
- [46] N. Shaker, G. Smith, and G. Yannakakis, “Evaluating content generators,” in *Procedural Content Generation in Games: A Textbook and an Overview of Current Research*, N. Shaker, J. Togelius, and M. J. Nelson, Eds. Cham, Switzerland: Springer, 2016, ch. 12, pp. 215–224.



RAFAEL C. E SILVA received the B.A. degree in videogames from Lusófona University, Portugal, in 2021. He has presented some of his work—games, papers, and posters—in conferences while in high school and university. Throughout his studies, he worked in several projects with game engines, VR, games with a focus on accessibility for the disabled, and Q&A analysis for game prototypes. His research interests include the areas of procedural content generation and in-engine tooling to improve game development workflows.



NUNO FACHADA received the B.Sc., M.Sc., and Ph.D. degrees in electrical and computer engineering from IST/Universidade de Lisboa, Portugal, in 2005, 2008, and 2016, respectively. After a year as a Postdoctoral Researcher at IST, he moved to Lusófona University, where he is currently an Assistant Professor at the School of Communication, Arts and Information Technology—teaching in the Videogames B.A. and Computer Science Ph.D. programs—and a Researcher at COPELABS. His main research interests include artificial intelligence, machine learning, modeling and simulation, research software, high-performance computing, and computer science education.



DIOGO DE ANDRADE received the B.Sc. degree in software engineering from IST/Universidade de Lisboa, Portugal, in 2000. He is currently pursuing the Ph.D. degree at the Computer Science Doctorate Program, Lusófona University. Since 2018, he has been a Lecturer at Lusófona University’s Videogames B.A. He has programmed and designed several games for companies, such as Spellcaster Studios, B5 Pocket Studio, and Fun Punch Games, having worked extensively in the DigitalTV industry, developing end-to-end systems and MPEG PSI/SI streamers. His research interests include procedural content generation, agent-based simulation, and computer graphics.



NÉLIO CÓDICES received the B.Sc. degree in computer science from FCT/NOVA University, Portugal, in 2004. He is currently pursuing the Ph.D. degree in computer science and engineering at IST/Universidade de Lisboa on the topic of procedural content generation in videogames. He lectures, since 2015, several game development courses at Lusófona University’s Videogames B.A., and IADE/Universidade Europeia’s Games Development B.Sc. He is the Co-Founder and the Director of the videogame development studio BATTLESHEEP.

• • •