

Automatic Construction of High Quality Roadmaps for Path Planning

D.Nieuwenhuisen

A.Kamphuis

M.Mooijekind

M.H.Overmars

institute of information and computing sciences, utrecht university

technical report UU-CS-2004-068

www.cs.uu.nl

Automatic Construction of High Quality Roadmaps for Path Planning

D.Nieuwenhuisen A.Kamphuis M.Mooijekind M.H.Overmars
email: {arnok,dennis,mmooijek,markov}@cs.uu.nl

December 24, 2004

Abstract

Path planning plays an important role in many virtual worlds, like computer games. Currently the motion of entities is often planned using a combination of scripting, grid-search methods, local reactive methods, flocking and crowd behavior. In this paper we describe a new approach, based on a technique from robotics, that computes a roadmap of smooth, collision-free, high-quality paths. This roadmap can be used to obtain instantly good paths for entities. We also describe applications of the technique for planning the motion of groups of entities and for creating smooth camera movements through an environment.

1 Introduction

In many virtual worlds, entities like world inhabitants, non playing characters (NPCs), and vehicles, must plan their motions between locations in the virtual world. Currently this is typically achieved using a combination of scripting, A^* -based grid search, local reactive methods, flocking and crowd behavior.

In scripting, the designer explicitly described the paths that can/must be followed by the entities. This is normally part of level design. Scripting is a time consuming process for the designer. In addition, it can lead to repetitive behavior that is easily observed by the user. Scripting gets increasingly complicated when many entities move in the same space.

Grid based methods divide the world in a grid of cells and plan motion using an A^* -based search on the free cells (see e.g. [5, 21]). If the virtual world is small and grid-like, this technique is useful. However, when the world becomes complicated and large and many entities must move around, grid based methods take a large amount of computer time. Pruning the search reduces the time but might lead to wrong paths. Also, motions created by grid search tend to be unnatural because, in general, for grid searches a smooth path needs to be created in the post-processing phase resulting in expensive queries.

Reactive methods adapt a previously computed motion to obstacles found near the path that were not taken into account during initial planning; for example other entities or small, movable objects (see e.g. [13, 22, 1, 18]). Even though theoretically reactive methods can be used to compute full paths, this normally leads to dead-lock situations in which the entities no longer know where to move (e.g. they get stuck in a corner of the room). Another problem with reactive methods is that it is often difficult to adapt the internal animation of the character to the motions produced.

In robotics, many other path planning approaches have been developed that might be applicable to path planning in virtual environments, like games or simulators. In robotics though, the emphasis


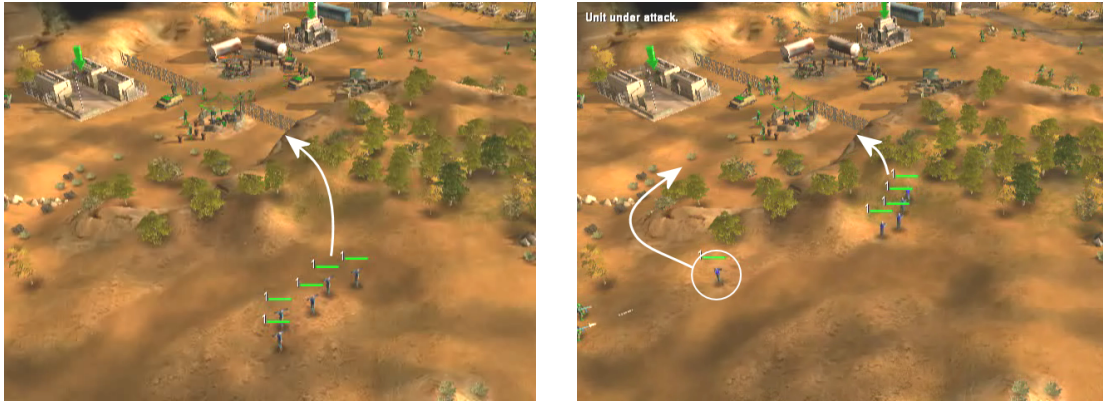


Figure 1: An example of a smooth roadmap computed by our technique.



(a) A group of five characters should attack the site pointed to by the arrow. (b) The group inappropriately splits up, losing some of its troops.

Figure 2: One of the problems with the current techniques for motion planning for multiple units is that the group splits up to reach the goal. This scene was taken from Command and Conquer: Generals from EA Games.

is often on the motion of a complicated robotic system in a relatively simple environment. In most virtual world applications the opposite is true. From a path planning perspective, the entity can often be modeled as a simple vertical cylinder, while the environment can be very complicated with tens of thousands of obstacles.

One popular path planning technique in robotics is the Probabilistic Roadmap Method (PRM). It has been studied by many authors, see e.g. [12, 25, 11, 7]. In a pre-processing phase this method builds a roadmap of possible motions of the robot through the environment. When a particular path planning query must be solved, a path is retrieved from this roadmap using a simple and fast graph search. The PRM approach is suited for very complicated environments. Unfortunately, the roadmap produced by the method can be rather wild, leading to low quality paths, consisting of straight line segments, that require a lot of time-consuming smoothing in order to be useful for virtual world applications.

In this paper we will describe a new path planning approach, building on the PRM method, that can be effectively applied in virtual world applications. The approach also constructs a roadmap of possible motions but guarantees that the paths are short, have enough clearance from the obstacles, and are C^1 continuous, leading to natural looking motions. See Figure 1 for an example of a roadmap computed using our approach. After constructing the roadmap, which can be done in the pre-processing phase, paths can be retrieved almost instantaneously, and do not require any post-processing.

Besides the standard application, in which the roadmap is used for planning the paths for individual entities, we describe two additional applications. First, we consider the motion of groups of entities. This problem is often solved using a combination of grid-based planning and flocking [19, 20]. Unfortunately, this can lead to unwanted behavior where the group of entities splits up (see Figure 2 for an example). We will use the smooth paths computed by the new planning approach as a backbone path and then use a social potential field approach to guide the flock through a corridor around this path (extending our earlier work in [10]). This results in a natural motion in which the group is guaranteed to stay together.

In many virtual world applications, moving the camera through an environment is one of the most fundamental operations. Currently the camera is often under direct control of the user or under indirect control of the user. Direct camera control is difficult, it easily leads to motion sickness due to redundant motions, and is often not required. Building on our earlier work in [15] we describe a

method in which the user only specifies positions of interest and the camera automatically moves to such positions using a smooth, collision free motion. For computing the camera path we will use the new planning approach described. This is then combined with techniques to control the view direction and the speed of the camera to obtain a camera motion that is pleasant to watch.

2 Roadmap generation

In this section we will describe how, in a pre-processing phase, a roadmap of possible motions for the entities can be computed. A roadmap is normally represented as a graph in which the nodes correspond to placements of the entity and the edges represent collision-free paths between these placements. A standard technique for automatic roadmap creation is the probabilistic roadmap method (PRM).

Unfortunately, the PRM method leads to low quality roadmaps that can take long detours. This is due to the random nature of the PRM method. Techniques exist to improve paths in a post-processing stage but this is time-consuming and can still lead to long detours. Here we present a variant of the PRM method that creates short, smooth and high quality roadmaps in the pre-processing phase. These roadmaps can then in real-time be used to solve path planning queries almost instantaneous using a simple shortest path graph search algorithm (for example Dijkstra’s shortest path algorithm).

In the pre-processing phase we create the roadmap graph, consisting of vertices (V) and edges (E). For the placement of the entity we only take its position (x, y) into account since these are the only parameters that are important for planning the path. Later, the other parameters (such as orientation) can be added depending on the application. The edges of the graph will represent straight line and circular paths between the vertices. Only vertices and edges that are collision free are allowed in the graph.

In order to be able to use the graph for as many different queries as possible, we need a good coverage of the space. Many improvements for the PRM method have been proposed in order to achieve this (see e.g. [2, 3, 17, 27, 8, 4, 9]), but all are based on the same underlying concept and lead to roadmaps consisting of straight line segments only that result in low path quality.

For a roadmap that is used to steer entities, we can formulate the following criteria:

- The paths generated by the roadmap should always keep some minimum amount of clearance from the obstacles in the scene.
- The paths should be smooth i.e. it should be C^1 continuous.
- A path needs to be created very fast (not delaying the motion) and should be short, not taking any significant detours.

2.1 Creating samples on the Voronoi diagram

The Voronoi diagram of a scene defines for every obstacle a set of points in the free space that are closer to this obstacle than to any other obstacle in the scene; together these points form the Voronoi diagram (see Figure 3 for an example).

Here, we propose a new variant of the PRM method that uses the Voronoi diagram as a guide. The method works as follows. In every iteration of the algorithm we randomly pick a sample (placement) of the entity, we call it c . Then we check whether c is collision free for the entity. If this is the case, we continue by retracting it to the Voronoi diagram using the following procedure. We calculate the closest point on an obstacle from c , we call this point c_c (Figure 4(a)). Next, another sample c' is

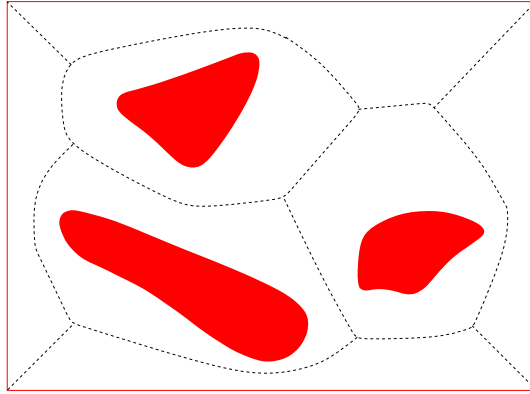


Figure 3: An example of a Voronoi diagram. We treat all four boundaries of the workspace as separate obstacles.




Figure 4: Retracting samples to the Voronoi diagram.

moved from c in the opposite direction of c_c using as a step size the distance between c and c_c (Figure 4(b)). We proceed until the closest obstacle to c' changes. We now have two samples c and c' , both having another closest obstacle. The above procedure guarantees that the Voronoi diagram passes through a point between c and c' .

We continue by using binary search between c and c' , with precision ϵ until we have found a sample that has its two closest obstacles at the same distance. This sample, called c_v , is at most a distance ϵ away from the Voronoi diagram (Figure 4(c)). Now, we add c_v to the list of vertices V in the roadmap graph.

After adding a sample as a vertex to the graph, we determine its neighbor vertices. The set of neighbor vertices of vertex v , called N_v is defined as all vertices V that are closer to v than some chosen maximum neighbor distance. For each vertex v_n in N_v we test whether the straight line connection between v and v_n is collision free. If this is the case, then we add the connection (v, v_n) as an edge to the set of edges E of the graph. If two vertices are already connected in the graph (via other vertices), then we only add the new edge if the path between v and v_n is shortened considerably by at least some constant K (for more details see [16]).

2.2 Retracting edges

We have retracted the nodes of the graph to the Voronoi diagram (within a certain boundary ϵ) but when connecting the samples with edges, these edges are usually not on the Voronoi diagram and can get very close to obstacles (Figure 5(a)). In order to solve this problem, edges are retracted to the

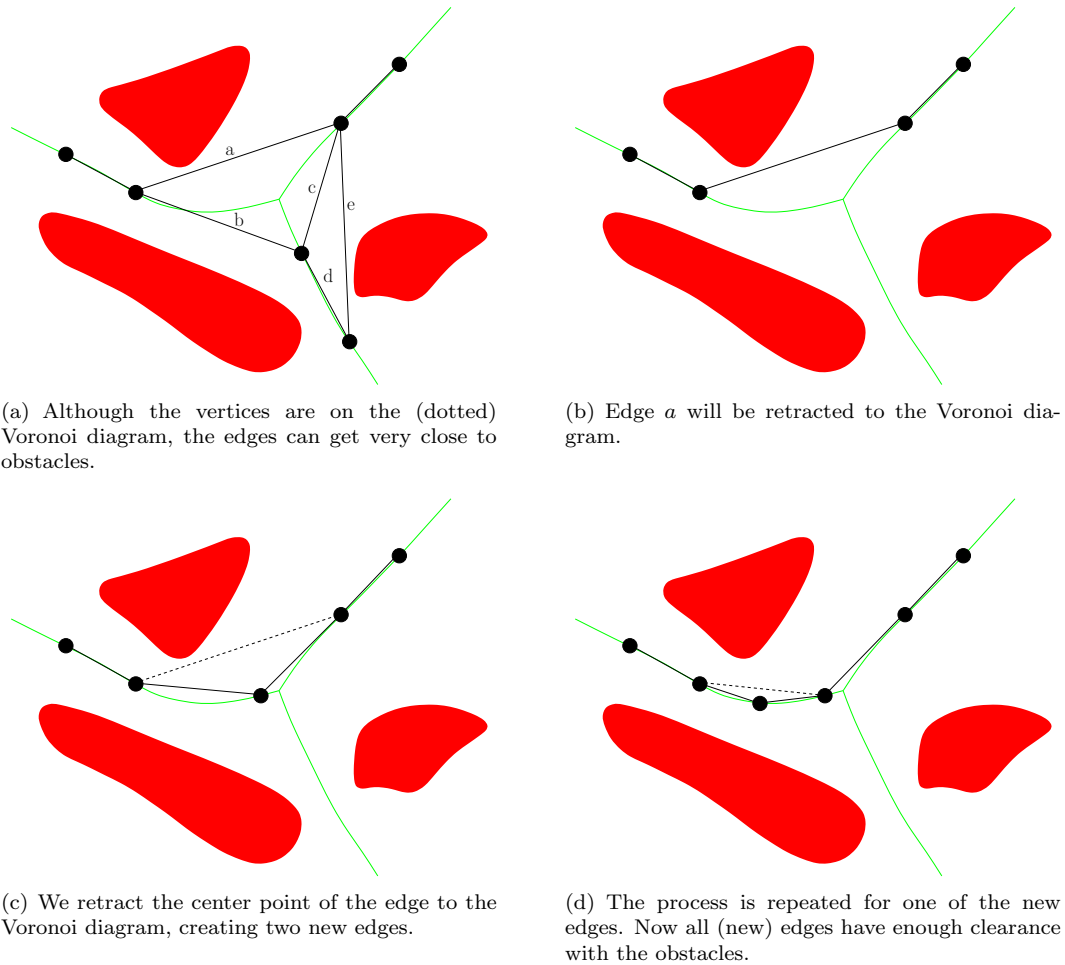


Figure 5: Retracting an edge to the Voronoi diagram.

Voronoi diagram until every part of the edge is at least some pre-specified distance away from the obstacles. We achieve this by proceeding in the following manner: if (a part of) an edge is too close to an obstacle, this edge is split in two equal length parts and the middle point is retracted to the Voronoi diagram using the same procedure as described in the previous section. This procedure is recursively repeated for the two new edges until every edge has enough clearance with the obstacles. An example of this procedure is shown in Figure 5. In some cases (when the edge passes through a very narrow corridor), the clearance threshold will never be reached and the edge will be split an infinite number of times. In order to prevent this, we stop retracting edges if their length is shorter than some predefined value.

Retracting edges to the Voronoi diagram may result in some edges overlapping each other. For example in Figure 5(a) if edges c and d are already retracted, then retracting e will result in overlap. Fortunately, detecting overlapping edges is easy. For every pair of edges e_i and e_j , we check how far their endpoints are away from the other edge. If this distance is smaller than some predefined distance, then we try to project the vertices of e_i on e_j and vice versa. If at least one of these projections is successful, we call the two edges overlapping and we can join them. We can distinguish four different kinds of overlapping edges. In Figure 6 these are shown together with the situation after removing the overlap.




Figure 6: Merging two overlapping edges e_0 and e_1 . Four different cases can be distinguished (a..d). The results after merging are shown in (e..h).

2.3 Improving the roadmap

In the previous sections, we optimized the clearance in the roadmap. It can be improved further though. In particular we would like to make sure that a path runs through every corridor between the obstacles. This can be achieved by applying some of the known techniques for finding paths in narrow corridors (see e.g. [6, 3, 8]). On the other hand, allowing cycles in the graph in a controlled manner can lead to the same result [16].

If a vertex v has only two neighbor vertices, it may be possible to directly connect those two neighbors, bypassing v . This would lead to longer edges and, thus, less (unnecessary) rotations. We remove these vertices if the merged edge is collision free and has enough clearance.

2.4 Circular blends

After retracting the vertices of the graph to the Voronoi diagram and after adding some minimal amount of clearance to the edges we still end up with a graph that consists of straight line segments. If an entity follows such a path it will have C^1 discontinuities in its motion at the vertices that cause sudden directional changes. In order to solve this problem we will replace parts of the straight line edges by circular blends.

The degree of a vertex is defined as the number of edges that is connected to this vertex. If a vertex has degree 1, it is an endpoint of a path segment, and no circular blend needs to be added. If a vertex has degree 2, the addition of the circular blend is straightforward. We find the centers of the two edges, and use these to create a circle arc that touches both edges. Now, we use this to replace a part of the path (see Figure 7(a)). If the degree of a vertex v is higher than 2, we find the centers of all incoming edges. We now add a blend for every pair of these centers (see Figure 7(b) for an example of a vertex with degree 3).




Figure 7: Creating the circular blends.

In the previous section, we retracted the edges of the graph to the Voronoi until they had at least some predefined clearance. Adding circular blends may decrease this clearance. Since we do not want the clearance to be lower than some predefined value, we check the minimum clearance of each circular blend. If it is too low, then we replace the blend by another blend that has a smaller radius. We repeat this until the blend has enough clearance. This procedure is shown in Figure 7(c).

Figure 1 shows an example of a typical roadmap graph created with this method. Computing this roadmap took about 1 second on a Pentium IV 2.4 GHz. Realize though that roadmaps can be computed during the creation of the scene and can easily be stored with it.

2.5 Answering path queries

Whenever an entity has to move to a new location, it can search the graph in real-time to plan its route through the environment. Because of the properties of the roadmap, paths will be short, have enough clearance from the obstacles, and are smooth. To obtain a path we first need to connect the current position of the entity to the graph. This can easily be done by finding the closest vertex in the graph and connecting the current placement to this vertex using circular blends. We proceed in the same way for the goal placement of the entity. Now we can use a shortest path algorithm in the graph to find the path between the current and goal positions. See Figure 8 for an example of such a path.

Computing paths is extremely fast. Even for a large roadmap graph consisting of 1000 vertices and 3000 edges, the calculation of the shortest path takes less than 10ms on a Pentium IV 2.4GHz.

3 Path planning for groups

Virtual worlds are often populated with a large number of moving entities. The entities should often behave as a coherent group rather than as individuals. For example, when one needs to simulate the behavior of whole army divisions. Current techniques solve the problem of path finding on the entity level, i.e. they plan the motion of individual entities, using techniques like flocking to keep the entities together. However, in cluttered environments this often leads to loss of coherence. There is no guarantee that the entities will stay together, albeit that 'staying together' is not well defined. Even though the entities all have a similar goal, they try to reach this goal without real coherence. This results in groups splitting up and taking different paths to the goal, for example as in Figure 2.

We will briefly describe a novel technique in which groups are guaranteed to stay together. More details can be found in [10]. We are given a virtual environment in which a group of entities must move from a given start to a given goal position. The entities must avoid collisions with the environment and with

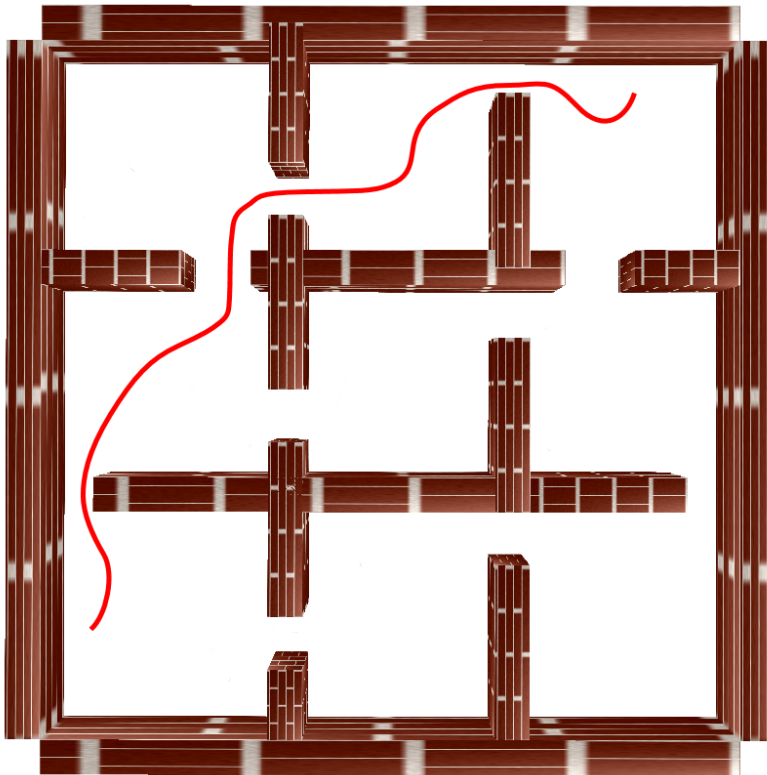


Figure 8: An example of a resulting path.




Figure 9: The corridor

each other, and should stay together as one group. The entities are modeled as discs (or cylinders) and are assumed to move on a plane or terrain. Later, the resulting paths for the cylinders can be used to animate avatars, e.g. sprites or motion captured human-like avatars.

The method works as follows: First, a so-called *backbone path* for a single entity is computed. This path defines the homotopic class used by all entities. Two paths P_0 and P_1 are said to be in the same homotopic class only if P_0 can be continuously deformed into P_1 without intersecting any obstacle. Next, a *corridor* is defined around the backbone path in which all entities must stay. Finally, the movement of the entities is generated using force fields with attraction points on the backbone path. By limiting the distance between the attraction points for the different entities, coherence of the group is guaranteed.

3.1 Backbone path planning and the corridor

The first phase of the approach consists of finding the backbone path. Since every entity should be able to traverse the path, the clearance on the path should be bounded by a some minimum value, namely the radius of the enclosing circle/cylinder of the largest entity. The backbone path can thus be defined as follows: A backbone path is a path in the 2D workspace, where the clearance at every point on the path is at least the radius of the enclosing circle/cylinder.

Although a minimum clearance of the radius of the enclosing circle is sufficient to find a path, we prefer a larger clearance, since a larger clearance leads to behavior that is more coherent. Also we prefer the paths to be smooth and short. Hence, the paths in the roadmap created with the method described above are very well suited for this application.

From the backbone path, a corridor is created. For this we use the clearance around the path. On every point on the path, the clearance is defined as the radius of the largest circle around this point that does not intersect with the environment. The value of the clearance is upper bounded by the maximum group width, i.e. the clearance can never exceed the maximum group width. The union of all the upper-bounded clearance circles forms a corridor around the backbone path. Figure 9(a) shows a path that is too close to obstacles, resulting in artificial narrow passages. In contrast, the path in Figure 9(b) was generated with the roadmap approach described above, and lies far from obstacles. The resulting corridor is much more natural.

3.2 Generating the motion inside the corridor

Once the corridor is created, we need to use it to generate the motion of the individual entities. The approach used is an artificial force field technique. Forces are defined that act on the entities and influence their movement. Every entity in the group has a corresponding attraction point on the backbone path. This attraction point is selected as the maximum advanced point p along the backbone such that the entity is still inside the circle centered at that point p with radius equal to the clearance at p (see Figure 9(c)). The attraction points make the entities move forward and keep the entities inside the corridor. The entities also repulse each other to avoid collisions between them. Additional forces could be incorporated, for example to accomplish formations.

3.3 Keeping coherence in the group


In order to keep the group coherent, the dispersion should be upper bounded. Due to the manner in which the corridor is constructed, the lateral dispersion (dispersion perpendicular to the backbone path) is automatically upper bounded by the group width. However, the longitudinal dispersion (in the direction of the backbone path) is not yet bounded in this approach. To correct this, the distance along the path from the least advanced attraction point to the most advanced attraction point is limited. This results in the entities in front waiting for the entities at the back.

3.4 Results


The behavior of the group can be controlled by adjusting the coherence parameters, *lateral dispersion* and *longitudinal dispersion*. Figures 10(a) to 10(c) show a group of 50 entities moving through an environment. In these pictures the lateral and longitudinal dispersion is varied, resulting in a longer, more stretched group (10(a)) or more compact group (10(c)).

Figure 10(d) shows the same group moving through the environment from the left lower corner to the right upper corner. The most advanced entities, i.e. the entities that passed the narrow passage earliest, wait for the last entity to pass the passage.


We tested the performance of the approach to show that the technique is usable in real-time virtual world applications such as computer games. For this, we developed a typical implementation. In this implementation we created numerous paths. The processor usage during the creation of the paths was minimal. For groups of 50 to 100 entities the processor usage did not exceed 1 percent. More efficient implementations could further decrease the processor usage.




(a) The group with high longitudinal dispersion and very low lateral dispersion.



(b) The group with medium longitudinal dispersion and medium lateral dispersion.



(c) The group with low longitudinal dispersion and larger lateral dispersion.



(d) The group moves through a narrow passage, making the first wait for the last ones (movement is from the lower-left to the upper-right).

Figure 10: A group of 50 entities moving in a virtual world. These paths and behaviors are created with the same approach, only by varying the parameters.




Figure 11: A speed diagram. The left image shows the path, the middle image shows the maximal speed allowed at each position. The right image shows the actual speed, taking acceleration and deceleration bounds into account.

4 Planning camera motions

Every virtual world has a camera through which the user views the world. Usually this camera moves depending on user input and place of action. A camera motion directs the camera from one position to another while controlling camera speed and view direction. There are many situations in which an automatic camera motion could be of great use. Think for example about an architectural walk through of a new housing project. The inexperienced virtual environment user (which are many future new home owners) would spend much time in the virtual environment just navigating, unintentionally neglecting the real goal of viewing their new future house. By automating the movement of the camera, by allowing the user to click on interesting locations, the user can enjoy the virtual worlds without the problems of getting lost or queasy.

The roadmap from the previous section can easily be used to steer a camera. Using this roadmap, the camera is guaranteed to keep a certain amount of clearance from the obstacles and the circular arcs make sure that the camera motion is gentle. The roadmap alone however is not enough to create a smooth camera motion. Camera theory [14, 26] shows that we need to take care of two more variables. First, the speed of the camera should be adapted according to the curvature of the path. Otherwise, objects will move too fast through the view, similar to the effect of a fast-forwarding movie. Second, the user should get cues about where the camera is going. In particular the viewer should be able to anticipate a camera rotation. We will resolve these issues in the next two sections.

4.1 Adapting the camera speed

Smoothness of the path is not enough for a smooth camera motion. The speed of the camera along the path should be adapted according to the curvature of the path. Also there should be a maximum acceleration and deceleration for the camera in order to prevent too abrupt speed changes.

Since our path consists of straight lines and circle arcs, we can adapt the speed of our camera by making use of the radius of the arcs. The smaller the radius, the lower the camera speed. When the camera leaves an arc with a small radius, we accelerate until we have reached the maximum speed of the current arc or straight line. If, on the other hand, the next arc requires a lower speed than the current camera speed, we must start decelerating on time, such that when we reach the next arc, our speed is sufficiently low. A speed diagram can be computed efficiently that satisfies both the constraints on the maximal speed for each arc and the bounds on acceleration and deceleration. See Figure 11 for an example of such a speed diagram for a simple path.

4.2 Smoothing the viewing direction

Intuitively one might think that the viewing direction should be equal to the direction of the camera motion. As stated before however, camera theory states that the user should be given cues about where the camera is heading or else the user may experience discomfort. We can achieve this by always looking at the position the camera will be in a short time. Experiments show that about 1 second is the right amount. Note that, as we fix the time we look ahead, the distance we look ahead changes depending on the speed of the camera. This is exactly what we want to achieve as in sharp turns we want to look at a nearer point than in wide turns. Looking ahead has another important effect. If we would look in the direction of motion and the camera reaches a circular arc, then it suddenly starts rotating at the start of the arc. Stated more formally, the rotation of the camera is only C^0 continuous. It can be proved that looking ahead solves this issue by making the camera rotation C^1 continuous.

4.3 Results

We implemented our approach in a walk-through system for virtual worlds. See Figure 12 for a screenshot. Rather than letting the user steer the camera directly, we display a map in the top left corner. By clicking on the map the user indicates the position she wants to move to. A smooth camera motion is then calculated in the way described above and the user can focus on the environment instead of the camera control. The processor time required for this is minimal. Experiments indicate that this is a pleasant way to inspect the environment.

5 Conclusions

In this paper we have described a new technique for automatic construction of high-quality roadmaps in virtual environments that can be used for computer controlled entities to quickly find their routes. We have also shown how this technique can be used to plan the motion for groups of entities and the camera through which we observe the world. Due to the nature of the roadmap, the approach also works in large environments, where grid-based methods typically fail (requiring too much storage and processing time).

We described our method as a 2-dimensional approach in which entities move on a ground plane. It is though easy to extend it to e.g. terrains and even motions in buildings in which the roadmap would automatically follow the corridors and stairs.

Roadmap construction is best seen as being part of the construction of the virtual world. It is easy to incorporate special requirements from the level designer. For example, the designer can add fake obstacles to force the path to e.g. stay on the sidewalks of a road. Also the designer can easily manipulate the roadmap graph by manually adding, changing, or removing nodes. Moreover, weights can be added to the graph to e.g. indicate preferred routes.

In some applications the virtual environment is not completely known at design time because it contains for example a door or a moving car. Also in such situations our technique can be applied by constructing the roadmap in the absence of the dynamic objects. During runtime, recent techniques for handling dynamic changes to the environments [23, 24] can then be used to take these additional objects into account.



Figure 12: An implementation of the techniques. The user can click on a location in the map at the left top and the program creates a smooth camera motion to that location.

6 Acknowledgements

This research was supported by the Dutch Organization for Scientific Research (N.W.O.). This research was also supported by the IST Programme of the EU as a Shared-cost RTD (FET Open) Project under Contract No IST-2001-39250 (MOVIE - Motion Planning in Virtual Environments). Part of this research has been funded by the Dutch BSIK/BRICKS project.

References

- [1] S. Baert. Motion planning using potential fields. gamedev.net, 2000.
- [2] R. Bohlin and L.E. Kavraki. Path planning using lazy prm. In Proc. IEEE Int. Conf. on Robotics and Automation, pages 521–528, 2000.
- [3] V. Boor, M.H. Overmars, and A.F. van der Stappen. The gaussian sampling strategy for probabilistic roadmap planners. In Proc. IEEE Int. Conf. on Robotics and Automation, pages 1018–1023, 1999.
- [4] M. Branicky, S. Lavalle, K. Olson, and L. Yang. Quasi randomized path planning. In Proc. IEEE Int. Conf. on Robotics and Automation, 2001.
- [5] Mark DeLoura, editor. Game Programming Gems 1. Charles River Media, 2000.
- [6] R. Geraerts and M. Overmars. A comparative study of probabilistic roadmap planners. In Algorithmic Foundations of Robotics V, Springer Tracts in Advanced Robotics 7, pages 43–57. Springer-Verlag Berlin Heidelberg, 2004.
- [7] C. Holleman and L. Kavraki. A framework for using the workspace medial axis in prm planners. In Proc. IEEE Int. Conf. on Robotics and Automation, volume 2, pages 1408–1413, 2000.
- [8] David Hsu, Tingting Jiang, John Reif, and Zheng Sun. The bridge test for sampling narrow passages with probabilistic roadmap planners. In Proc. IEEE Int. Conf. on Robotics and Automation, 2003.
- [9] P. Isto. Constructing probabilistic roadmaps with powerful local planning and path optimization. In IEEE/RSJ Int. Conf. on Intelligent Robots and Systems, pages 2323–2328, 2002.
- [10] Arno Kamphuis and Mark H. Overmars. Finding paths for coherent groups using clearance. In Eurographics/ACM SIGGRAPH Symposium on Computer Animation (2004), page to appear, August 2004.
- [11] L. Kavraki and J.-C. Latombe. Randomized preprocessing of configuration space for fast path planning. In Proc. IEEE Int. Conf. on Robotics and Automation, pages 2138–2139. IEEE Press, San Diego, CA, 1994.
- [12] L. Kavraki, P. Švestka, J.-C. Latombe, and M.H. Overmars. Probabilistic roadmaps for path planning in high-dimensional configuration spaces. IEEE Transactions on Robotics and Automation, 12:556–580, 1996.
- [13] F. Lamiroux and O. Lefebvre D. Bonnafous. Reactive path deformation for nonholonomic mobile robots. In IEEE Transactions on Robotics, page to appear, 2004.
- [14] G. Millerson. TV Camera Operation. Focal Press, London, 1973.
- [15] D. Nieuwenhuisen and M.H. Overmars. Motion planning for camera movements. In Proc. IEEE Int. Conf. on Robotics and Automation, pages 3870–3876. IEEE Press, San Diego, CA, 2004.

- [16] D. Nieuwenhuisen and M.H. Overmars. Useful cycles in probabilistic roadmap graphs. In Proc. IEEE Int. Conf. on Robotics and Automation, pages 446–452. IEEE Press, San Diego, CA, 2004.
- [17] C. Nissoux, T. Siméon, and J.-P. Laumond. Visibility based probabilistic roadmaps. In Proc. IEEE Int. Conf. on Intelligent Robots and Systems, pages 1316–1321, 1999.
- [18] M. Pinter. Toward more realistic pathfinding. gamasutra.com, March 2001.
- [19] C.W. Reynolds. Flocks, herds, and schools: A distributed behavioral model. Computer Graphics, 21(4):25–34, 1987.
- [20] C.W. Reynolds. Steering behaviors for autonomous characters. In Game Developers Conference, 1999.
- [21] S. Russell and P. Norvig. Artificial Intelligence: A Modern Approach. Prentice Hall, 1994.
- [22] W.B. Stout. Smart moves: Intelligent path-finding. Game Developer, October 1996.
- [23] Jur P. van den Berg, Dennis Nieuwenhuisen, Léonard Jaillet, and Mark H. Overmars. Creating robust roadmaps in changing environments. Technical Report UU-CS-2004-069, Institute of Information and Computing Sciences, Utrecht University, The Netherlands, 2004.
- [24] Jur P. van den Berg and Mark H. Overmars. Roadmap-based motion planning in dynamic environments. In Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems, pages 1598–1605, 2004.
- [25] P. Švestka and M.H. Overmars. Coordinated path planning for multiple robots. Robotics and Autonomous Systems, 23:125–152, 1998.
- [26] M. Wayne. Theorising Video Practice. Lawrence and Wishart, London, 1997.
- [27] S.A. Wilmarth, N.M. Amato, and P.F. Stiller. Maprm: A probabilistic roadmap planner with sampling on the medial axis of the free space. In Proc. IEEE Int. Conf. on Robotics and Automation, pages 1024–1031, 1999.