PROCEDURAL CREATION OF 3D SOLUTION CAVE MODELS

Matt Boggus, Roger Crawfis
The Ohio State University
395 Dreese Laboratories
2015 Neil Avenue
Columbus, OH 43210-1277
United States of America
boggus, crawfis @cse.ohio-state.edu

ABSTRACT

There are a number of challenges in creating a three dimensional model of a cave. Data acquisition is difficult due to the high surface detail of rock so existing geometric data is often poorly sampled. We propose methods to create 3D cave models based on the type of cave passage as well as a set of cave patterns. Given a curve that represents flow of water, we can create a coarse level of detail model for a cave passage. Given a surface as input we can generate cave pattern images, which are then used to create a model of an entire cave system.

KEY WORDS
Procedural Modeling, Cave Passages, and Cave Patterns

1. Introduction

Visual modeling of large, detailed systems is a difficult task for practitioners of computer graphics. The challenge compounds when the availability of spatial and textural data is limited, as is the case for caves. One coping mechanism is procedural modeling: creating models with an algorithm. Recently, methods have been developed to generate man made creations such as buildings and cities [1]. There is also a broad body of work on modeling naturally occurring phenomena such as plants [2], fire [3], and water [4], yet caves remain relatively unexplored. Despite this, caves are often featured as environments in movies and computer games. Caves are common in earlier games such as Metroid and Lemmings as well as more recent series including Halo, Fable, and Crysis. Virtual 3D caves could also be useful for training simulations, as safety is a high concern for cave explorers.

Caves are natural underground voids large enough for a human to enter. They are found all over the world and can be carved out of many different materials, including limestone, granite, and glacial ice. Between ten and twenty percent of the Earth's land area consists of karst landscapes, where dissolution of bedrock is a primary feature of terrain [5]. This process is influenced by the underlying caves and fluid transport. Caves formed by rock being dissolved by acidic water are referred to as solution caves. In this paper we focus on solution caves as they account for the majority of observed caves on earth. They also are the most prominent type of cave in popular culture, featuring elaborate formations like stalagmites and stalactites, more formally referred to as speleothems.

1.1 Related Work

Maps are a standard method of displaying two-dimensional spatial information. Since caves incorporate depth information, they are more difficult to visualize. In addition to the overhead, or plan view, of a cave, an additional perspective from the side, or extended profile, is often included as part of a cave map. An example of a plan view of a cave is presented in Figure 1 and Figure 2 depicts a cave viewed from its side.

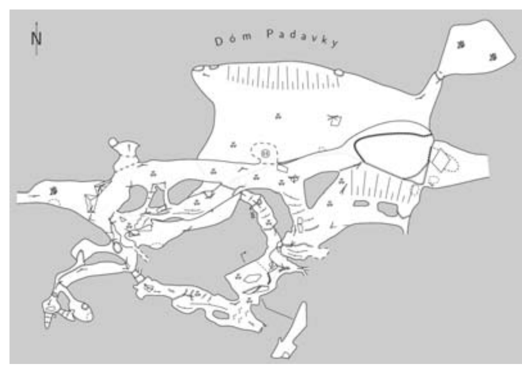

Figure 1: Cave map of Padavka Chamber (taken from [6])

Figure 2: Extended profile of Rabbit Cave (taken from [6])

There are a few existing software applications capable of creating cave maps [6] and [7]. Most spatial cave data is limited. During a cave survey, several locations in the cave are selected and measurements of the distances to the rock above, below, and the sides are recorded. Using this information, 3D line drawings can be made [8]. Some programs make polygonal models [6], [8], [9], and [10], but due to the limitations in the data they cannot capture the visual detail of an actual, physical cave. A Virtual Reality Modeling Language model of a cave is shown in Figure 3. Note that the model lacks a high amount of geometric detail.

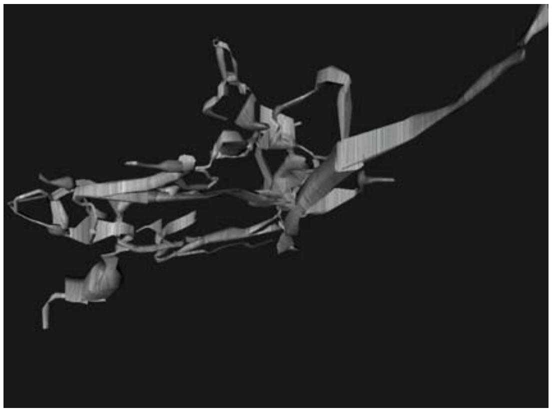

Figure 3: Screenshot of a model of Fulford Cave using Compass [9]

More accurate data can be acquired by making an excursion into the cave along with scanning equipment [11] and [12], but there is no publicly distributed data of this variety. Scanning is a very time consuming process. A 3D model of Wakulla Springs, an underwater cave approximately 6,409 meters in length, was creating using sonar data that took 76.1 hours to collect [13]. There is no standard for cave data, so there is also no standard data structure for representing caves. Gong, Lin, and Yin describe an object oriented model, representing the top and bottom surfaces of cave passages [14]. The vertices on each surface are stored in a triangulated irregular network, or TIN. This format allows only one height value for each 2D coordinate, which restricts the types of volumes that can be represented. A TIN cannot model a surface that includes an arch or overhang, features that are common to canyons and rocky terrain.

Geologists have discovered that caves form in distinctive, but reoccurring patterns. Currently, there is no work that capitalizes on this information, but similar ideas have been used in other areas of procedural content generation. Grammar based generation methods, such as L-systems, have been used to generate city street plans [1]. Given a set of curves on a 2D manifold, mazes are created by running a simulation that alters point positions [15]. There are also techniques to alter the creation of a maze based on an input image [16].

1.2 Our Approach

In this paper we discuss cave formation with a focus on the primary features of solution caves in order to define the plausibility of a procedurally created cave model. We also introduce methods of creating models of individual cave passages. Our last contribution is automatic generation of cave pattern images and their application in creating models of cave systems.

2. Solution Cave Morphology

Water absorbing carbon dioxide from the atmosphere or dead plants within the soil forms a weak carbonic acid. It descends through porous material until it reaches the water table. The water table is the surface where pressure from ground water is equal to atmospheric pressure. The rock below the water table is completely saturated with ground water. After dissolving soluble rock such as limestone, if the solution is carried away due to groundwater flow, the resulting void gradually opens over time to form a cave. Approximately 99% of caves form due to expansion of large pre-existing fractures or partings as opposed to expansion of tiny pores [5]. This is why caves often open along joints, which are gaps in rock, and bedding planes, which are the divisions between the individual strata of sedimentary rock [17].

There are two types of cave passages formed by dissolution: phreatic and vadose [5]. Phreatic passages are round and horizontal in appearance. An example of a phreatic passage is pictured in the top left of Figure 4. Over time, they widen in all directions. This is because they form below the water table; the region is fully saturated with water so dissolution of rock occurs evenly. Vadose passages are typically canyon-like in appearance. The bottom left of Figure 4 illustrates a vadose passage. They form in the region above the water table by water that is rapidly flowing downward towards the water table. Because the amount of water flowing fluctuates over time and the passage depth decreases over time due to dissolution of the floor, the width of the passage varies as it deepens. Depending on fault lines in the rock structure and changes in the water table, a combination of the two passages may form, with the top structure phreatic and the bottom vadose or vice versa. The two possible combinations are shown on the right side of Figure 4.

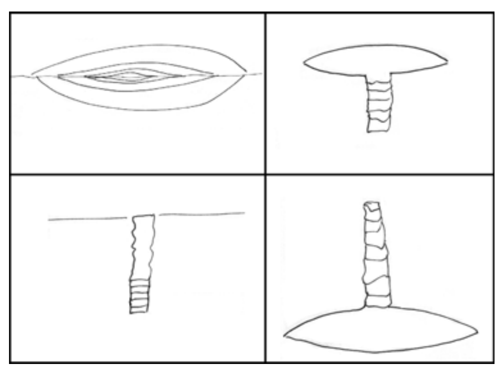

Figure 4: Cave passage cross sections

(Top left) Phreatic passage: The horizontal line indicates the bedding plane the passage is opening and the series of curves indicates removal of material over time.
(Bottom left) Vadose passage: The various horizontal lines within the passage indicate the removal of material over time.
(Right) Combination passages.

Figures 5 and 6 are photographs of cave passages in Mammoth Cave National Park. Inside a phreatic passage, in Figure 5, the path curves to the left, indicating the direction of the water flow that shaped the passage. The ceiling is a wide, rounded arch which is mirrored by the floor. The vadose passage in Figure 6 looks similar to a canyon.

Figure 5: Phreatic passage in Mammoth Cave National Park

Figure 6: Vadose passage in Mammoth Cave National Park

Even though cave formation depends heavily on underlying rock structure, as passages open further and begin to intersect with each other, patterns form [5]. These cave patterns are clearly visible in the plan view map of a cave, and a particular cave may contain multiple patterns. Branchwork patterns, as seen in the top row of Figure 7, resemble tributaries that converge into a single confluence further downstream. The streams can meander as on the left or they can be angular as on the right. The next category is maze caves, which contain multiple closed loops. As with branchwork, the maze cave passages may be curved or straight, with the former called anastomotic and the latter referred to as network. The middle row of Figure 7 depicts an anastomotic pattern on the left and a network pattern on the right. Maze caves are typically the result of periodic flooding that opens several paths at once. Spongework caves are formed by expanding pores and other small openings in rock that contains no major fractures. A final pattern, ramiform, consists of a central region with smaller, irregular sponge-like passages branching outward. The bottom row of Figure 7 displays an example of each of the final two patterns.

Cave passages can be considered as a coarse level of detail for a 3D cave model. Their shapes and connections are influenced by rock structure and the flow of water. Water also sculpts the fine details of a cave, creating cave scallops. Cave scallops are asymmetrical, spoon-shaped depressions formed by eddies flowing near cave walls [5]. They can also be made on floors and ceilings. Their size can vary greatly, from a diameter smaller than an inch to larger than a foot. A scallop's size is determined by the velocity of the water that formed it: the quicker the flow, the smaller the depression.

Speleothems are secondary mineral deposits that form in caves that are above the water table [18]. Water containing sediment may precipitate, slowly forming shapes based on how the water is moving. There are dozens of common speleothem formations including stalactites, stalagmites, columns, flowstone, and rimstone dams.

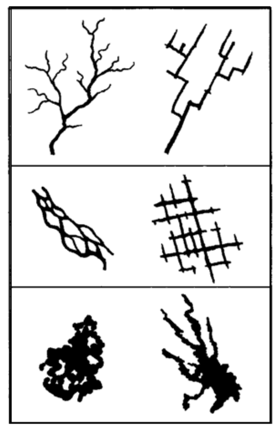

Figure 7: Cave Patterns (taken from [18])
(Top row) Two examples of branchwork patterns.
(Middle row) Two examples of maze patterns.
(Bottom row) Examples of spongework and ramiform patterns.

3. Procedural Cave Generation

3.1 Surface Generation

There are many methods that generate surfaces given one or more curves as input. Rotating a curve around an axis forms a surface of revolution. Projecting a curve along a vector creates a surface of extrusion. Specifying a curve of extrusion instead of a vector makes a swept surface.

In addition to synthetic surfaces, a plethora of real world terrain surface data is available thanks to the efforts of The National Geospatial-Intelligence Agency (NGA). Digital Terrain Elevation Data, or DTED, files hold values representing height (the z dimension) at regular intervals arranged in a grid (the x and y dimensions). Heightmaps are commonly used to store and visualize this data. The surface is constructed by connecting neighbouring values to form a triangular or quadrilateral mesh. There are also many methods specifically dedicated to procedurally creating synthetic terrain. Fractal methods are often used to create detail on demand [19], but there are also methods to make a base terrain such as the fault and particle deposition algorithms [20]. Thermal and hydraulic erosion algorithms can provide details to otherwise static terrain [21].

3.2 Surface Pair Generation

Given methods to create a single heightmap, it is easy to create a pair by running the methods again. Using different parameters creates two independent surfaces. Alternatively, we can reuse the existing data by reflecting it across the x-y plane. The advantage to using reflection is that the number of points above the z value for reflection is the exact same as the number of points on the surfaces that overlap. The amount of overlapping points is inversely proportional to the value of z. Figure 8 demonstrates the effect of lowering the plane of reflection. In the top image, there is no intersection. In the bottom image, there is overlap on the extreme right as well as on the left near the border.

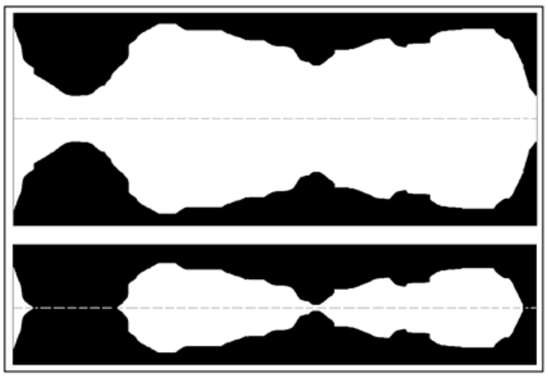

Figure 8: Cross section of heightmap reflection
The plane of reflection is indicated by the dotted line. When the plane is lower as in the bottom, there is more overlap between the two surfaces.

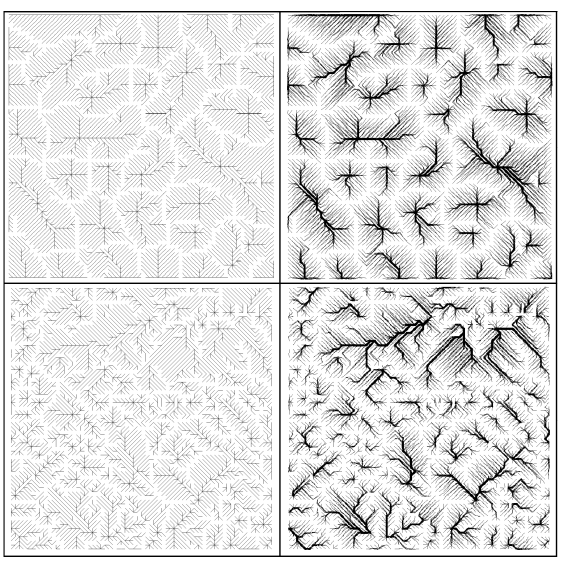

Figure 9: Generated branchwork cave patterns
The top two patterns were created from a synthetic heightmap while the bottom two were made from a DTED of a region in Montana. The patterns on the left are discrete whereas the ones on the right include noise to create a meandering effect.

3.3 Cave Patterns

Our goal here is to create images that share the same characteristics as cave patterns. Ultimately, formation of solution caves is driven by the movement of water. This is especially apparent for caves fitting the branchwork pattern. To procedurally create these patterns, we begin by creating or loading a heightmap. Then we select a point to represent a source of water. This selection can be random or deterministic. Next, we move the point to a neighbouring grid point according to the maximal downward slope, similar to the process of gradient descent. This models the tendency of water to follow the path of least resistance. The descent process repeats until the point is outside of the grid or has reached a local minimal value. The source selection and descent simulation is repeated, either for a specified number of iterations or until every grid point has been used as a source. This process results in a set of lines that are similar to branchwork patterns. Two results of this simulation are shown on the left of Figure 9. Since we allow movement only to the gridded positions in the heightmap, the patterns are angular. Adding noise smaller than the distance between points in the grid to the point locations creates a meandering effect as shown in the right side of Figure 9. Using the original point location sequences as input in approximation or interpolation algorithms can also break up the angular pattern.

Since spongework patterns tend to be more random, we begin by creating a heightmap with white noise values. Then we alter the data by setting each height value to the average of itself and its neighbours, which is done in parallel. This is the same as treating the heightmap as an image and applying a smoothing 2D convolution mask. After averaging several times, the resulting surface is smooth and resembles a hilly landscape. We select the height of the plane of reflection as the maximal value in the heightmap. Initially, the only overlap is one point, but as we decrease the height of the reflection plane, more overlap occurs between the surfaces. The volume below the surface of a heightmap is solid, so the volume above the reflected surface is solid as well. This means that only areas where the reflected surface is above the original surface are open, indicating the presence of a cave. As the reflected surface drops, more areas of the cave become closed.

The plane of reflection may be placed manually or using the following automated method. In order to determine an optimal height for the reflection plane, we compute the connected components of the regions with and without overlap of the surfaces. Ideally, we would like to have one large connected component for the areas without overlap to represent the cave area. After each decrease of height of the reflection plane, we recompute the connected components. Once the largest component of the cave is less than some user specified threshold of the total area we set the plane of reflection at its height in the previous iteration. This placement attempts to retain a large amount of area accessible while retaining the porous, but not completely open, structure that a spongework pattern has. Figure 10 depicts a view of the two surfaces at the end of this process. The darker areas are where the reflected heightmap has not yet descended to the original, which corresponds to the dark regions in a cave pattern image as in Figure 7.

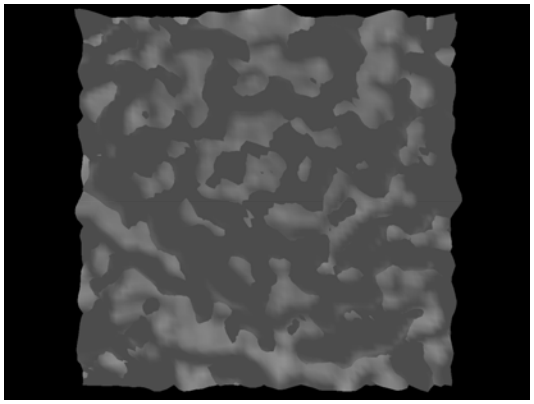

Figure 10: Spongework cave pattern
The lighter areas display where the original surface overlaps with the reflected surface, closing off parts of the cave.

3.4 3D Models from Cave Patterns

The resulting cave pattern images can be used in the creation of synthetic 3D cave models. White space on a cave map indicates the presence of a wall. When considering the heightmap reflection pair approach, this indicates a height value above the plane of reflection. In order to create a 3D model corresponding to the cave pattern, we begin with a heightmap holding real or synthetic terrain data. Since whitespace is the only region where walls are present, the plane of reflection should be placed above the highest point in the field. Using the cave pattern image as a mask, any point in the heightmap that pairs with a white value is raised to the plane of reflection, ensuring overlap that results in a wall. Figures 11 and 12 provide views inside of two virtual caves created using this method.

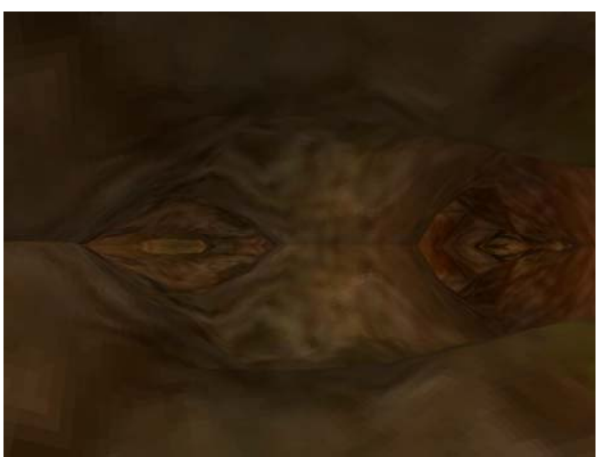

Figure 11: First person view inside the model from Figure 10

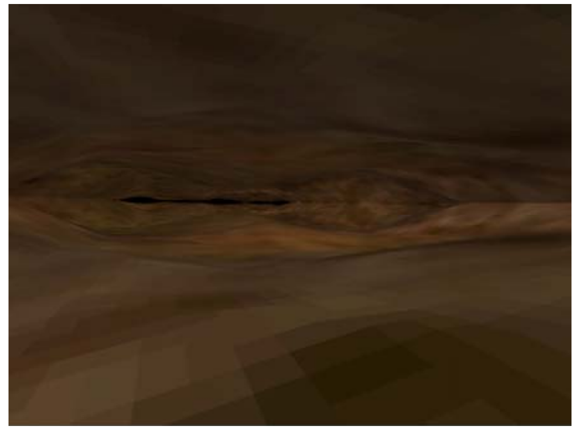

Figure 12: View inside of a procedural cave model

3.5 Single Passage Caves

Many caves consist of only a single passage. In this case we can create a 3D model based on the characteristics of its passage type. To make a model of a phreatic passage, after creating the rounded ceiling surface, we reflect it to form the floor. Vadose passages are more complex.

Since the width of vadose passages varies substantially as the height changes, two surfaces aligned to the x-y plane are not capable of modeling them. Instead, we align the two surfaces along the x-z plane to form the walls of the passage. Additional geometry is needed for the floor and ceiling, which can be as simple as connecting the two walls at specified heights. While this surface reflection technique does not match the characteristics of a vadose passage, because they are symmetrical to the flow of water rather than a plane, the method can still be used to create interesting artificial cave passages.

Modeling of passages should also include cave scallops. Since all scallops are small with respect to passage size, it is best to incorporate them only when a high level of detail is needed. Bump mapping and displacement mapping can be used to modify the appearance of a surface to display scallops.

4. Conclusion

With these methods, we can create synthetic 3D cave models with similarities to caves in the real world. However, these models do not include every feature of caves. Models generated from cave patterns will not feature phreatic and vadose passages. Creation of speleothems and populating them in caves is beyond the scope of this paper. Assuming prebuilt models exist for speleothems, they could be placed inside cave passages akin to the process of placing secondary features like vegetation on terrain.

We have demonstrated methods of generating 3D cave models via cave patterns. Existing programs modeling 3D caves do not take into account the processes that form cave passages and the overall patterns that they coalesce into. Accurate cave passage modeling requires information that is difficult or impossible to obtain. Raw geometric data is time consuming to obtain, finding the precise structure and weaknesses within a volume of rock is even more difficult, and determining accurate weather, precipitation, and fluid flow over the last several thousand years seems like a hopeless task. For caves, creation of synthetic models is a useful alternative.

Acknowledgements

We would like to thank Professor Edwin Scott Bair for his photos of Mammoth Cave National Park and information on Cave Geology.

References

[1] Y. I. Parish & and P. Müller, Procedural modeling of cities, *Proceedings of the 28th Annual Conference on Computer Graphics and interactive Techniques SIGGRAPH*, New York, NY, 2001, 301-308.

[2] O. Deussen, P. Hanrahan, B. Lintermann, R. Mech, M. Pharr, & and P. Prusinkiewicz, Realistic modeling and rendering of plant ecosystems, *Proceedings of the 25th Annual Conference on Computer Graphics and interactive Techniques SIGGRAPH*, New York, NY, 1998 275-286.

[3] A. R. Fuller, H. Krishnan, K. Mahrous, B. Hamann, & K. I. Joy, Real-time procedural volumetric fire. *Proceedings of the 2007 Symposium on interactive 3D Graphics and Games*, Seattle, WA, 2007, 175-180.

[4] R. Bridson & M. Müller-Fischer, Fluid simulation: SIGGRAPH 2007 course notes. *ACM SIGGRAPH 2007 Courses*, San Diego, CA, 2007, 1-81.

[5] A. N. Palmer, Origin and morphology of limestone caves, *Geological Society of America Bulletin, 103*(1), 1991, 1-21.

[6] Therion, http://therion.speleo.sk/. Accessed 2/10/2009.

[7] WinKarst, http://www.resurgentsoftware.com/winkarst.html Accessed 2/10/2009.

[8] The Survex Project, http://survex.com/ Accessed 2/10/2009.

[9] Compass, http://fountainware.com/compass/ Accessed 2/10/2009.

[10] P. Schuchardt & D. A. Bowman, The benefits of immersion for spatial understanding of complex underground cave systems. *Proceedings of the 2007 ACM Symposium on Virtual Reality Software and Technology*, Newport Beach, CA, 2007, 121-124.

[11] K. A. Robson Brown, A. Chalmers, T. Saigol, C. Green, & F. d'Errico, An automated laser scan survey of the Upper Palaeolithic rock shelter of Cap Blanc. *Journal of Archaeological Science 28*(3), 2001, 283-289.

[12] W. I. Sellers & A. T. Chamberlain, Ultrasonic cave mapping. *Journal of Archaeological Science 25*(9), 1998, 867-873.

[13] B. am Ende, 3D Mapping of Underwater Caves. *IEEE Computer Graphics Applications, 21*(2), 2001, 14-20.

[14] J. Gong, H. Lin, & X. Yin, Three-dimensional Reconstruction of the Yaolindong, *Cartography and Geographic Information Science*, *27*(1), 2000, 31-40.

[15] H. Pedersen & K. Singh, Organic labyrinths and mazes. *Proceedings of the 4th international Symposium on Non-Photorealistic Animation and Rendering*, Annecy, France, 2006, 79-86.

[16] J. Xu & C. S. Kaplan, Image-guided maze construction. *Proceedings of the 34th Annual Conference on Computer Graphics and interactive Techniques SIGGRAPH*, New York, NY, 2007, 301-308.

[17] M. Filipponi & P.-Y. Jeannin, What makes a bedding plane favourable to karstification? - The role of the primary rock permeability. *Proceeding of the 4th European Speleological Congress*, Vercors, France, 2008 32-37.

[18] A. N. Palmer, *Cave Geology* (Dayton, OH: Cave Books, 2007).

[19] G. S. Miller, The definition and rendering of terrain maps. *Proceedings of the 13th Annual Conference on Computer Graphics and interactive Techniques SIGGRAPH*, New York, NY, 1986, 39-48.

[20] M. A. DeLoura, *Game Programming Gems* (Rockland, MA: Charles River Media, 2000).

[21] F. K. Musgrave, C. E. Kolb, & R. S. Mace, The synthesis and rendering of eroded fractal terrains. *Proceedings of the 16th annual conference on Computer graphics and interactive techniques SIGGRAPH*, New York, NY, 1989, 41–50.