TERRAIN SYNTHESIS BY-EXAMPLE

John Brosz
Email: brosz@cpsc.ucalgary.ca

Faramarz F. Samavati
Email: samavati@cpsc.ucalgary.ca

Mario Costa Sousa
Email: mario@cpsc.ucalgary.ca

University of Calgary
2500 University Drive NW, Calgary, Canada

Abstract: Synthesizing terrain or adding detail to terrains manually is a long and tedious process. With procedural synthesis methods this process is faster but more difficult to control. This paper presents a new technique of terrain synthesis that uses an existing terrain to synthesize new terrain. To do this we use multi-resolution analysis to extract the high-resolution details from existing models and apply them to increase the resolution of terrain. Our synthesized terrains are more heterogeneous than procedural results, are superior to terrains created by texture transfer, and retain the large-scale characteristics of the original terrain.

1 INTRODUCTION

Terrain synthesis is the process of creating the artificial terrains used in games, movies, and simulations. Artificial terrains are necessary whenever terrain information is not available, lacks resolution and detail, or must meet specific feature criteria. Creating artificial terrain is not easy as large, high-resolution terrains often involve millions of data points.

There are two major applications of terrain synthesis. The first is in creating a synthetic terrain guided by a user who provides a rough outline of the desired characteristics. In this case the goal is to retain the attributes of the given outline, while adding realistic features to fill out the terrain. The second application is in increasing the resolution of an existing terrain for close-up viewing. Although any type of surface interpolation can increase resolution, it will also make the terrain look overly smooth. Extra information in the form of deviation from the smooth surface is necessary to make the terrain interesting and realistic. Terrains synthesized for both applications should appear visually realistic, meet desired topology requirements, and be inexpensive in terms of modeling time (Roettger and Frick, 2002).

Presently in commercial systems like Bryce (Bryce, 2005) and Terragen (McLusky, 2005), artificial terrains are created through a mixture of procedural methods and user painting of height fields (Roettger and Frick, 2002). The procedural methods, usually based on fractal subdivision, are controlled by user specified parameters. These methods have major shortcomings in that they lack erosion features and create homogeneous terrains. To compensate for this, erosion processes can be simulated on these surfaces, adding missing erosion features. Unfortunately erosion simulation is slow and introduces many more parameters for the user to control.

In general, creating any new model, especially a detailed and complex model is difficult. However, if an example with similar attributes to what is being created can be used for inspiration, the task becomes easier. It is more natural for a user to specify that they want their terrain to look like a particular example, but with their specified features, rather than by providing a large number of parameters. Moreover, high resolution terrains are available to the public through the internet (U.S.G.S., 2005). The major goal of this work is to replace procedural and erosion synthesis parameters with use of an example terrain to make it easier to synthesize realistic, heterogeneous terrain.

1.1 System Overview

There are two different terrains that are used to create a terrain by example. The first we call the *base* terrain. This is the terrain that is either a rough estimate of the desired terrain's large-scale characteristics, or it is a real terrain where we would like to increase the resolution. The other is the *target* terrain. This is the terrain that displays the high frequency, small-scale characteristics that the user desires to have in their synthetic terrain. The target terrain must have a resolution greater than or equal to the desired resolution of the synthetic terrain in order to provide small-scale information of the correct resolution. Our goal is to extract the small-scale characteristics from the target terrain and apply them to the large-scale characteristics of the base terrain.

Our system has three major steps. The first step is to extract the details from the base terrain. The next step is to find the areas in the base and target terrains that best match one another and produce a map of these relationships. The last step is to re-organize the details to match the mapping and then add these details through subdivision, refining the base terrain.

To extracting these details and copying them between terrains, we use the subdivision and reverse subdivision filters created by Samavati and Bartels

(Samavati and Bartels, 1999). We also explore another method of extraction where we estimate the local fractal scale, using this estimate to create multi-fractal terrains by-example.

When copying the details from target terrain to the base terrain it is important that information coming from one type of feature, gets applied to that same type of feature. For example the peak of a mountain should get applied to another mountain peak. To achieve this mapping we have developed two methods. The first is interactive, relying upon the user to select matching areas. The other method is automatic, based on a texture synthesis technique.

1.2 Contributions & Organization

This work makes the following contributions: (1) use of multi-resolution analysis to extract and apply high-resolution terrain information, (2) two techniques for matching similar areas of the example terrain and the terrain being synthesized are presented, (3) adaptation and use of Image Quilting (Efros and Freeman, 2001) to automatic matching of terrain features, and (4) adaptation of Losasso and Hoppe's residuals measurement (Losasso and Hoppe, 2004) to estimate fractal scale when creating multi-fractals.

The remainder of this paper is organized as follows. Section 2 reviews work related to terrain synthesis. The data we use is described in Section 3. Section 4 discusses multi-resolution modeling and our reasons for using multi-resolution analysis. Section 5 presents the terrain by-example system. Section 6 discusses terrain rendering. Lastly, Section 7 shows results and Section 8 presents conclusions and future work.

2 BACKGROUND

Fournier et. al (Fournier et al., 1982) and Lewis (Lewis, 1987) developed fractal and general stochastic subdivision methods for creating synthetic terrains. These two techniques succeed in creating a variety of terrain-like objects. Although the terrains produced by these systems look interesting and impressive, there are problems with the created terrains. The first is that they lack the characteristics of erosion. The second is that it is difficult to control the fractal or stochastic functions to create a specific terrain. The last is that these terrains tend to be homogeneous (i.e., they are similarly rough, or smooth, over the entire terrain).

Miller (Miller, 1986) improved Fournier et. al's midpoint insertion technique by replacing linear interpolation with third order B-spline based subdivision. This removes artifacts associated with the original midpoint technique.

There exist several works where the resolution of existing terrain models is increased through estimating the fractal dimension of the terrain (Pumar,

1996) (Brivio and Marini, 1996) (Losasso and Hoppe, 2004). This estimate is used to provide fractal values to fill in the new data points. Our multi-fractal by-example technique differs from these works in that we estimate the size of the fractal displacement on the target terrain at high resolution, rather than by the displacements present in the base terrain.

Musgrave et. al (Musgrave et al., 1989), Kelley (Kelley et al., 1988), Benes and Forsbach (Benes and Forsbach, 2001), and Nagashima (Nagashima, 1997) are among the researchers that have developed erosion simulations to increase the realism of procedurally created terrains. Such systems reproduce the effects of water, thermal, and other types of erosion.

However, there are problems with these processes. The first is that these methods usually take a large amount of time to simulate the erosion. The second problem is that these processes introduce a large number of new parameters (e.g., number of time steps, rainfall patterns, soil conditions, wind patterns, etc.) that must be accurate to achieve realistic and predictable results. These parameters are difficult to select accurately (Nagashima, 1997). Another problem is the evaluation of accuracy. All the proposed systems admit to being mostly empirical models that aim to capture the essential behavior of a few of the most noticeable erosion processes. Lastly, it has not been determined how erosion effects can be simulated in scenarios where the user wishes to increase the resolution of real terrain data. In this situation we do not want the erosion to eliminate existing features.

Musgrave in Ebert's book (Ebert, 1994) discusses a variation on fractal synthesis known as multi-fractals. This technique improves upon fractal terrains by varying the fractal scale over the terrain, resulting in heterogenous features. Unfortunately, automatic application of multifractals is difficult and can easily produce undesired effects as is shown in (Ebert, 1994).

Chiang et. al (Chiang et al., 2005) present an interactive approach to terrain synthesis. In their system the underlying shape of the terrain is created from simple geometric objects (e.g., prisms, cones, etc). These primitives are matched to a *terrain units* by cross-section, mountain ridge, or contour similarity. Terrain units are extracted from a database of terrain data that has been manually segmented to contain one mountain, ridge, or other feature (Chiang et al., 2005).

3 TERRAIN DATA

The terrain data we use is Digital Elevation Model (DEM) data. DEMs are composed of a regularly spaced, 2D array of floating point elevation values. The DEMs used in this work come from the United States Geological Survey (U.S.G.S., 2005). There are several different sets of elevation data available, we have used the two that were available for the most

areas within North America. The first set is Shuttle Radar Topology Mission (SRTM) data. This data is canopy based and is available in 90m and 30m resolutions (U.S.G.S., 2005). The other set is the National Elevation Dataset (NED). This is composed of bare ground readings from a variety of sources including aerial photographs and physical measurements (U.S.G.S., 2005). NED data comes in several resolutions including 30m, 10m, and 3m resolutions.

4 MULTI-RESOLUTION

Multi-resolution modeling is important for our synthesis technique since these techniques describe the differences between high and low-resolution versions of models. We interpret these differences as being characteristic of a terrain's small-scale features that we use to add small-scale features to the base terrain's large-scale features. In particular, we have chosen to use multi-resolution analysis (MRA) to extract and copy these small-scale features. MRA is available for curves and arbitrary topology models, however since DEMs can be treated as tensor product surfaces we use only the curve application.

MRA operations fulfill four necessary conditions for our system: (1) the ability to coarsen terrain models, (2) separates and retains the high-frequency information lost in coarsening, (3) has capacity to refine terrain with high frequency information from other terrains, and (4) has the ability to perform these operations efficiently in both time and memory space.

We describe each row or column of terrain as a set of sequenced points: c_1^k, c_2^k, ..., c_n^k, and present these points in the column $C^k = [c_1^k, c_2^k, ..., c_n^k]^T$. Subdivision produces a refined, higher resolution set of n' points, C^{k+1}, through multiplication by a matrix:

$$C^{k+1} = P^k C^k \qquad (1)$$

where P^k is a $n' \times n$ subdivision matrix that doubles the resolution of the curve. To reverse this process we start with a fine set of n points, C^k. A coarse set of n' points, C^{k-1}, is produced by application of a decomposition matrix:

$$C^{k-1} = A^k C^k \qquad (2)$$

where A^k is the $n' \times n$ matrix that halves the resolution. Since C^{k-1} has fewer points, some high frequency information (details) about the curve C^k have been lost. These $n - n'$ details, D^{k-1}, are captured by applying another matrix on the original points:

$$D^{k-1} = B^k C^k \qquad (3)$$

where B^k is a $(n-n') \times n$ detail extraction matrix that is dependent on A. These details, D^{k-1}, are vectors that describe the differences between C^k and PC^{k-1}. The last component of MRA, Q, is a $n \times (n'-n)$ matrix that adds the contributions of the details, D, when refining the coarse points into our original data set.

So, if we decompose a curve using equations 2 and 3, we can reconstruct the initial curve exactly using:

$$C^k = P^k C^{k-1} + Q^k D^{k-1}. \qquad (4)$$

The matrices A^k and B^k should produce details that are as small as possible. If the details are small, then C^k and PC^{k-1} are almost the same, loosely indicating that C^{k-1} is a good approximation of C^k. We also desire small details since this ensures that most of the contribution to the final terrain is from the base terrain, rather than the target terrain. This preserves the large-scale features of the base terrain.

We chose to use the local Chaikin matrices derived by Samavati and Bartels (Samavati and Bartels, 2004). These matrices are derived from a local least squares analysis of Chaikin subdivision, based on third order B-splines. We chose Chaikin over higher order methods due to its simplicity and compactness and over lower order methods, such as Haar or Faber, because it yields surfaces with greater continuity and is better suited for smooth surfaces (Samavati and Bartels, 1999). Chaikin also underlies the fractal synthesis results obtained by Miller (Miller, 1986). Readers interested in the values of the matrices should refer to Samavati and Bartels (Samavati and Bartels, 2004).

Although many multi-resolution methods exist, we are only concerned with surface simplification methods that create coarse approximations of models. The goal of these techniques is to create coarse models that retain the shape of the original model. We can also limit our discussion to methods that can exactly reconstruct the original representation from the coarsened version. For exact reconstruction some form of details are necessary and the construction of the details is important to our technique.

One such method is Image Pyramids. This techniques use a hierarchy of images formed with the original image as the bottom level, and at each higher level the image has half the resolution of the preceding image. The reduction in resolution is obtained with a low-pass filter and sub-sampling. There are a variety of filters and MRA filters are obvious candidates. We chose not to use filters unique to Image Pyramids because they oversample the image or only approximately reconstruct the original image (Simoncelli et al., 1992).

Triangle decimation (Schroeder et al., 1992) removes triangles that make little contribution to the model. Surface Simplification using Quadrics (Garland and Heckbert, 1997), Progressive Meshes (Hoppe, 1996), and other techniques use edge contraction to remove edges, coarsening the model. These techniques are inappropriate for our purposes as they do not describe a uniform change in resolution over the entire object and because they produce triangulated irregular networks (TINs) rather than DEMs.

Furthermore, the behavior and effects of the details generated by these methods have not been researched to the same extent that MRA details have.

5   TERRAIN BY-EXAMPLE

To reduce the resolution and extract the details of a terrain, we apply the reverse subdivision matrices A and B (Equations 2 and 3) to each column of height values. This gives us an array of coarse points with almost half the height of the original as well as a set of details. We refer to this first set of details as the *column details*. Then we use the filters on each row, leaving us with a coarsened terrain of half the resolution and another set of details, the *row details*.

To refine our base terrain to double its resolution we use the subdivision matrices P and Q (Equation 4) on each row of the array with the target's row details generating an array with double the width. Then P and Q are used on each column in conjunction with captured column details, resulting in a terrain with double the resolution.

First consider the case where the target terrain is exactly double the resolution of the base terrain. Our first step is to extract the high-frequency noise from the target with reverse subdivision. The resulting row and column details are our high-frequency noise. The simplest method of using these extracted details with our base is to simply apply them via subdivision to the base terrain. Unfortunately, simply applying these details in the order they were extracted leads to problems. As shown in Figure 1, the details extracted from a terrain correspond to features of that terrain. If we apply these details in the same order that we extracted them, we could apply the wrong details to the features in the base terrain. E.g., we might apply the details from a mountain in the target to a plateau in the base terrain. Section 5.2 describes two methods of matching features in the base to features in the target terrain.

Once a mapping has been specified we rearrange the target's row and column details into the order determined by the mapping. Subdivision is used to apply the details to the rows and then the columns of the base, doubling its resolution, producing our synthesized result. This process is referred to as applying one level of details or as performing one iteration of subdivision.

Now consider the case where the target terrain has a much higher resolution than the base terrain. In this case we perform reverse subdivision to the target several several times, until the target has approximately the same resolution as the base. As before our next step is to perform matching between the base and target terrains, and then subdivide the base terrain using the last set of details extracted from the target. We use the last set since these are the details that will describe the difference between the base's current reso-
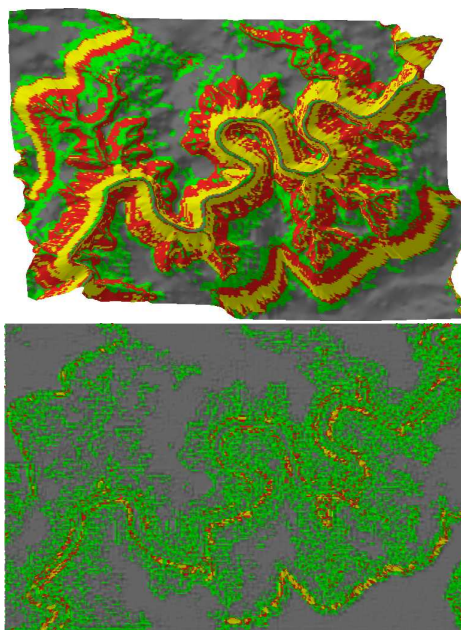

Figure 1: Top is a valley in Nevada, USA. Bottom is the result of extracting one level of details and applying them to a flat surface. The features in the terrain clearly correspond to features in the details. Colored by slope steepness.

lution and double this resolution. We can then repeat this matching and subdividing until we have used up all the details in reverse order and our synthesized terrain has the same resolution as the original target.

5.1   Multi-Fractals By-Example

Another way of creating terrain by-example is with multi-fractals. This technique has drawbacks but provides another method of creating terrains by-example.

The difference between fractal terrains and multi-fractal terrains is that the fractal scale changes over the terrain producing terrains that are less homogeneous. To produce multi-fractals by example we need to capture the fractal scale from the target terrain. To do this we must determine an estimate of the fractal scale and use this parameter on the feature-mapped areas of the base terrain. Losasso and Hoppe (Losasso and Hoppe, 2004) have obtained fractal scale estimates for entire terrains by using the variance of *residuals*. The residuals used are essentially the details resulting from linear interpolation.

In our work we estimate the fractal scale in a similar manner. We use Miller's Chaikin based terrain fractals (Miller, 1986), consequently the fractal scale becomes the variance of the displacements created by the Chaikin details. To implement our by-example multi-fractals we use this estimate of fractal scale over small areas of the target and then use these to decide the fractal scale on mapped areas of the base terrain.

With multi-fractals by-example, instead of creating the detail quilts, we estimate residuals over the target

terrain's selected blocks (the small area over which the fractal scale is calculated) and store the fractal scales to be used on the base terrain. The mapping procedure (matching similar pieces of terrain) is still performed, making this algorithm execute at approximately the same speed as MRA by-example.

Multi-fractals by-example have a major drawbacks in applying small-characteristic features. To understand why we must examine the generation of the fractal displacement. For our multi-fractals we are taking the variance of the details and using this to scale Gaussian noise. By doing this we are discarding an accurate and real distribution of details from the target, and replacing this with a Gaussian distribution. In essence we are discarding most of the noise we have gone to the trouble of extracting from the target. This is one of the reasons why multi-fractal by-example results tend to resemble the target less than MRA by-example results.

5.2 Mapping Terrain Features

The goal of this section is to find a mapping between similar areas and features on the target and base terrains. We have made use of two existing techniques for handling this problem. The first is interactive, relying on human interpretation to solve the matching problem. The other technique is automatic and based on Texture Quilting (Efros and Freeman, 2001). We chose this technique because of its speed and quality.

5.2.1 Interactive Mapping

The simplest technique of solving the mapping problem is to leave it up to the user. We replace the guessing at assorted parameters required for fractal synthesis or erosion simulation with a more intuitive mechanism. This mechanism is having the user perceive similar areas in the base and target terrains.

In this system the user selects a section of the target terrain and then a section of the same size of the base terrain. These sections may have any shape; we have chosen to use rectangles in our system for speed of calculation and ease of implementation. If mapped-to areas on the base terrain already posses details we replace the existing details. The user can also specify the number of levels of details to copy. The number of levels to copy is dependent on the desired resolution but in our tests we have found that copying three levels of detail provides good results. An example of this process can be seen in Figures 2 and 3.

Performing the MRA is very fast (Samavati and Bartels, 1999), allowing users to easily and quickly see the results of their specified mapping. This interactive method can also be used to edit the details of an existing terrain. By specifying a new mapping of an area, the base terrain's details are replaced, changing the small-scale characteristics of the selected area.

This method also easily allows more than one target terrain to be used to refine a single base terrain.
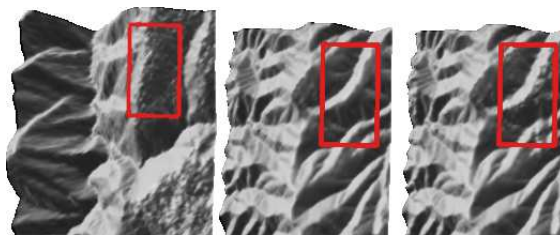

Figure 2: The interactive system copied the bumpy details from the highlighted area of the target (left) changing a smooth area of the base terrain (middle) into a more rough area (right).
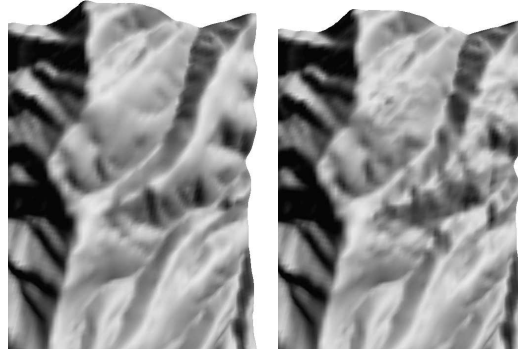

Figure 3: Close-up of the terrain in Figure 2, before (left) and after (right) synthesis. The red lines have been removed to allow examination of the borders.

5.2.2 Automatic Mapping: Terrain Quilting

In order to reduce the amount of user interaction we have also created an automatic mapping system. This automatic system is based on the Image Quilting algorithm (Efros and Freeman, 2001) used for texture synthesis and texture transfer.

The first step is to break the base terrain into blocks. We do this by dividing the base terrain into a grid of overlapping square blocks. By choosing a smaller block sizes, smaller features are matched and better mappings can be produced; however, this also increases the time required for the algorithm to run. Our default setting uses blocks that are approximately one-fifteenth of the average of the base terrain's width and length, however, the block size is adjustable by the user. For the best possible results the block size should be determined by the resolution of the small-scale characteristics the user desires to capture.

For each block in the base terrain we attempt to find a similar block in the target terrain. This is done by testing the similarity of some number of randomly selected blocks in the target. The number of candidate blocks we test is the smaller of forty or the number of blocks that fit equally into the target. This gives a good trade off between speed and quality, however if poor matches occur this setting is user adjustable.

In the Image Quilting algorithm (Efros and Freeman, 2001), the difference (or conversely the similarity) between two blocks is described by the sum of squares between the pixel brightness values. For terrains, the values we are comparing are elevations rather than pixel brightness so we more concerned with features (e.g., a mountain peak in the base terrain should be mapped to a mountain peak in the target terrain). Consequently, we need to examine the shape portrayed by the block, rather than the absolute height values of the block.

The metric we created to calculate the difference in shape between two blocks, i from the base, and j from the target, where we have indexed the height values in the blocks with a single index, is:

$$ e_{i,j} = \sum_{k=1}^{n} ((i_k - m_i) - (j_k - m_j))^2 $$

where n is the number of height values per block, i_k is the k^{th} height value of block i, and m_i is the mean height value of block i. Unlike Image Quilting, when calculating the difference between blocks we ignore the blending between blocks (since we are copying details, not pieces of terrain) so this does not contribute to the difference calculation.

5.2.3 Constructing the Detail Quilts

Once we have found the best match for each base terrain block, we construct a quilt of the target's details. This is done for both the row and the column details. We quilt the details, rather than the actual height values, because we want a new synthesized terrain, not a terrain that was created through mixing and matching pieces of the target. A mix and match terrain would suffer block blending issues and could alter large-scale characteristics of the base terrain. Instead, the system quilts the details so that the resulting terrain contains merely the addition of the details from the target. In Image Quilting a minimum error boundary cut is used to blend overlapping areas. In our work we use a linear blend of the details between the blocks. We do this because details are small in respect to the curve presented by the low-resolution data for the reason described in Section 4.

6 TERRAIN RENDERING

Rendering in our system is not just a matter of simply displaying terrains. Rather it is a tool that can allow us to more easily evaluate, compare, and contrast synthetic terrains. This is especially important when presenting synthesis results in one or two images. We use four rendering techniques that provide easier discrimination between and evaluation of terrains.

Our standard rendering method uses Gouraud shading on grey colored triangles to maximize contrast and make terrain evaluation possible when the images are printed in black and white.

The second rendering style uses Gouraud shading but with triangles colored based on elevation. From lowest elevation to highest triangles are colored brown, green, red, orange, yellow, and white. This coloring makes the terrain easier to understand, especially when the terrain is viewed from the top.

We also desired to make it easier to determine how smooth (or rough) a terrain is. To do this we color based on slope steepness. This is the rate of change in elevation making it a good measure to indicate the smoothness of a terrain (Shary et al., 2002). In our slope steepness figures, from most steep to least steep, the triangles are colored yellow, red, green, or grey. We have found this gives good insight into the amount of heterogeneity of rendered terrains and can emphasize features such as the riverbed in Figure 1.

Our last technique uses a non-photorealistic rendering technique to place small ink strokes. This rendering technique uses silhouette stability edges with the direction of the edge changed to point in the tangent direction as described in Brosz's work (Brosz, 2005). For evaluating terrains, this rendering technique emphasizes the differences between rough and smooth areas of synthesized terrains.
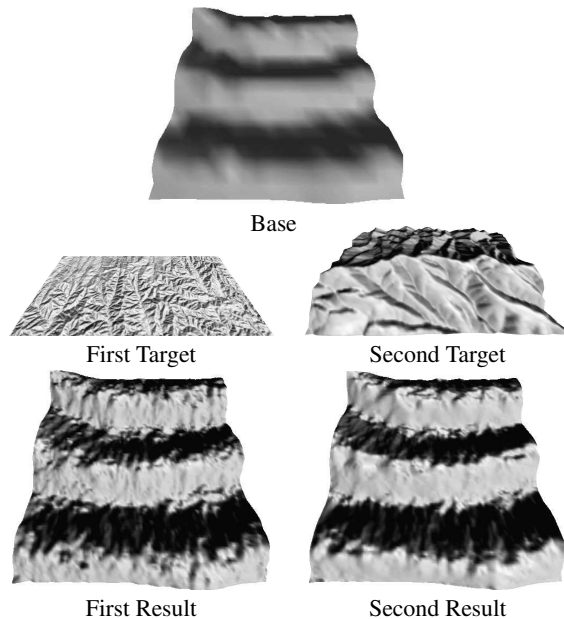

Figure 4: The first trial. The base is a user created 20×20 terrain. The results are 146×146, produced by copying three levels of details from their respective target terrains.

7 RESULTS AND DISCUSSION

There are several criteria for the success of our system. The most important is that the synthesized terrain should appear more detailed than if the base was subdivided without details. Our next expecta-

tion is that the terrain gains detail in a manner that is more predictable and realistic than current procedural methods. A related point is that we expect the synthesized terrain to show small-scale characteristics similar to those of the target. Another important goal for the success of the system is in the simplicity of the user interaction.

7.1 Providing a Control Synthesis

To compare the results of the terrain by-example system to the results of a fractal-based synthesis, we implemented fBm subdivision as described by Miller (Miller, 1986) with Voss' successive random additions (Peitgen and Saupe, 1988) to further reduce fractal artifacts. We estimate the initial fractal scale based on our residual calculation. H, the parameter controlling the change in fractal scale (Fournier et al., 1982), is 0.85 as we found this value produces good results.

We have not considered erosion techniques due to two problems. The first is that they involve many parameters that must be provided. These parameters depend on the type of terrain as well as many environmental conditions. The other is that these erosion simulations are likely to remove or significantly alter previously existing features of the base terrain that are real features and desired in the resulting model.

7.2 Experimentation

In this subsection we present two trials using Chaikin MRA. In the first trial we synthesized a terrain from a small, low resolution, hand-made terrain. The second trial compares a real, high-resolution terrain with a synthesized terrain that uses low resolution data from the same area as its base terrain. In both trials we used the default block size. Timings were performed on an AMD 2600XP+ with 1.5GB RAM and an ATI 9800 128MB video card. The base and target terrains used in our first trial, as well as the terrains resulting from our synthesis, are shown in Figure 4. The first target is a 30m NED from Kansas, USA, the second is a 30m NED from Utah, USA. The base terrain's resolution was increased from approximately 240m to 30m resolution. Figure 5 compares the result of fBm fBm synthesis to a result achieved with terrain by-example synthesis. The terrain by-example synthesis in this trial took approximately 3 seconds while the fBm synthesis required less than 1 second.

In comparing our two synthesized results, we can see that changing the target significantly changes the synthesis result. The first result has more jagged features, similar to the jagged hills in the Kansas target whereas the second result is smoother, more like the Utah target. Looking at our results, especially the sides of the hills, we can see that they have interesting bumps added while also appearing smooth, much like the target terrains. In contrast, the fBm result is very rough all over. The fBm synthesis, because it changes

the height at random is uniformly noisey, whereas the by-example results are more coherent.

Figure 5: Trial 1 comparison between by-example (top) and fBm synthesis (bottom). Notice the homogeneous nature fBm result when compared to the by-example result.

Our second trial uses a 30m NED of an area within the Utah salt flats as its base terrain. We chose a 10m NED from another portion of the Utah salt flats as our target (shown in Figure 6). The base and the results of by-example and fBm synthesis are shown in Figure 7. Figure 8 provides comparison between 10m NED data of the same area and the two synthesized 7.5m results. The by-example result (2794×394), obtained from two levels of copied details, was created in 8 seconds; the fBm result required 2.5 seconds.

In Figure 7 our synthesized terrain shows very little introduced roughness. This is due to the smoothed nature of the NED as well as the smooth nature of the salt flats. Close inspection reveals small irregularities added to several of the mountain ridges and a few small bumps in the flat area. In the fBm terrain we see that the introduced noise obscures some of the streambeds in the middle of the terrain. The most noticeable difference is that the fBm synthesis adds a large number of undesirable bumps to the flat area.

This trial, with the mountainous terrain at one end and the very flat terrain at the other, emphasizes the drawback of fBm synthesis. The estimated fractal scale is the average residual from the entire terrain, resulting in terrain that is too rough at the flat end and too smooth at the other. Use of multi-fractals could fix this, but even then the scale is still estimated over some area and, if that area has heterogeneous features, the scale can again be wrong. MRA however, is based on wavelets that are able to capture and represent these different scales effectively.

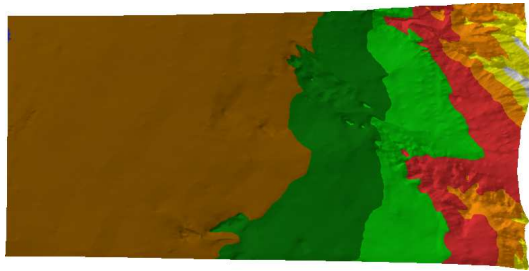It is important to note that the flats cannot be syn-

Figure 6: The target terrain used in the second trial, colored by elevation. This is a 10m NED from the Utah Salt Flats.
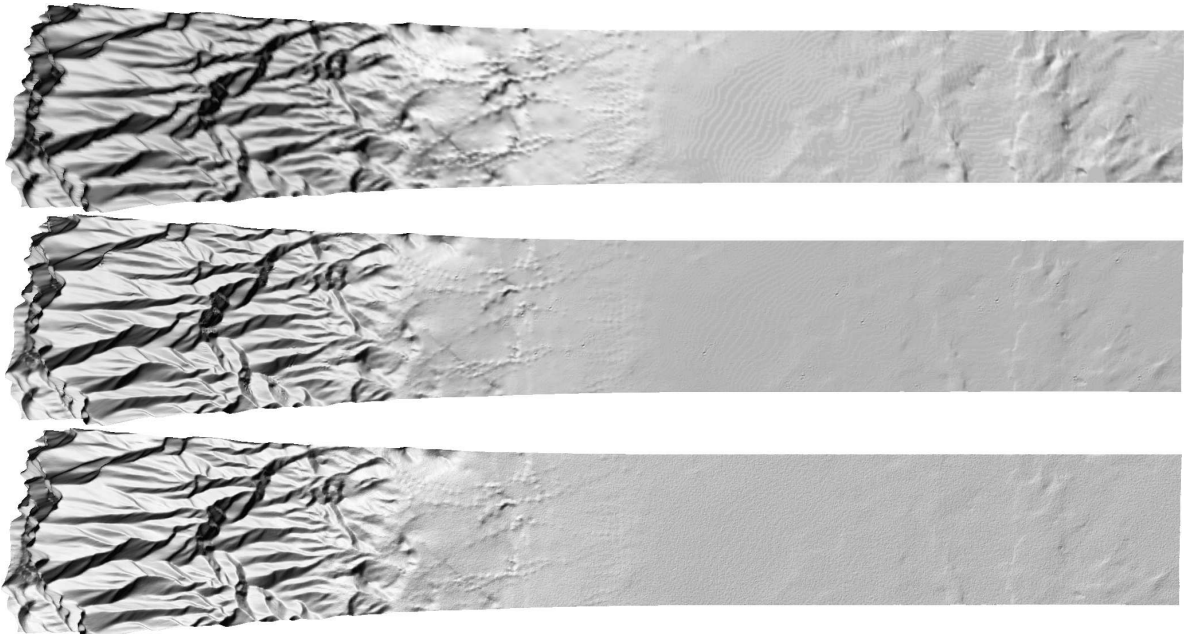

Figure 7: The second trial. The top image is the base terrain (30m NED data from the Utah Salt Flats), the middle is the Chaikin, terrain by-example result, and the bottom is the result of fBm synthesis.
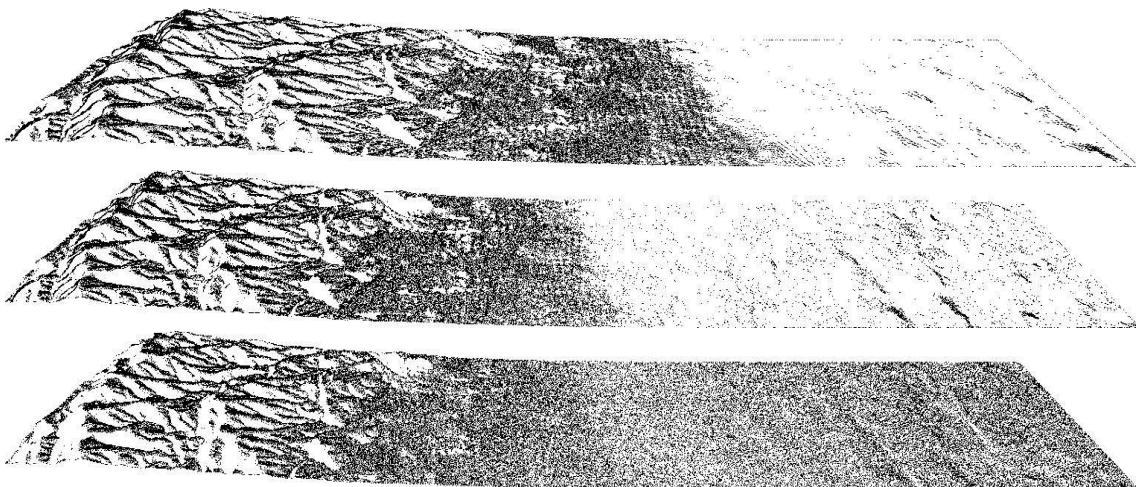

Figure 8: Non-photorealistically rendered comparison between the real, 10m resolution NED (top), our second trial's 7.5m by-example result (middle), and the 7.5m result of fBm synthesis (bottom). Notice the roughness of the salt flat fBm result.

thesized in a better fashion by simply adding erosion to the fBm synthesis. Thermal weathering erosion (Musgrave et al., 1989) may somewhat reduce the size of the fractal bumps, but not completely because the bumps are small and have very little slope. Hydraulic erosion techniques can smooth these bumps out but are designed to, and will, introduce new rivers and valleys, destroying the features of the base terrain.

7.3 Chaikin vs Multi-Fractals

In Section 5 we have discussed two different techniques of extracting and applying details. Chaikin MRA, is useful because it is based on third order B-Splines that effectively capture the smooth underlying nature of terrains. Multi-fractals by-example provide a close link to existing procedural methods. To compare these techniques we have repeated the first trial using multi-fractals. Figure 9 shows the results.

In Figure 9 it is clear that the multi-fractals are not as successful in capturing the character of the target terrain. It is not clear that the left multi-fractal result came from the ridge-filled Kansas target while the right came from Utah target. The Chaikin MRA results show these differences.
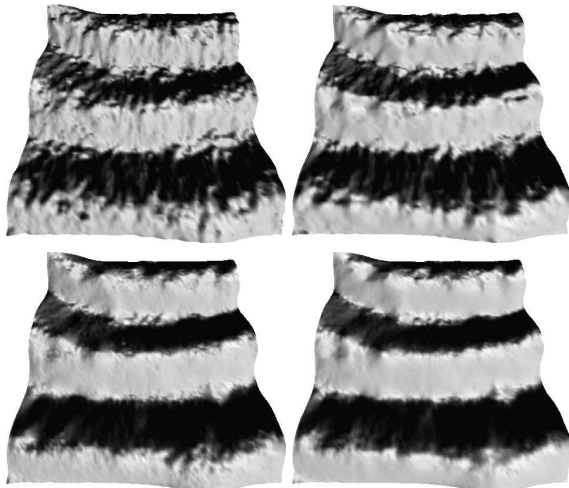

Figure 9: Repeat of Trial 1 (Figure 4) with Chaikin MRA (top) and multi-fractals by-example (bottom). In the left column are the results that used Kansas target, the right column presents the results of the Utah target.

7.4 Texture Transfer Compared

An important topic is why current texture transfer techniques are not suitable for synthesizing terrain. In some senses they are suitable. Can a target terrain be broken into pieces and then sewn together to form a realistic terrain or, as in Image Analogies, two different resolutions of the target terrain could provide a filter to increase the resolution of the base terrain?

The difference between our method and texture transfer approaches is that we do not copy pieces of terrain; instead we copy details. This ensures that our result is composed mostly of the features of the base terrain. Texture transfer rearranges pieces of the target to fit the base, making these methods unsuitable for increasing the resolution of existing terrain. By using details, since they are such small contributions to the result, we also avoid the blending artifacts that appear in texture transfer methods.

To test this we compare a terrain synthesized by-example with results of two texture transfer systems: Image Quilting (Efros and Freeman, 2001) implemented by Robert Burke (Burke, 2005) and the Image Analogies system (Hertzmann et al., 2001). Figure 10 shows the base, target, and our synthesis.

To perform the Image Quilting texture transfer, we convert the DEM to a height encoded image (i.e., an image where the highest areas are white and the lowest are black). Next, we resample the 30m base to a 10m resolution. This is akin to subdividing without adding details, producing a smooth, high-resolution terrain. We then apply texture transfer between the target and the resampled base producing the result shown in Figure 11. It is clear that directly using the original terrain has introduced blending problems, failing to produce a realistic terrain.

For the Image Analogies texture transfer we used height encoded images of the 30m and 10m NED target terrain as the learning pair. As both of these images had to be the same size, we resampled the 30m image to achieve 10m resolution. Similarly the 30m base image was also resampled. The chossen matching method was Hertzmann et. al's recommended approximate nearest neighbor search. Figure 11 presents our best result with this system. More results of our testing can be seen in Brosz's work (Brosz, 2005). The result of the Image Analogies system does not reproduce the flat area of the upper left base terrain and generally produces terrains with a speckled appearance. The hillsides have new ridges and lines, giving an interesting but unwelcome terraced look. It does not appear to be a bad result unless one considers that neither the base nor the target has these features. Both terrains have ridges but these ridges are mostly smooth, not noisy like synthesis result. Since neither the base nor the target terrains have speckles, these must be artifacts created by blending problems when placing pixels.

The results of two texture transfer systems show that they have problems in maintaining flat terrain due to the problems of blending images (as opposed to blending details). It is also due to the fact that they are designed to preserve the source image (the target terrain) rather than to keep the features of the target image (the base terrain). The Image Analogies system performed better than the Image Quilting system but overall we feel it is clear that terrain by-example is superior for synthesizing terrain.
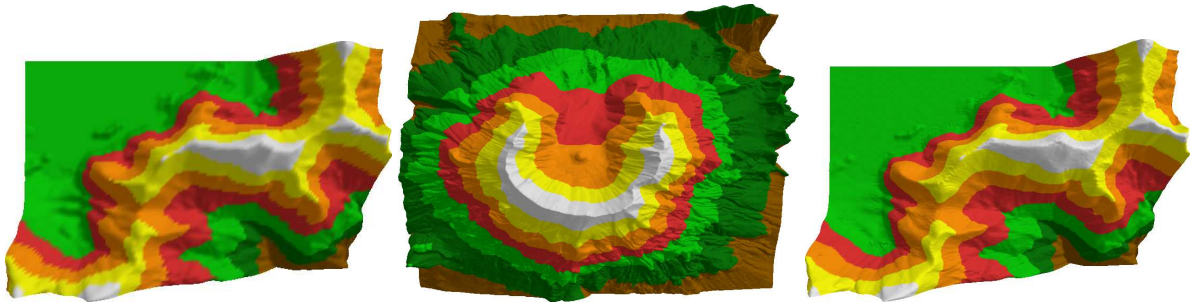
Figure 10: The left image is of the base terrain (30m NED of Harmony Flats, Washtingon, USA), the middle is the target (10m NED of Mount St. Helens, Washington, USA) and the bottom is the by-example synthesis result.
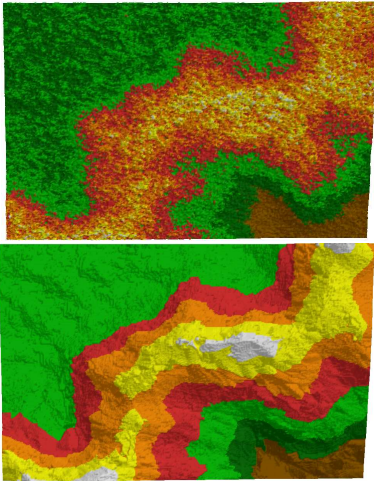

Figure 11: Top is the result of the Image Quilting algorithm using settings of 3 iterations, initial block size of 16, 17% overlap, and 50 candidates. Bottom is a result of the Image Analogies algorithm using steered filters and the default settings. Both are rendered with elevation coloring.

8 CONCLUSIONS

In this work a pipeline to synthesize terrains based on an example terrain has been presented. This system simplifies user interaction, replacing procedural and erosion parameters with selection of a target terrain. The resulting terrains show influence from the target terrain and have heterogeneous features. Lastly the synthesis is completed quickly, achieving synthesis in seconds. The availability of high resolution terrain data for use as targets makes this system possible.

In examining our criteria for success, we can see that our synthesised terrains are more detailed than the base terrain. Moreover, the terrains gain details that seem realistic. This is shown in all the trials, but particularly in Trial 2 where the by-example result maintained the flatness of the salt flats. The interactive system copies details from the target to the base in a predictable manner. The automatic system also results in terrains that show characteristics of the target terrain (Figure 4). The by-example results appear similar to even if they do not exactly match the real, high resolution data of the same area. Our system simplifies user interaction, producing user-specified terrain without input of a plethora of parameters. By-example synthesis can create completely synthetic terrains as well as add levels of detail to existing terrain.

Multi-resolution analysis provides an efficient and effective method for capturing and copying the character of a terrain. Multi-fractals present another method of extracting details from the target terrain, however this technique is unable to capture a target's character as well as Chaikin MRA.

Both the interactive and the automatic methods provide good starting points for feature matching in this pipeline. The techniques work well for most terrains, however, both methods have drawbacks. The interactive method requires more input from the user. For large terrains this could become time consuming and tedious. However, this system offers the potential for editing of existing details in terrains that proves an interesting and potentially useful feature.

The automatic feature mapping system has several limitations. Our random block matching technique, adapted from Image Quilting, is not ideal because it does not guarantee that the best matching block in the target terrain is found. An exhaustive search or variation of simulated annealing could ensure better matching. From our examination of the matches made by our system it seems that the current system works well without modification. Other techniques for comparing shapes could increase the effectiveness or speed of automatic matching.

There is also the issue of how the block size should be chosen. Small block sizes can achieve more accurate matches, but when the block size is too small, neighboring information is lost, causing bad matches. The values we have used produce qualitatively good results, however, a more quantitative examination would be useful. Adding a contribution from the neighboring area to the error metric might also alleviate the small block problem.

8.1 Future Work

Introducing a painting interface to specify specific mappings could increase the functionality of the by-

example system. This could be similar to the paint by numbers used with Image Analogies (Hertzmann et al., 2001). This could make it easier for users to indicate mappings and to specify different target terrains. E.g., half a base terrain could be painted yellow, with yellow mapped to a desert target, and the rest could be mapped to a mountainous target.

There exist methods (Pumar, 1996), (Brivio and Marini, 1996) that use estimates of fractal dimension (opposed to our estimate of fractal scale) to increase the resolution of terrains. It is possible that these alternative estimates may improve the quality of the terrain synthesized with multi-fractals by-example. It is our intuition however, that the poor performance of multi-fractals is more greatly due to the randomness of displacements as opposed to the size of those displacements. Consequently, we are not optimistic of improving the multi-fractal by-example technique.

Triangulated Irregular Networks (TINs) abandon the regular grid format of DEMs, reducing the number of data points necessary to represent terrains. TINs are able to better represent sharp features, such as ridges, by aligning triangle edges with feature edges (DEM's edges seldom coincide with features). A more sophisticated shape analysis technique in conjunction with the details extracted with progressive meshes (Hoppe, 1996), arbitrary topology MRA, or another multi-resolution technique could create a TIN by-example system.

REFERENCES

Benes, B. and Forsbach, R. (2001). Layered data representation for visual simulation of terrain erosion. In *Proceedings of SCCG '01*, pages 80–86.

Brivio, P. A. and Marini, D. (1996). A fractal method for digital elevation model construction and its application to a mountain region. *Computer Graphics Forum*, 12(5):297–309.

Brosz, J. (2005). Terrain modeling by example. Master's thesis, University of Calgary, Department of Computer Science.

Bryce (2005). *Bryce 5.0 User Manual*. Daz Software. http://bryce.daz3d.com/Bryce5_Manual_DAZ.pdf.

Burke, R. (2005). Image quilting, texture transfer, and wang tiling implementation. http://www.heroicsalmonleap.net/mle/wang/.

Chiang, M., Huang, J., Tai, W., Liu, C., and Chang, C. (2005). Terrain synthesis: An interactive approach. In *International Workshop on Advanced Image Tech.*

Ebert, D. S., editor (1994). *Texturing and Modeling: A Procedural Approach*. AP Professional.

Efros, A. A. and Freeman, W. T. (2001). Image quilting for texture synthesis and transfer. In *Proceedings of SIGGRAPH '01*, pages 341–346.

Fournier, A., Fussell, D., and Carpenter, L. (1982). Computer rendering of stochastic models. *Commun. ACM*, 25(6):371–384.

Garland, M. and Heckbert, P. S. (1997). Surface simplification using quadric error metrics. In *Proceedings of SIGGRAPH '97*, pages 209–216.

Hertzmann, A., Jacobs, C. E., Oliver, N., Curless, B., and Salesin, D. H. (2001). Image analogies. In *Proceedings of SIGGRAPH '01*, pages 327–340.

Hoppe, H. (1996). Progressive meshes. In *Proceedings of SIGGRAPH '96*, pages 99–108.

Kelley, A., Malin, M., and Nielson, G. (1988). Terrain simulation using a model of stream erosion. In *Proceedings of SIGGRAPH '88*, pages 263–268.

Lewis, J. P. (1987). Generalized stochastic subdivision. *ACM Transactions on Graphics*, 6(3):167–190.

Losasso, F. and Hoppe, H. (2004). Geometry clipmaps: terrain rendering using nested regular grids. *ACM Transactions on Graphics*, 23:769–776.

McLusky, J. (2005). Terrain dialog description. http://www.planetside.co.uk/terragen/guide/dlg_terrain.html.

Miller, G. S. P. (1986). The definition and rendering of terrain maps. In *Proceedings of SIGGRAPH '86*, pages 39–48.

Musgrave, F., Kolb, C., and Mace, R. (1989). The synthesis and rendering of eroded fractal terrains. In *Proceedings of SIGGRAPH '89*, pages 41–50.

Nagashima, K. (1997). Computer generation of eroded valley and mountain terrain. *The Visual Computer*, pages 456–464.

Peitgen, H.-O. and Saupe, D., editors (1988). *The Science of Fractal Images*. Springer-Verlag.

Pumar, M. A. (1996). Zooming of terrain imagery using fractal-based interpolation. *Computers and Graphics*, 20(1):171–176.

Roettger, S. and Frick, I. (2002). The terrain rendering pipeline. In *Proc. of East-West Vision '02*, pages 195–199.

Samavati, F. F. and Bartels, R. H. (1999). Multiresolution curve and surface representation by reversing subdivision rules. *Computer Graphics Forum*, 18(2):97–120.

Samavati, F. F. and Bartels, R. H. (2004). Local bspline wavelets. In *Proceedings of the International Workshop on Biometric Technologies 2004*.

Schroeder, W. J., Zarge, J. A., and Lorensen, W. E. (1992). Decimation of triangle meshes. In *Proceedings of SIGGRAPH '92*, pages 65–70.

Shary, P. A., Sharaya, L. S., and Mitusov, A. V. (2002). Fundamental quantitative methods of land surface analysis. *Geoderma*, 107(1-2):1–32.

Simoncelli, E. P., Freeman, W. T., Adelson, E. H., and Heeger, D. J. (1992). Shiftable multi-scale transforms. *IEEE Trans. on Informations Theory*, 38(2):587–607.

U.S.G.S. (2005). Seamless data distribution system. http://seamless.usgs.gov.