

# A Declarative Approach to Procedural Modeling of Virtual Worlds

R.M. Smelik<sup>a</sup>, T. Tutene<sup>b</sup>, K.J. de Kraker<sup>a</sup>, R. Bidarra<sup>b</sup>

<sup>a</sup>*Modelling, Simulation & Gaming Department, TNO, The Netherlands*

<sup>b</sup>*Computer Graphics Group, Delft University of Technology, The Netherlands*

---

## Abstract

With the ever increasing costs of manual content creation for virtual worlds, the potential of creating it automatically becomes too attractive to ignore. However, for most designers, traditional procedural content generation methods are complex and unintuitive to use, hard to control, and generated results are not easily integrated into a complete and consistent virtual world.

We introduce a novel *declarative modeling approach* that enables designers to concentrate on stating *what* they want to create instead of on describing *how* they should model it. It aims at reducing the complexity of virtual world modeling by combining the strengths of semantics-based modeling with manual and procedural approaches. This article describes two of its main contributions to procedural modeling of virtual worlds: *interactive procedural sketching* and *virtual world consistency maintenance*. We discuss how these techniques, integrated in our modeling framework SketchaWorld, build up to enable designers to create a complete 3D virtual world in minutes. Procedural sketching provides a fast and more intuitive way to model virtual worlds, by letting designers interactively sketch their virtual world using high-level terrain features, which are then procedurally expanded using a variety of integrated procedural methods. Consistency maintenance guarantees that the semantics of all terrain features is preserved throughout the modeling process. In particular, it automatically solves conflicts possibly emerging from interactions between terrain features.

We believe that these contributions together represent a significant step towards providing more user control and flexibility in procedural modeling of virtual worlds. It can therefore be expected that by further reducing its complexity, virtual world modeling will become accessible to an increasingly broad group of users.

## Keywords:

virtual worlds, declarative modeling, semantic modeling, consistency maintenance, procedural methods, procedural sketching

---

## 1. Introduction

3D virtual worlds have grown tremendously in size, detail and visual realism, their use being widespread far beyond entertainment games: they are found in games for training, movies, simulations, visualizations, online social spaces, etc. The downside of this, however, is that an ever increasing amount of content has to be created to fill up these worlds.

Current modeling systems offer designers total *control*, requiring the virtual world, with all its objects, to be modeled entirely by hand, on a low level of abstraction. With skill and dedication, designers can literally create any world exactly the way they want it, modeling every single aspect in detail. Eventually, the workload and repetitiveness of creating large worlds this way will become unbearable.

Virtual world models in manual modeling systems are also quite rigid: once completely constructed, they are hard to modify, e.g. major changes in the landscape may result in the designer effectively having to start from scratch. Furthermore, because manual modeling systems are complex and require extensive 3D modeling experience, the diversity of their user group is limited, typically excluding e.g. gaming enthusiasts creating new levels, and training instructors designing a tailored curriculum.

### 1.1. Procedural modeling

There is an urgent need for new modeling alternatives with both higher productivity and reduced complexity. Automatic procedural modeling seems an attractive alternative that promises such a productivity gain and a seemingly endless variation in content. It has been an active research topic for over thirty years, resulting in high-quality procedures for specific terrain features, such as landscapes [1, 2], rivers [3, 4, 5], plant models [6] and vegetation distribution [7], road networks [8], urban environments [9], and building facades [10, 11, 12].

However, as concluded in our recent survey [13], traditional procedural methods are not directly a suitable alternative to manual modeling. They typically lack any user control other than a set of, often unintuitive, input parameters, which not always have a clear, predictable effect on the output, the generated content. To some extent, using traditional procedural methods in a modeling system comes down to *trial and error*. In addition, because the runtime of these algorithms often is far from interactive, this process becomes even more cumbersome. Furthermore, each procedure generates one specific type of content. Fitting all this content together into a consistent virtual world involves a large amount of manual effort. These causes explain why procedural modeling is hardly used so far in mainstream content design. A hybrid solution combining the strengths of procedural methods with manual control would be ideal.

## 1.2. Extending traditional procedural methods

In recent years, the research direction in procedural modeling is shifting to user controllable and interactive procedures, thereby successfully addressing some of the downsides of traditional methods. Here we survey some noteworthy examples.

The lack of user control in procedural generation of elevation maps was the first issue addressed by several researchers. Their methods vary in interactivity and level of control, from coarse to fine-grained. Some of the proposed extensions provide a way to constrain the generation process in a non-interactive manner by new forms of user input. Stachniak and Stuerzlinger [14] propose a method that integrates constraints expressed as mask images. It employs a search algorithm that finds an acceptable set of deformation operations to apply to a random terrain in order to obtain a terrain that conforms to these constraints. Zhou et al. [15] describe a technique that generates terrain based on an example input height-map and a user line drawing that defines the occurrence of large-scale curved line features, such as mountain ridges. Features are extracted from the example height-map, matched to these curves and seamed in the resulting height-map. Saunders [16] proposes a method that synthesizes a height-map based on Digital Elevation Models (DEM) of real-world terrain. A user draws a 2D map of polygonal regions, each of which is marked to have a certain elevation profile. A height-map is instantiated using a genetic algorithm, which selects DEM data that matches the requested elevation profile. Kamal and Uddin [17] present a constrained mid-point displacement algorithm that creates a single mountain according to such properties as elevation and base spread. Belhadj [18] introduces a more general method where a set of known elevation values constrain the mid-point displacement process. Doran and Parberry [19] propose a different constraint-based approach using agents, each creating a specific landform (e.g. coastline, beach, mountain).

With the evolution of the GPU as a device for general purpose parallel processing, interactive user control in elevation map generation has become feasible. Schneider et al. [20] introduce a setup in which the user interactively edits the terrain by painting grayscale images, which are used as the base functions of their noise generator. Using an efficient GPU-based hydraulic erosion algorithm, Stava et al. [21] propose an interactive way for users to modify terrain using several types of hydraulic erosion. To provide users with more control over the exact appearance of mountain ranges, Galin et al. [22] introduce a sketch-based height-map generation method in which users sketch the silhouette and bounds of a mountain in a 3D interface, and the generator creates a matching mountain using noise propagation. Even more fine-grained control over the shape of mountains is provided by the interactive procedural brushing system introduced by de Carpentier and Bidarra [23]. These GPU-based procedural brushes allow users to interactively sculpt a terrain in 3D using several types of noise.

Extensions of traditional procedural methods are also proposed for other terrain features, for instance for interactively defining road networks used in city models. Chen et al. [24] propose interactive modeling of road networks using tensor fields that can create common road patterns (grid, radial, along

a boundary) and blend these in a plausible way. McCrae and Singh [25] present a method for converting strokes to 3D roads that are automatically fit in the terrain. Their system also creates junctions and viaducts for crossing roads. An A\* based road generation method proposed by Galin et al. [26] uses an elaborate cost function to encode the influence of terrain slope, water bodies and vegetation on the trajectory of the road.

As districts, blocks and parcels are defined by the city's road network, Kelly and McCabe [27] propose an interactive method to generate secondary roads and house blocks based on the primary roads the user manipulates. A similar system by de Villiers and Naicker [28] lets users create a road network and city blocks using sketch strokes, and interprets a set of sketch gestures that modify the properties of the city blocks (e.g. population size, function). Weber et al [29] present an interactive simulation system for cities growing over time, by expanding streets in the city's road network. A dynamic system that connects geometrical with behavioral modeling is proposed by Vanegas et al. [30]. Here, users paint behavioral variables like employment density, which automatically leads to changes in the population distribution and, thereby, the city geometry. Shape grammars are often employed for automatic creation of building facades (see e.g. [11]). However, defining a suitable shape grammar is complex and requires much experience. Addressing this, Lipp et al. [31] propose a shape grammar editing system, in which the effects of new rules are interactively visualized.

Although these extensions have clearly contributed to procedural modeling research, their downside is that they are primarily designed to generate one specific aspect of virtual worlds. Seldom has attention been given to the integration of separately generated terrain features into a complete virtual world. Early work by Amburn et al. [32] already formulated the problem of fitting roads with terrain: on a coarse level, the road follows the elevation profile of the terrain and on a fine level, the terrain must be modified to match locally with the road embankment profile. This specific integration problem was recently addressed by Bruneton and Neyret [33], who propose a shader-based system for real-time integration of Geographic Information Systems (GIS) vector features, such as road and rivers, into a DEM. They create a road profile texture based on footprint geometry, and integrate the profile by blending this texture with a height-map texture. The discussed work by Galin et al. [26] extends this by also removing any vegetation along the road. Although this research successfully addresses an important integration issue, managing the consistency of all terrain features as well as their relationships and dependencies remains an open problem.

A more generic approach to the problem of integrating terrain features and maintaining their relationships is to enrich their description with additional *semantic* information. Semantics-based modeling has been successfully applied to several fields, as surveyed in [34], including CAD/CAM [35], smart object behavior [36], and automatic layouts of interior scenes [37, 38]. However, the role of semantics in procedural modeling has, until now, been limited, which hinders the systematic integration of all the different procedural methods.

Commercial procedural modeling tools typically also focus on generating one specific feature, such as height-maps (e.g.

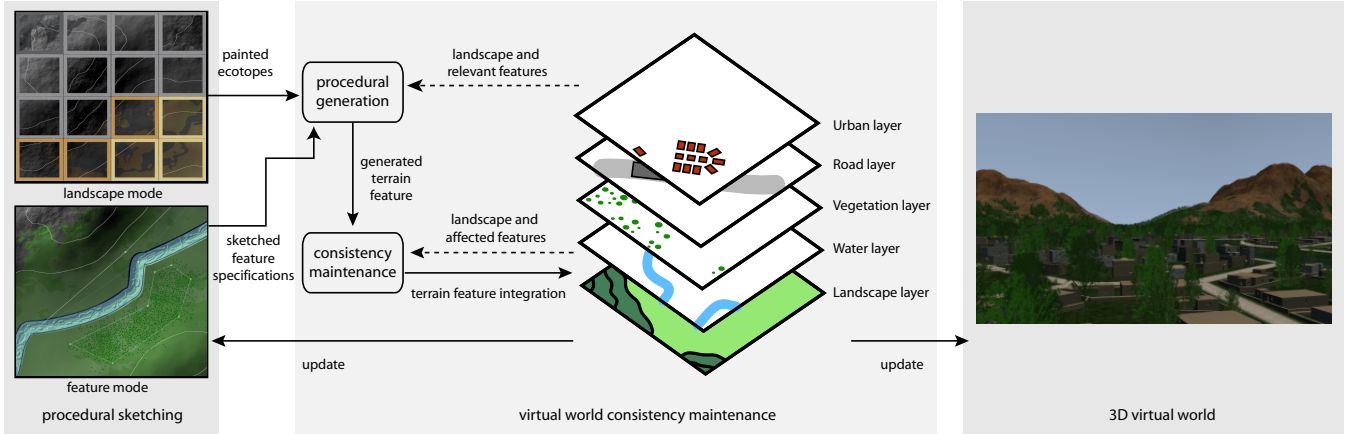


Figure 1: An overview of the workflow of the declarative modeling approach: using procedural sketching (Section 2), designers interactively create the virtual world, of which each feature is automatically generated, integrated and maintained within the virtual world model (Section 3). From this semantic model of the virtual world, other representations are derived, such as a 3D geometric model.

L3DT [39], and many others), vegetation (e.g. XFrog [40]) or city models (e.g. CityEngine [41]). A noteworthy exception is CityScape [42], which allows for a hybrid of manual and procedural modeling, although it provides a somewhat limited and narrowly focused set of procedural operations.

### 1.3. Declarative modeling of virtual worlds

As stated above, procedural modeling research is no longer solely focused on the generation of individual models of high quality. Researchers realize that more intuitive input, improved user control, and automatic integration of results are instrumental to the acceptance of procedural methods in mainstream virtual world development. However, to date, no research method or commercial tool provides such an integrated and flexible solution that allows designers to procedurally model a complete virtual world that matches their intent.

Our *declarative modeling approach* aims at contributing to these open issues, by combining the strengths of manual and procedural modeling, and providing a more productive and less complex workflow to model virtual worlds. Basically, this approach lets designers concentrate on *what* they want to create instead of on *how* they should model it. Designers state their intent using simple, high-level constructs. The declarative approach builds upon established research results on parameterized procedural generation, constraint solving and semantic modeling, in order to automatically translate these statements into a matching 3D virtual world. The consistency of this world is automatically maintained using a *semantically rich* model of all its features and their relations, analogous to the automatic maintenance of interior scenes based on object semantics [37].

This article discusses the two main contributions of our declarative modeling approach:

1. an intuitive and accessible user interaction method called *procedural sketching* (Figure 1 left hand);
2. automatic *virtual world consistency maintenance* through generic methods for resolving interactions between terrain features and the landscape (Figure 1 middle).

We are developing a prototype modeling framework, called SketchaWorld, that demonstrates the feasibility of this declarative approach (see Figure 1). Its goals are:

1. to increase designers' productivity, while still allowing them to work in an iterative manner and exercise control over the generated results;
2. to provide an intuitive way for people without special modeling expertise to create virtual worlds that meet their requirements;
3. to facilitate the application of results from research in procedural methods in an integrated modeling approach.

The remainder of this article is structured as follows. Section 2 presents interactive procedural sketching. Automatic virtual world consistency maintenance is explained in Section 3. The implementation and results of the prototype framework are shown in Section 4, including an example modeling session. Section 5 discusses the advantages and current limitations of our approach. Section 6 summarizes ongoing and future work.

## 2. Interactive procedural sketching

User input and control is provided by *procedural sketching* (Figure 1 left hand). With easy to use editing tools, designers create a 2D digital sketch: a rough layout map of the virtual world. Procedural sketching provides two interaction modes:

**Landscape mode** Designers paint a top view of the landscape by coloring a grid with ecotopes (an area of homogeneous terrain and features). These ecotopes encompass both elevation information (elevation ranges, terrain roughness) and soil material information (sand, grass, rock, etc.). The grid size is adjustable and the brushes used are similar to typical brushes found in image editing software, including draw, fill, lasso, magic wand and transition pattern brushes (e.g. from ocean to shore).

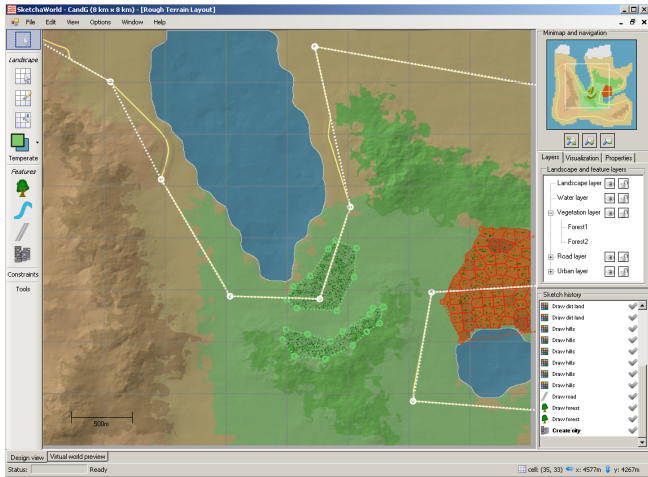


Figure 2: User interface for procedural sketching (feature mode), also showing editing tools (left hand) and navigation, layers, and edit history (right hand).

**Feature mode** Designers specify features like forests, lakes, rivers, roads, and cities on the landscape using vector lines and polygon tools. This resembles vector drawing software: placing and modifying lines and polygons is done by manipulating control points.

Each sketched element is procedurally expanded to a corresponding terrain feature (*procedural generation* box in Figure 1). To directly see the effect of an edit action on the virtual world model (e.g. drawing ecotopes, rerouting the path or modifying the shape of a feature, removing a feature), users sketch on a 2D top view of the generated virtual world. This view updates immediately as new results are generated. Depending on the interaction mode, an overlay is displayed representing relevant elements of the user sketch. Figure 2 shows the user interface for procedural sketching in feature mode. By keeping the user interface and interaction modes simple and clear, procedural sketching is more accessible for people without special modeling expertise; see also [43].

A *short feedback loop* between a designer’s edit action and the visualization of the generated results is essential to allow designers to model virtual worlds iteratively. This requires each edit action to be executed separately and the results of an action to be displayed immediately. Such an interactive setup allows designers to quickly see the effect of their edit operations and work towards the desired end result.

Several challenges have to be overcome in order to provide designers with this interactive and iterative workflow. The solutions for this are detailed in [44]. Although improvements in hardware and new approaches such as GPU computing significantly alleviate the execution time of procedural methods, operations affecting large regions or requiring complex algorithms (e.g. city generation) may still execute at non-interactive rates. Therefore an asynchronous setup was implemented, explained in Figure 3. It separates the *user interaction* and the actual *execution* of edit actions. This allows designers to continuously work on the virtual world without being hindered by

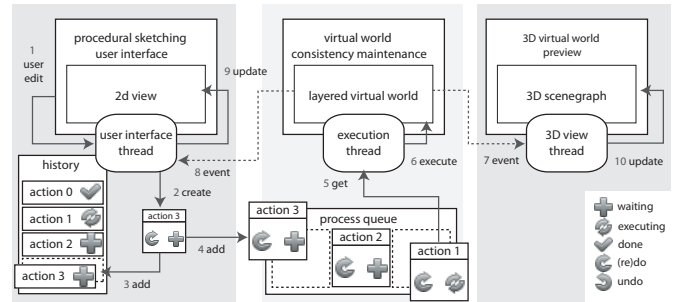


Figure 3: Diagram showing the execution flow of interactive procedural sketching. In this situation, a designer has just edited (1) the landscape or a certain terrain feature. For this, the *user interface thread* creates an edit action (2) with id number 3 and appends it to the history (3). Furthermore, a copy of the action is enqueued (4) in the *process queue*. Simultaneously, the *execution thread* polls the queue (5) and obtains action 1, which is executed (6), modifying the virtual world. These modifications trigger events indicating changes (7, 8) that are received by the user interface thread, which updates the 2D view (9) and the 3D view thread, which modifies (10) and renders the 3D scenegraph.

long procedural operations.

The ability to undo and redo any modeling action is one of the main requirements for an iterative modeling workflow. For this, the familiar edit history is provided. Because of memory constraints inherent to modeling large virtual worlds, it is far more efficient to implement undo and redo by (partial) regeneration, instead of storing all intermediate modeling states. At the expense of some additional computation time, designers are provided with unlimited undo and redo facilities. Furthermore, because designers expect regenerated results to be exactly the same as before when redoing an action, we need to carefully manage the state of the random number generator used in procedural methods, as explained in [44].

### 3. Virtual world consistency maintenance

Using procedural sketching, designers declare the properties of the landscape and specify its terrain features. In the declarative approach introduced in Section 1.3, these features are not only generated according to the user specifications, but they are also properly integrated in order to form a consistent and lifelike virtual world. Each terrain feature introduced in a virtual world affects in some way the existing landscape and nearby features, and vice versa, e.g. the feature fits itself to local constraints, it affects the structure of a nearby feature, it modifies the elevation profile, or it forms some sort of connection with a compatible feature. One can imagine the amount of tedious manual modeling work if the responsibility of solving these interactions and keeping the virtual world consistent would be left to the designer. Because the semantics of terrain features and their relations are encapsulated in the layered virtual world model, we are able to maintain the consistency in an automated manner. This section summarizes the generic resolution methods with which the consistency of the virtual world model is maintained.

### 3.1. Virtual world semantic model

The virtual world’s semantic model is a layered data structure (Figure 1 middle). All terrain features are grouped in logical layers of the virtual world model, which are inspired from Geographic Information Systems. This ordering has been chosen because of the semantic similarity and relations of the terrain features within each layer. We distinguish five specific layers, stacked as follows:

1. Urban layer: e.g. cities, districts, parcels, buildings
2. Road layer: e.g. highways, local roads and streets, bridges
3. Vegetation layer: e.g. natural forests, planted vegetation
4. Water layer: e.g. rivers, canals, lakes, oceans
5. Landscape layer: elevation profile and soil material

Each terrain feature is defined at several *levels of abstraction*, giving structure to its layout (e.g. city topology) and its objects (e.g. buildings). Procedural generation of each terrain feature can therefore be described as a top down process, starting from a coarse user specification, and refining this specification from abstract structures to concrete objects in several refinement steps, resulting in the complete feature representation. Moreover, the objects integrating a terrain feature do not only have a geometric description (e.g. polygon, spline, 3D shape), but also a *semantic* description, which includes their relevant attributes, as well as the logical connections and geometric and functional constraints involving them [37]. The number of levels of abstraction varies according to the complexity of the feature. At least, the following three levels can be discerned per feature type:

1. Specification level: user-sketched coarse outline and input parameters (e.g. a forest specification);
2. Structural level: the layout of the feature and the area it encompasses (e.g. the contour of the forest);
3. Object level: all individual semantic objects making up the feature that will result in concrete, geometric objects (e.g. the set of individual trees).

On the structural level, features only have a semantic representation; this structure helps to layout their individual objects, and to preserve their logical and functional structure (e.g. which districts make up a city center). At the object level, further structuring is provided by connections (e.g. street connectivity) and constraints (e.g. minimum distance between certain objects) between semantic objects.

### 3.2. Consistency maintenance basic notions

We first introduce several preliminary concepts and notions before discussing the actual resolution methods. In each refinement step in a terrain feature’s generation process, we can discern two phases (shown in Figure 1 middle):

1. Starting from the specification level, the next level of abstraction of the feature is procedurally generated;
2. This level of the feature is fit with its surroundings.

The generation procedure is steered by the provided specification and influenced by relevant nearby terrain features. In order to fit with its surroundings and place its objects, a feature often requires a certain area of the landscape to be clear of any other features. Moreover, the local elevation profile might be unsuitable for a feature to attach its objects; therefore, it might require the landscape to adhere to a profile constraint (e.g. regarding slope, regularity or flatness). For these requirements, features are provided with two types of requests:

**Claim** A feature can make a claim for an area of terrain in order to use it exclusively. A *claim* is a request by feature  $f_x$  to reserve an area of terrain  $a$  (defined as a complex 2D polygon) for exclusive use at the given level of abstraction  $l$ . A claim can be either granted or rejected. A feature can claim an area of terrain on either the structural or the object level of abstraction. This allows a feature to disregard irrelevant, small objects when laying out its structure, and to consider only concrete, existing objects when placing its integrating objects.

**Modification** A feature can request a local modification of the landscape (elevation profile or soil material) in order to properly fit with its requirements (see Section 3.3).

A *feature interaction* is said to occur when two claims are made for the same area of terrain. More precisely, a *feature interaction* is an overlap area  $a$  between the area claimed by feature  $f_x$  and the area already granted to another feature  $f_y$ , due to a claim at the same level of abstraction  $l$ . Because area ownership is by definition exclusive, feature interactions always have to be solved; we do this in one of two possible ways:

1. An interaction is solved in a cooperative way when it is possible to introduce a connecting structure. A *connection* formed at a level of abstraction  $l$  links the feature  $f_{lose}$ , for which the claim is rejected, to feature  $f_{win}$  over the disputed area  $a$ . The claim of  $f_{win}$  on this area  $a$  is or remains granted. Connections made on the structural level are abstract (e.g. a highway is linked to a city), and results in concrete connection objects (e.g. a road junction object), once made on the object level. Examples of connection objects include bridges, tunnels, road junctions, estuaries, etc. Of course, connections cannot be defined for every possible pair of feature types, as there might not be a sensible real world equivalent (e.g. between a lake and a forest). Furthermore, for some pairs of terrain feature types, several alternative connection types may be defined.
2. An interaction is solved in a conflictive way when one of the features,  $f_{win}$ , prevails and the other,  $f_{lose}$ , can no longer make any use of the disputed area  $a$ . If the conflict occurs at the structural level of abstraction, this entails that the feature layout of  $f_{lose}$  cannot include this area  $a$  and as such has to be restructured. At the object level, this means that the objects of  $f_{lose}$  cannot be placed within this particular area.

The decision mechanism for resolving interactions is steered by *priorities*. A priority is a ranking score or weight of a feature

---

**Algorithm 1** Landscape interaction resolution method

---

```
// solve all interactions between landscape and terrain features
// a - area of the landscape that is changed
SolveLandscapeInteractions(area a):
  // find all affected terrain features
   $F = \{ f \mid f \in \text{features}, a \cap f.\text{area} \neq \emptyset \}$ 
  sort  $F$  according to highest  $\text{priority}_{\text{claim}}(f, \text{level}_{\text{structural}})$ 
  // let each affected feature handle the landscape interaction
  for all feature  $f$  in  $F$  do
     $f.\text{restructure}(a, \text{level}_{\text{structural}})$ 
  end for
```

---

at the structural or object level. Priorities are divided in three categories:

- $\text{priority}_{\text{claim}}$ (feature  $f_x$ , level  $l$ ) is the priority of feature  $f_x$  for claiming an area of terrain at the level of abstraction  $l$ ;
- $\text{priority}_{\text{connect}}$ (feature  $f_x$ , feature  $f_y$ , area  $a$ , level  $l$ ) is the priority of feature  $f_x$  for forming a connection with feature  $f_y$  at area  $a$  and at abstraction level  $l$ ;
- $\text{priority}_{\text{conflict}}$ (feature  $f_{\text{lose}}$ , area  $a$ , level  $l$ ) indicates the *threshold* at which feature  $f_{\text{lose}}$  rejects any connection.

For each feature type and for each level, a single user-defined value serves as the default priority value for every feature instance. However, for  $\text{priority}_{\text{connect}}$  and  $\text{priority}_{\text{conflict}}$ , features can define *context-dependent* functions that compute the actual priority value, for instance a cost function for road connections. Furthermore, each specific feature instance may override any of the priority values of its type.

### 3.3. Landscape interaction

Before discussing how multiple features interact with each other, we analyze the basic interaction between the landscape and a single terrain feature. The landscape plays a special role in our virtual world model in the sense that it forms its omni-present basis on which all other terrain features attach. Therefore, the landscape does not compete with terrain features in the form of claims, and no priorities need to be defined for the landscape. A local change to the landscape affects all terrain features in that area, each to an extent that is defined by its particular semantics. This notion is captured in the landscape interaction resolution method, summarized in Algorithm 1.

As follows from this algorithm, the outcome of the landscape interacting with any feature  $f$  present in area  $a$  is always that  $f$  needs to adapt its structure. However, features decide to what extent to restructure, if at all, which typically depends on the scope of the changes to the landscape. Drastic changes in the elevation profile will probably cause features like roads or cities to strongly revise their structure, whereas changes in soil material will probably affect vegetation the most. Depending on the feature’s procedure implementation, restructuring could entail that part of the former structure and objects of the feature are removed from the respective terrain layer(s), and that corresponding area claims are abandoned, thus allowing other

---

**Algorithm 2** Feature interaction resolution method

---

```
// solve all interactions between a feature and existing features
//  $f_x$  - feature which has made a new claim
//  $a_x$  - area of terrain claimed by  $f_x$ 
//  $l$  - level of abstraction
SolveFeatureInteractions(feature  $f_x$ , area  $a_x$ , level  $l$ ):
  // find all interacting features with granted claims
   $F = \{ f_y \mid f_y \in \text{features}, a_x \cap f_y.\text{area} \neq \emptyset \}$ 
  sort  $F$  according to highest  $\text{priority}_{\text{claim}}(f_y, l)$ 
  // handle all feature interactions
  for all feature  $f_y$  in  $F$  do
    if  $\text{priority}_{\text{claim}}(f_x, l) > \text{priority}_{\text{claim}}(f_y, l)$  then
      SolveInteraction( $f_y, f_x, a_x \cap f_y.\text{area}, l$ )
    else
      SolveInteraction( $f_x, f_y, a_x \cap f_y.\text{area}, l$ )
    end if
  end for

// solve an interaction between a pair of terrain features
//  $f_{\text{lose}}$  - feature for which the claim is rejected
//  $f_{\text{win}}$  - feature for which the claim is granted
//  $a$  - disputed area
//  $l$  - level of abstraction
SolveInteraction(feature  $f_{\text{lose}}$ , feature  $f_{\text{win}}$ , area  $a$ , level  $l$ ):
  // determine whether connection can and should be formed
  if  $\text{connectionDefined}(f_{\text{lose}}.\text{type}, f_{\text{win}}.\text{type})$  and
     $\text{priority}_{\text{connect}}(f_{\text{lose}}, f_{\text{win}}, a, l) > \text{priority}_{\text{conflict}}(f_{\text{lose}}, a, l)$  then
    // interaction is solved with a connection
     $f_{\text{lose}}.\text{connectTo}(f_{\text{win}}, a, l)$ 
  else
    // interaction is solved by restructuring the losing feature
     $f_{\text{lose}}.\text{restructure}(a, l)$ 
  end if
```

---

features to reclaim these areas. We handle interactions ordered by  $\text{priority}_{\text{claim}}$ , as its value provides a good heuristic of the impact a feature will have on other features.

Although features adjust their structure to the landscape profile, to be able to attach to the landscape, features can have specific requirements for the local elevation profile. For instance, a building could require an area of flat terrain. To fulfill these requirements, features can act on the landscape through *modification* operations, which they can issue after each refinement phase. A *modification* is a request by a feature  $f_x$  for a certain area  $a$  of the landscape to be constrained to the given elevation or soil material profile  $p$ . This profile is to be integrated in the landscape according to a blending fall-off mask  $m$ . This mask allows for smooth transition zones. Such a modification request is combined with the overlapping requests of nearby features; claimed areas cannot be affected by other feature’s modify requests.

### 3.4. Feature interaction resolution

Terrain features compete with each other to claim areas of the landscape for their own use. A generic interaction resolution method has been devised to handle potentially conflicting claims between terrain features. This resolution method is outlined in Algorithm 2, and works as follows. A new terrain feature  $f_x$  can make a claim for a certain area of terrain  $a_x$ , on either the



structural or the object level  $l$ . The set of interacting features is found, and is sorted according to priority<sub>claim</sub>, where features with higher priorities are resolved first, as these features will most likely have the highest impact on the new claim.

For each  $f_y$  interacting with feature  $f_x$  on an area  $a$  and on a feature level  $l$ , the claim priorities for  $f_x$  and  $f_y$  are compared, resulting in either  $f_x$  losing the claim to  $f_y$  or vice versa.

Each interaction is then resolved either by forming a connection to the winning feature  $f_{win}$  or, if that is unavailable or has a lower priority, by a conflict where the losing feature  $f_{lose}$  has to restructure itself to avoid the lost area of terrain. More concrete, a connection must be defined between the feature types of  $f_{lose}$  and  $f_{win}$  and it must have a higher priority over conflict, which may not be the case in specific scenarios where a given connection might be deemed too costly or impractical.

Together, these resolution methods offer a generic and automated way to overcome interactions, while still providing the designer with control through setting and overriding the different types of priorities.

#### 4. Framework implementation and results

To validate our approach we implemented representative terrain features for each of the five terrain layers. Their generation procedures are often based on (combinations of) existing procedural methods, but they have been modified to fit in the framework, i.e. to consider their surroundings and context, and to employ the claim and modification methods of the interaction resolution scheme of Section 3. These features form a basic set useful for many scenarios, to support other types of scenarios; additional features might be integrated such as railways, lakes, canals, harbors, levees, farming fields, etc. For each of these terrain features, we briefly discuss their procedure and conclude with an example session using all features.

##### 4.1. Landscape layer: elevation and soil

The Landscape layer provides the foundation for placing terrain features. Its generation procedure is detailed in our previous work [43], and is outlined in Algorithm 3. The input is derived from the grid of painted ecotopes (see Figure 1). The definition of each ecotope includes, among other things, a minimum and maximum elevation and a roughness percentage, describing how rough or smooth its terrain should be. From this definition, each cell in the ecotope grid is assigned a randomized, *local* variation of these ranges. These local values in the grid are smoothed using a small Gaussian kernel, to obtain natural changes in elevation. For each point  $(x, y)$  in the landscape, the coarse grid  $g$  is interpolated with Catmull-Rom splines to obtain the ranges  $\bar{r}$  at the desired spatial resolution (e.g. 1 meter). Note that the coordinate  $(x, y)$  is first perturbed in 2D to decrease the regularity of this interpolation, resulting in  $(x', y')$ . A combination of several "flavours" of fractal noise, such as ridged multi-fractal noise [45], mixed according to the roughness factor, are used to determine the elevation value within the range  $\bar{r}$ . The distribution of soil material is based on the ecotope value at  $(x', y')$ , and by mapping the elevation value to a lookup table. The procedure

results in  $(result_{elevation}, result_{soil})$  being stored at position  $(x, y)$  in a height-map data-structure.

---

#### Algorithm 3 Landscape layer generation method

---

```
// g - coarse grid of ecotopes (elevation ranges, roughness, etc.)
GenerateLandscape(ecotope grid g, random seed s):
  for y = 0 to height do
    for x = 0 to width do
      // perturb x and y locally
      p = perturb(x, y, s) // in range [-π ... π]
      x' = x + offset_perturb cos(p), y' = y + offset_perturb sin(p)
      // obtain 4x4 grid region for Catmull-Rom interpolation
      x_g = (x' - gridcelldim(g)/2) / gridcelldim(g) // grid x
      y_g = (y' - gridcelldim(g)/2) / gridcelldim(g) // grid y
      x_1 = ⌊x_g⌋, x_f = x_g - x_1, x_0 = x_1 - 1, x_2 = x_1 + 1, x_3 = x_2 + 1
      y_1 = ⌊y_g⌋, y_f = y_g - y_1, y_0 = y_1 - 1, y_2 = y_1 + 1, y_3 = y_2 + 1
      // interpolate 4x4 grid cells based on fractions (x_f, y_f)
      // interpolation result r.xyz = (min, max, roughness)
      r̄ = spline(g, x_0, x_1, x_2, x_3, y_0, y_1, y_2, y_3, x_f, y_f)
      // generate 3 noise values, mix based on roughness
      n̄ = noiseValues(x', y', s) // all in range [-1 ... 1]
      f̄ = mixFactors(r̄.z) // all in range [0 ... 1], |f̄| = 1
      v = f̄ · n̄ // combined noise value
      // result for x, y: elevation and soil values
      result_elevation = r.x + (1/2 + v/2)(r.y - r.x)
      result_soil = distribute(x', y', result_elevation, ecotope(g, x_g, y_g))
    end for
  end for
```

---

The advantage of this combination of fractal noise and coarse grid interpolation is that it allows to evaluate any point independently of its neighbors. This makes it efficient to execute the procedure in parallel (see Section 4.7). Furthermore, the elevation and soil maps can be split in conveniently sized tiles, where each tile is computed and stored independently. This allows us to manage even relatively large landscapes at a high spatial resolution (e.g. 2 or 4 meters).

##### 4.2. Water layer: river feature

The generation of a river is constrained by the feature specification polyline and by the local elevation profile. The actual river path is determined by iteratively finding a sub-path for each pair of control points  $\bar{p}_{src}$  and  $\bar{p}_{dst}$  in the specification. Starting from the highest elevated point  $\bar{p}_{cur} = \bar{p}_{src}$ , each iteration generates  $n$  new candidate points  $\bar{p}_{can}$  on a circle with range  $r_{step}$  at an interval  $[\alpha - \alpha_{dev} \dots \alpha + \alpha_{dev}]$ , where  $\alpha$  is the angle from  $\bar{p}_{cur}$  towards  $\bar{p}_{dst}$  and  $\alpha_{dev}$  is an angle range of deviation from  $\alpha$  (e.g. 36°). Each of these generated candidates is scored according to the following weighted sum of scores:

$$S_{elevation} = \frac{\bar{p}_{cur}.z - \bar{p}_{can}.z}{\|\bar{p}_{can} - \bar{p}_{cur}\|}$$

$$S_{curve} = 1 - \cos^{-1}\left(\frac{\|\bar{p}_{dst} - \bar{p}_{cur}\|}{\|\bar{p}_{dst} - \bar{p}_{cur}\|} \cdot \frac{\|\bar{p}_{can} - \bar{p}_{cur}\|}{\|\bar{p}_{can} - \bar{p}_{cur}\|}\right) / \alpha_{dev}$$

$$S_{can} = w_{elevation} S_{elevation} + w_{curve} S_{curve} + w_{feature} S_{feature}$$

The terms  $S_{elevation}$  and  $S_{curve}$  denote the scores for local elevation difference and river curvature, respectively. The feature

score term  $s_{feature}$  is negative if the current segment crosses a feature resulting in a conflict that the river feature would lose. If the candidate with the highest score  $s_{can}$  does not adhere to the constraint of monotonously decreasing elevation, the elevation value of the candidate is forced to an elevation slightly lower than its predecessor. The determined river bank is claimed for exclusive use, with possible connections with the various types of roads, other rivers and the sea. The concrete river object is created by segmenting a Catmull-Rom spline along the path, with small variations in width according to the local slope. Using a modification request, the elevation profile is locally adapted to fit with the river’s bed and bank, and the soil is changed to a more fertile ecotope, where applicable.

The advantage of this procedure is the generation speed and the fact that it closely follows the designer’s specification. It generates good results as long as the control points are reasonably well placed, otherwise they can be interactively adjusted.

#### 4.3. Vegetation layer: forest feature

The forest feature is declared by its coarse area, vegetation density and species that occur in this forest. The input area is slightly perturbed using standard noise techniques to obtain natural tree lines. The refined forest area is claimed on the structural level. An iterative procedure based on the method introduced by Deussen et al. [7], which simulates the competition of plants for natural resources as space and water, determines a distribution of trees of different species in this forest (see [43] for specific details on this procedure). Additional input for this procedure is the local elevation profile and soil material, as the definition of a tree species includes preferences for certain soil types, as well as elevation and slope ranges. For each generated tree object, its position is claimed. The forest feature *restructure* operation removes any trees in the areas lost to other features.

#### 4.4. Road layer: primary road feature and road network

A primary road (e.g. highway, interstate) is declared by coarsely specifying its path. Its generation procedure is based on the path plotting algorithm introduced by Kelly et al. [27]. This algorithm iteratively finds a smooth path between a set of control points of a polyline defined on an elevation map. It prefers an even change in elevation from start to end, while guaranteeing all control points to be visited and the path to deviate only within a limited range from the specification. The algorithm was extended to avoid unacceptably sharp turns and slopes, to connect to existing features, such as rivers, if necessary, and to avoid potential feature conflicts with negative consequences for the road. Although the procedure uses a scoring mechanism to determine a path, it is not optimizing a cost function, as for instance the A\*-based path finding method by Galin et al. [26], and therefore is not guaranteed to find an optimal path in all cases. However, our procedure typically runs more interactively while staying close to the coarse path sketched by the designer, thereby providing more fine-grained user control.

Similarly to the river feature procedure explained in 4.2, a path is constructed between subsequent control points by generating and testing intermediate nodes. The partial procedure

outlined in Algorithm 4 explains that the scoring of candidates is again composed of three factors: change in elevation, road curvature and crossed terrain features. However, as in [27], an even change in elevation is preferred, whereas the river prefers the steepest slope. Furthermore, the weights reflect the natural preferences of the river ( $w_{elevation} = 0.7, w_{curve} = 0.1, w_{feature} = 0.2$ ) versus the road ( $w_{elevation} = 0.35, w_{curve} = 0.15, w_{feature} = 0.5$ ).

---

#### Algorithm 4 Iterative path planning for the road feature

---

```

for all subsequent points  $(\bar{p}_{src}, \bar{p}_{dst})$  in  $\bar{p}_{start} \dots \bar{p}_{end}$  do
  while !reached do
     $c = \text{generateCandidates}(\bar{p}_{src}, n_{can}, r_{step}, \alpha, \alpha_{dev})$ 
    for all  $\bar{p}_{can}$  in  $c$  do
      //  $s_e$ : local elevation change ratio / global change ratio
       $s_e = 1 - (\frac{|\bar{p}_{src.z} - \bar{p}_{can.z}|}{\|\bar{p}_{dst} - \bar{p}_{src}\| - \|\bar{p}_{dst} - \bar{p}_{can}\|} - \frac{|\bar{p}_{end.z} - \bar{p}_{start.z}|}{\|\bar{p}_{end} - \bar{p}_{start}\|})$ 
      //  $s_c$ : angle deviation from straight road
       $s_c = 1 - \arccos(\frac{\bar{p}_{dst} - \bar{p}_{cur}}{\|\bar{p}_{dst} - \bar{p}_{cur}\|} \cdot \frac{\bar{p}_{can} - \bar{p}_{cur}}{\|\bar{p}_{can} - \bar{p}_{cur}\|}) / \alpha_{dev}$ 
      //  $s_f$ : lost conflict(s) =  $-\infty$ , connection(s) = 0, none = 1
       $s_f = \text{score}(\text{findInteractions}(\bar{p}_{src}, \bar{p}_{can}))$ 
       $s[\bar{p}_{can}] = w_e s_e + w_c s_c + w_f s_f$ 
    end for
     $\bar{p}_{best} = \text{selectCandidate}(c, s)$ 
    if  $\|\bar{p}_{dst} - \bar{p}_{best}\| \leq r_{snap}$  then
      // snap to destination control point
      reached = true,  $\bar{p}_{best} = \bar{p}_{dst}$ 
    end if
     $\bar{p}_{src} = \bar{p}_{best}$ 
  end while
end for

```

---

After the 3D path has been planned, the elevation profile of this path is fit to a defined slope constraint and smoothed using a Gaussian smoothing kernel. Catmull-Rom interpolation results in a 3D spline. By segmenting and sweeping along the road spline, we obtain the road surface. The area of the road is then claimed on the structural level. Connections have been defined with other road features (junction), river features (bridges), and city features (junctions with local road network’s outer ring). Using a modification request, the elevation profile is cut and filled to match with the desired embankment profile of this road. This embankment profile consists of the inner road area, the embankment and the verge. For the verge, a blending zone is established to create a natural smooth transition between the road and the surrounding landscape.

A city’s road network consists of secondary and tertiary roads. It is composed of road patterns, as described by Sun et al. [8], where the secondary roads follow the population centers and the tertiary roads follow the grid, radial or mixed patterns. The procedure iteratively places new roads within the city, while checking validity constraints concerning angle, slope and length, and applying snapping rules, such as snapping a new road to existing nearby junctions (see, e.g. [9]). The primary road and a city’s road network can form connections by road junctions.

#### 4.5. Urban layer: city feature

A designer defines a city feature by coarsely sketching its outline and selecting a template describing its historical background,



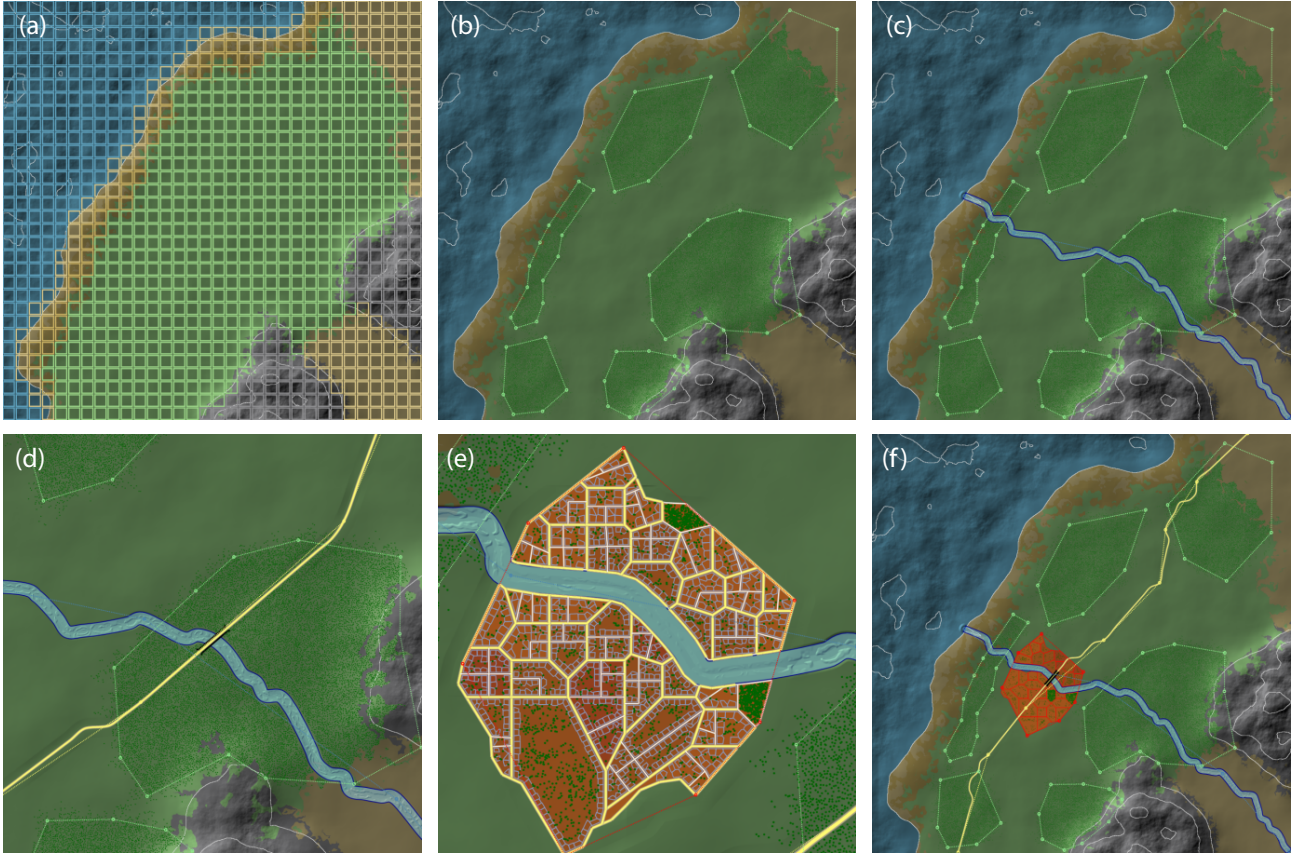


Figure 4: A procedural sketching session: (a) basic landscape, defined by brushing ecotopes, (b) several forest features, (c) river flowing towards the sea, (d) road feature crossing this river, (e) city created along the river, (f) road rerouted to run through the city.

which defines how the city core is structured and includes a statistical distribution of types of districts typically found in these kinds of cities. Template examples for a Western European city include the mercantile, the feudal or the absolutistic historic city. The designer can also declare the city's population size, affecting the number and type of residential districts that will be created.

The generation procedure is described in detail in [46]. The city area is first split in one or more clusters. A cluster is an area of landscape limited by either the city bounds or land that is either unsuitable for building or could not be claimed from other features. A distribution of districts within these clusters is determined by an iterative district placement procedure, which uses constraints derived from established models of urban land use. In this procedure, terrain features have an attraction or repulsion influence on districts, depending on their type (commercial, heavy industry, high-class residential, etc.). For each district, candidate locations are generated within the city bounds and evaluated according to a suitability score, which is a weighted sum of scores for several factors, including the type and proximity of other, already placed districts, the location's soil material, the area within the city model (e.g. center, suburban), and the distance to nearby rivers and to primary roads. Each district is placed at the most suitable candidate location.

After districts have been placed, secondary roads and streets are placed to form the road network of the city (see Section

4.4). Within this network, parcels and building lots result from subdividing the available open space, and buildings are generated based on footprint shapes and district types (see [47] for details). Each generated street and building within the city issues a modification request to constrain the local elevation profile.

#### 4.6. Example

To give a more vivid impression of how one can currently design a virtual world in SketchaWorld, we present the (intermediate) results of an example session (see Figure 4), in which a designer creates, in a couple of minutes, a landscape with a city along a river. The example session also involves several instances of consistency maintenance, resulting in conflicts, connections and modifications to the landscape. Throughout this session, a 3D geometric model derived from the layered semantic virtual world model, is incrementally updated (Figure 1 right hand), allowing designers to preview the results in 3D.

Figure 4.a shows the basic landscape, sketched in *landscape mode* by brushing the ecotope grid: a coastline with some mountains and dry land in the east. On top of this landscape, using the tools of the *feature mode*, features of the landscape are added. The designer specifies several forest features using polygons, resulting in the vegetation distribution shown in Figure 4.b. Note that the trees do not grow on steep slopes, rocky terrain and barren land.




Figure 5: 3D virtual world, resulting from the example session of Figure 4: (a) natural environment with road crossing river, (b) final virtual world, (c) city close-up.

A river feature is sketched to run from the mountain lands in the south-east towards the ocean. After the river’s path is coarsely defined, a suitable course is plotted through the landscape. Through the default priorities, the claim for land is granted to the river, which leads to restructuring of the affected forests. Using a modification request, the river bed and banks are carved into the elevation profile (Figure 4.c). In Figure 4.d the designer adds a main road to the virtual world. Although here the river has a higher priority than the road, a connection is defined between the two features: a bridge is inserted at the river crossing connecting the road segments. Again, the landscape is modified, in this case to form a road embankment (except at the crossing which is still claimed by the river). The forests lose again some of their claimed terrain, now to this primary road. Figure 5.a shows the resulting virtual world.

Subsequently, the designer specifies a small city with a fairly low priority. As a result, its districts and secondary roads have to form around the river, as shown in Figure 4.e. Finally, the designer decides to reroute the primary road to now run through this city (Figure 4.f). This edit action affects the city’s claim, therefore, this city is restructured to include this main road and its bridge. Figure 5.b shows the final virtual world, and details of the city are shown in Figure 5.c.

#### 4.7. Implementation and performance

SketchaWorld is implemented in C# and C++ and uses OpenSceneGraph for the visualization of the resulting 3D virtual world. Several of its generation and integration procedures are highly suitable for a parallel implementation, and where therefore implemented on the GPU using CUDA, a C-like programming language for performing general purpose parallel computations on GPU’s. These GPU procedures are an order of magnitude more efficient than their CPU counterparts, thus making interactive procedural sketching feasible. Examples of procedures implemented in CUDA include the elevation and soil map generation procedure, and the procedures for landscape modification requests. In addition, it is used for creating images, such as textures and the 2D view on the virtual world.

Regarding the quality of the resulting 3D virtual world models, we have aimed at achieving *functional realism*, i.e. to generate a plausible and consistent structure and a logical placement of objects and features. Our goal is to be able to export created virtual worlds to a number of engines which differ very much

in their graphical capabilities and application (e.g. entertainment games, training simulators). Therefore, instead of focusing on graphics quality, we have concentrated on making new exports straightforward to develop, by clearly separating the virtual worlds semantic model from any derived representations (see Figure 1) and, particularly, by facilitating the replacement of currently used content (e.g. textures) with high-quality content. Currently supported export formats are OpenSceneGraph, COLLADA and GIS raster and vector data; we are developing additional exporters to widen the possible application of our prototype.

The current prototype’s performance is good for small and medium sized landscapes (up to 1000 km<sup>2</sup>). Most procedural operations typically take between 100 ms and 3 seconds. However, when modeling larger landscapes continuously updating the 3D geometry of the virtual world becomes a bottleneck that hinders the interactive feedback loop. To further improve the user experience of the prototype and to enable modeling larger landscapes, we plan to optimize some of the complex generation procedures and make more use of GPU computing.

## 5. Discussion

The combination of procedural sketching and virtual world consistency maintenance helps designers to concentrate on what they want to create and to express this intent at a high level of abstraction, without being bothered with low-level modeling tasks. We have had many informal user sessions with simulation and game development professionals, training instructors, game researchers and students, which provided valuable feedback on our modeling approach and current prototype. Typically, they found the procedural sketching approach very easy to use, requiring no 3D modeling experience, and were able to create a virtual world matching their intent within minutes.

Although the approach gives designers proper control at a high level, especially game designers commented that they miss even more fine-grained user control, for instance to tune generated objects (e.g. individual buildings or trees) to precisely match their intent or artistic vision. We therefore believe that the integration of both procedural generation and manual editing operations is a very promising direction for this research. It seems to combine the best of two worlds, but so far the topic is as good as unaddressed, and it provides numerous new challenges to

overcome (see [48]). For instance, manually rerouting a generated street through a city center might entail automatically moving and removing certain blocks of buildings and possibly creating new lots and individual buildings. However, during partial regenerations of such areas (for instance because of feature interactions), manual changes should be preserved as much as possible, which makes these situations especially complex.

Similarly, although designers can influence the consistency maintenance mechanism through setting priorities, either per feature type, or by overriding e.g. the claim priority of a specific feature instance (e.g. a historical wood within a city), we think this level of control may be too coarse to accommodate all modeling scenarios. We are investigating more fine-grained mechanisms for constraining the automatic consistency maintenance, which might include locking features or objects from modification, grouping objects together, introducing local geometric constraints, etc.

A major advantage of this approach is the *flexible* support for integrating both existing and new procedural methods within the same framework, especially since the semantics and, particularly, the interactions of terrain features are defined *independently* from their generation procedures. The framework provides a clean interface and a structured way to incorporate new procedural methods for terrain features as e.g. railways, lakes, canals, or replace existing methods with sophisticated procedures for e.g. complex urban environments. The interaction resolution mechanism is flexible enough to accommodate for interactions with newly introduced features, as well as new types of connections. However, it is primarily geared towards solving interactions that are either user-intended (e.g. a bridge connection) or unavoidable (e.g. removing all trees in the highway path). Naturally, to be integrated effectively in this framework, a procedure has to be made aware of its surroundings and be able to cope with losing claims. Furthermore, the procedure should be made compatible with interactive procedural sketching, which might not be straightforward for procedural approaches that are based on optimization or growth. For these procedures, a significant amount of work could be required to fit them in the framework, after which the respective feature interactions can be properly handled.

The additional execution time for consistency maintenance is in typical modeling cases hardly noticeable. However, generating or removing a large, high priority feature can, of course, have a significant impact on nearby features, resulting in a cascade of many conflicts and connections to be resolved. Such extreme situations can be alleviated by using proper heuristics or conflict avoidance during procedural generation (see e.g. Section 4.4).

The framework SketchaWorld is currently used in several research projects and for a number of simulators. We are investigating possibilities of making it available to a wider audience.

## 6. Conclusions

We identified the challenges currently faced by designers of virtual worlds, arising from limitations of both manual and procedural modeling techniques. From recent developments in procedural modeling research, we concluded that improvements

in user control, interactivity and integration of results are now not only feasible but also essential to increase the acceptance of procedural techniques in mainstream virtual world development.

In this article we introduced declarative modeling of virtual worlds, a novel approach that combines the integrated use of various procedural modeling techniques with a semantics-driven model to capture designers intent. This approach was illustrated using our research prototype, the modeling framework SketchaWorld. One of its main contributions is *procedural sketching*, which combines user interaction, data representation and procedural techniques in order to enable designers of virtual worlds to concentrate on stating *what* they want to create, instead of on describing *how* they should model it. This is achieved by means of familiar sketch actions, promoting interactive experimentation to quickly see the effects of procedural modeling operations.

This interaction method is complemented by *virtual world consistency maintenance*, which automatically solves conflicts due to spatial interactions occurring between terrain features and their surroundings. By maintaining in the virtual world semantic model more information than simply its geometric data, this technique assists designers in keeping it in a coherent state, in which the semantics of all terrain features is preserved. This encourages designers to freely experiment with design alternatives without the penalty of laborious manual editing, as feature integration and consistency are taken care of automatically. In traditional modeling systems, in turn, making large scale changes to a virtual world is mostly prohibitive, in terms of both time and manual effort.

Current research is focusing on performance optimization of individual procedures and on enhancements of the quality and diversity of the 3D output. However, our main research challenge ahead concerns new mechanisms for designers to interactively manipulate generated objects, and for controlling and handling the consequences of these manual operations. We believe that supporting such a fine level of editing will be crucial to have designers fully profit from procedural techniques, while remaining in control over their outcome.

Concluding, the *declarative modeling approach* presented here provides virtual world *designers* with the productivity gain of procedural generation techniques, while still allowing for user control; it provides *researchers* with a flexible platform to integrate new techniques from procedural modeling research within an interactive environment; and, finally, it provides whole new groups of *non-specialists* with a clear and accessible method to create complete 3D virtual worlds in minutes.

## Acknowledgments

This research has been supported by the GATE project, funded by the Netherlands Organization for Scientific Research and the Netherlands ICT Research and Innovation Authority. We would like to thank the reviewers for their constructive criticism.

- [1] G. S. P. Miller, The Definition and Rendering of Terrain Maps, in: SIGGRAPH '86: Proceedings of the 13<sup>th</sup> Annual Conference on Computer Graphics and Interactive Techniques, Vol. 20, ACM, New York, NY, USA, 1986, pp. 39–48.
- [2] F. K. Musgrave, Methods for Realistic Landscape Imaging, Ph.D. thesis, Yale University, New Haven, CT, USA (1993).



- [3] A. D. Kelley, M. C. Malin, G. M. Nielson, Terrain Simulation Using a Model of Stream Erosion, in: SIGGRAPH '88: Proceedings of the 15<sup>th</sup> Annual Conference on Computer Graphics and Interactive Techniques, ACM, New York, NY, USA, 1988, pp. 263–268.
- [4] P. Prusinkiewicz, M. Hammel, A Fractal Model of Mountains with Rivers, in: Proceeding of Graphics Interface '93, 1993, pp. 174–180.
- [5] S. T. Teoh, River and Coastal Action in Automatic Terrain Generation, in: CGVR 2008: Proceedings of the 2008 International Conference on CG & VR, CSREA Press, Las Vegas, Nevada, USA, 2008, pp. 3–9.
- [6] P. Prusinkiewicz, A. Lindenmayer, The Algorithmic Beauty of Plants, Springer-Verlag, New York, NY, USA, 1990.
- [7] O. Deussen, P. Hanrahan, B. Lintermann, R. Měch, M. Pharr, P. Prusinkiewicz, Realistic Modeling and Rendering of Plant Ecosystems, in: SIGGRAPH '98: Proceedings of the 25<sup>th</sup> Annual Conference on Computer Graphics and Interactive Techniques, ACM, New York, NY, USA, 1998, pp. 275–286.
- [8] J. Sun, X. Yu, G. Baciú, M. Green, Template-based Generation of Road Networks for Virtual City Modeling, in: VRST '02: Proceedings of the ACM Symposium on Virtual Reality Software and Technology, ACM, New York, NY, USA, 2002, pp. 33–40.
- [9] Y. I. H. Parish, P. Müller, Procedural Modeling of Cities, in: SIGGRAPH '01: Proceedings of the 28<sup>th</sup> Annual Conference on Computer Graphics and Interactive Techniques, ACM, New York, NY, USA, 2001, pp. 301–308.
- [10] P. Wonka, M. Wimmer, F. Sillion, W. Ribarsky, Instant Architecture, in: SIGGRAPH '03: Proceedings of the 30<sup>th</sup> Annual Conference on Computer Graphics and Interactive Techniques, ACM, New York, NY, USA, 2003, pp. 669–677.
- [11] P. Müller, P. Wonka, S. Haegler, A. Ulmer, L. V. Gool, Procedural Modeling of Buildings, in: SIGGRAPH '06: Proceedings of the 33<sup>rd</sup> Annual Conference on Computer Graphics and Interactive Techniques, ACM, New York, NY, USA, 2006, pp. 614–623.
- [12] D. Finkenzerler, Detailed Building Facades, IEEE Computer Graphics and Applications 28 (3) (2008) 58–66.
- [13] R. M. Smelik, K. J. de Kraker, T. Tuteneel, R. Bidarra, S. A. Groenewegen, A Survey of Procedural Methods for Terrain Modelling, in: Proceedings of the CASA Workshop on 3D Advanced Media In Gaming And Simulation (3AMIGAS), Amsterdam, The Netherlands, 2009.
- [14] S. Stachniak, W. Stuerzlinger, An Algorithm for Automated Fractal Terrain Deformation, Computer Graphics and Artificial Intelligence 1 (2005) 64–76.
- [15] H. Zhou, J. Sun, G. Turk, J. Rehg, Terrain Synthesis from Digital Elevation Models, IEEE Transactions on Visualization and Computer Graphics 13 (4) (2007) 834–848.
- [16] R. L. Saunders, Terrainsaurus: Realistic Terrain Synthesis Using Genetic Algorithms, Master's thesis, Texas A&M University (December 2006).
- [17] K. R. Kamal, Y. S. Uddin, Parametrically Controlled Terrain Generation, in: GRAPHITE '07: Proceedings of the 5<sup>th</sup> International Conference on Computer Graphics and Interactive Techniques in Australia and Southeast Asia, ACM, New York, NY, USA, 2007, pp. 17–23.
- [18] F. Belhadj, Terrain Modeling: a Constrained Fractal Model, in: AFRI-GRAPH '07: Proceedings of the 5<sup>th</sup> International Conference on Computer Graphics, Virtual Reality, Visualisation and Interaction in Africa, ACM, New York, NY, USA, 2007, pp. 197–204.
- [19] J. Doran, I. Parberry, Controlled Procedural Terrain Generation Using Software Agents, IEEE Transactions on Computational Intelligence and AI in Games 2 (2) (2010) 111–119.
- [20] J. Schneider, T. Boldte, R. Westermann, Real-Time Editing, Synthesis, and Rendering of Infinite Landscapes on GPUs, in: Vision, Modeling and Visualization 2006, 2006.
- [21] O. Stava, B. Beneš, M. Brisbin, J. Křivánek, Interactive Terrain Modeling Using Hydraulic Erosion, in: EG / SIGGRAPH Symposium on Computer Animation, Eurographics Association, Dublin, Ireland, 2008, pp. 201–210.
- [22] J. Gain, P. Marais, W. Strasser, Terrain Sketching, in: I3D '09: Proceedings of the 2009 Symposium on Interactive 3D Graphics and Games, ACM, New York, NY, USA, 2009, pp. 31–38.
- [23] G. de Carpentier, R. Bidarra, Interactive GPU-based Procedural Height-field Brushes, in: Proceedings of the 4<sup>th</sup> International Conference on the Foundations of Digital Games, Florida, USA, 2009.
- [24] G. Chen, G. Esch, P. Wonka, P. Müller, E. Zhang, Interactive Procedural Street Modeling, in: SIGGRAPH '08: Proceedings of the 35<sup>th</sup> Annual Conference on Computer Graphics and Interactive Techniques, Vol. 27, ACM, New York, NY, USA, 2008, pp. 1–10.
- [25] J. McCrae, K. Singh, Sketch-based Path Design, in: GI '09: Proceedings of Graphics Interface 2009, Canadian Information Processing Society, Toronto, Ontario, Canada, 2009, pp. 95–102.
- [26] E. Galin, A. Peytavie, N. Marchal, E. Gurin, Procedural Generation of Roads, in: Computer Graphics Forum: Eurographics 2010 Proceedings, Vol. 29, Eurographics Association, Norrköping, Sweden, 2010, pp. 429–438.
- [27] G. Kelly, H. McCabe, Citygen: An Interactive System for Procedural City Generation, in: Proceedings of the 5<sup>th</sup> Annual International Conference in Computer Game Design and Technology, Liverpool, UK, 2007, pp. 8–16.
- [28] M. de Villiers, N. Naicker, A Sketching Interface for Procedural City Generation, Tech. rep., University of Cape Town (November 2006).
- [29] B. Weber, P. Müller, P. Wonka, M. Gross, Interactive Geometric Simulation of 4D Cities, Proceedings of Eurographics 2009 28 (2009) 481–492.
- [30] C. A. Vanegas, D. G. Aliaga, B. Beneš, P. A. Waddell, Interactive Design of Urban Spaces Using Geometrical and Behavioral Modeling, ACM TOG: Proceedings of ACM SIGGRAPH Asia 2009 28 (5) (2009) 1–10.
- [31] M. Lipp, P. Wonka, M. Wimmer, Interactive Visual Editing of Grammars for Procedural Architecture, in: SIGGRAPH '08: Proceedings of the 35<sup>th</sup> Annual Conference on Computer Graphics and Interactive Techniques, ACM, New York, NY, USA, 2008, pp. 1–10.
- [32] P. Amburn, E. Grant, T. Whitted, Managing Geometric Complexity with Enhanced Procedural Models, in: SIGGRAPH '86: Proceedings of the 13<sup>th</sup> Annual Conference on Computer Graphics and Interactive Techniques, ACM Press, New York, NY, USA, 1986, pp. 189–195.
- [33] E. Bruneton, F. Neyret, Real-time Rendering and Editing of Vector-based Terrains, in: Computer Graphics Forum: Eurographics 2008 Proceedings, Vol. 27, Crete, Greece, 2008, pp. 311–320.
- [34] T. Tuteneel, R. Bidarra, R. M. Smelik, K. J. de Kraker, The Role of Semantics in Games and Simulations, ACM Computers in Entertainment 6 (2008) 1–35.
- [35] R. Bidarra, W. Bronsvort, Semantic Feature Modelling, Computer-Aided Design 32 (3) (2000) 201–225.
- [36] M. Kallmann, D. Thalmann, Modeling Objects for Interaction Tasks, in: EG Workshop on Animation and Simulation, 1998, pp. 73–86.
- [37] T. Tuteneel, R. Bidarra, R. M. Smelik, K. J. de Kraker, Using Semantics to Improve the Design of Game Worlds, in: Proceedings of AIIDE 2009 - 5<sup>th</sup> Conference on Artificial Intelligence and Interactive Digital Entertainment, Stanford, CA, USA, 2009.
- [38] T. Tuteneel, R. Bidarra, R. M. Smelik, K. J. de Kraker, A Semantic Scene Description Language for Procedural Layout Solving Problems, in: Proceedings of AIIDE 2010 - 6<sup>th</sup> Conference on Artificial Intelligence and Interactive Digital Entertainment, Stanford, CA, USA, 2010.
- [39] A. Torpy, L3DT, <http://www.bundysoft.com/L3DT/>.
- [40] Greenworks, XFrog, <http://www.xfrog.com>.
- [41] Procedural, inc., CityEngine, <http://www.procedural.com>.
- [42] PixelActive, Cityscape 1.8, <http://pixelactive3d.com>.
- [43] R. M. Smelik, T. Tuteneel, K. J. de Kraker, R. Bidarra, Declarative Terrain Modeling for Military Training Games, International Journal of Computer Game Technology (IJCGT) 2010 (2010) 1–11.
- [44] R. M. Smelik, T. Tuteneel, K. J. de Kraker, R. Bidarra, Interactive Creation of Virtual Worlds Using Procedural Sketching, in: Proceedings of Eurographics 2010: Short Papers, Eurographics Association, 2010.
- [45] F. K. Musgrave, C. E. Kolb, R. S. Mace, The Synthesis and Rendering of Eroded Fractal Terrains, in: SIGGRAPH '89: Proceedings of the 16<sup>th</sup> Annual Conference on Computer Graphics and Interactive Techniques, ACM, New York, NY, USA, 1989, pp. 41–50.
- [46] S. A. Groenewegen, R. M. Smelik, K. J. de Kraker, R. Bidarra, Procedural City Layout Generation Based On Urban Land Use Models, in: Eurographics 2009: Short Papers, Eurographics Association, Munich, Germany, 2009, pp. 45–48.
- [47] R. M. Smelik, T. Tuteneel, K. J. de Kraker, R. Bidarra, A Case Study on Procedural Modeling of Geo-Typical Southern Afghanistan Terrain., in: Proceedings of the IMAGE 2009 Conference, IMAGE Society, St. Louis, MO, USA., 2009, pp. 329–337.
- [48] R. M. Smelik, T. Tuteneel, K. J. de Kraker, R. Bidarra, Integrating Procedural Generation and Manual Editing of Virtual Worlds, in: PCGames '10: Proceedings of the 2010 Workshop on Procedural Content Generation in Games, ACM, New York, NY, USA, 2010, pp. 1–8.