Large-scale mapping and navigation in virtual worlds.
RSMG Report 4

Katrina Samperi
k.samperi@cs.bham.ac.uk

November 22, 2010

Supervisors: Russell Beale and Nick Hawes RSMG Members: Richard Dearden and Bob Hendley

Abstract

This report covers the work done since the final submission of the Thesis Proposal on the 1st September 2010. It covers work done on the framework between Second Life and Java and the work towards using trails to aid with the generation of probabilistic road maps.

1 Project summary

The focus of this PhD is to enable intelligent mobile virtual agents to autonomously generate maps in large scale, dynamic and potentially unstructured environments. Once this map has been generated the agent should be able to use it to plan and navigate routes. This combination of factors in the environment makes the mapping problem significantly harder than work done previously and will bring mapping closer to a state in which the real world can be described by mobile agents over a long period of time.

Since the submission of the previous report on September 1st, the main focus of work has been on the agent framework and trail maps. We planned to submit a paper to the Autonomous Agents and Multi-agent Systems conference (AAMAS) about using trails to improve probabilistic roadmap (PRM) generation. This plan proved rather ambitious however and while some work has been done towards the final result, it was not in publishable form.

2 Agent Framework

The agent framework was the first step identified in allowing us to solve the mapping problem. We chose to use Second Life as the test environment as it provides both a controllable environment and supports the creation of online agents. We can create virtual agents that connect to Second Life through the LibOpenMetaverse Library (Open Metaverse Foundation, 2011). This library was created by reverse engineering the packages sent to the virtual world, and as such exhibits strange behaviour and very little documentation. Despite these issues, we have successfully created an agent that is able to log onto the world, move from one place to another and respond to basic chat commands. The agent is also able to gather information about objects in its immediate vicinity and report this back to the client for interpretation.

We implemented this using the Internet Communications Engine (ICE) middleware (ZeroC, 2011). This allows the C# library to be called from a Java client. We wanted this option as it would then be possible to integrate the agent with existing frameworks and take advantage of any previously written code. We looked at creating this framework using web services, as previous java clients have used (Pedley, 2011), but this option is limited in its ability to provide transparent communication to the server and stability issues.

After basic functionality of the agent was implemented, we turned our attention to interpreting -the objects present in the environment. Due to the complexity of how Second Life stores and and renders objects there is little understanding and almost no examples of how the library works. In general, if a third party client wishes to render the environment, it uses the code from the standard viewer, rather than attempt to understand the objects itself. At first we attempted an offline approach to object interpretation, with the gathering step being entirely independent of the processing stage. This led to problems understanding how the objects were placed and rotated in the world and so to solve this an online approach was implemented such that the objects is processed as it is received by the client and added to the object map immediately.

A very simple agent framework between the Second Life servers and Java now exists. It needs work done to make it stable and report all information to the client, rather than the limited subset that we focused on to try and generate trail and probabilistic road maps.

3 User trails and Probabilistic Road maps

The work on trails was moved forward to September as the AAMAS conference was announced and had a suitable track on virtual agents. In hindsight this was an overly ambitious plan, but it has led to rapid progress on the framework and trail generation.

3.1 The hypothesis

The motivation behind this paper was to improve the generation of probabilistic road maps by using user trails as an additional source of information. Trails are defined as the route users take through the world, gathered by observation over time. The information gathered can then be used to generate a free space representation of the world and routes through this without the need for an underlying metric map. By using this information it is believed that we can improve the generation time of PRMs by requiring fewer calculations to be done, and so generate maps faster, especially in very large environments. This then relates to the larger problem I am investigating by allowing faster map generation and navigation in large scale worlds and by allowing the generation of road maps to continue despite the entire world not yet being explored.

3.2 Method

To test this we need a metric representation of the world, a set of trail information and both the old and new version of the probabilistic road map algorithm.

The metric representation of the world was intended to be created by gathering information on the objects present in Second Life. This information could be used to decide which of these objects were placed at ground level and generate a simple object map of the environment. However, this step was underestimated, as even a simple environment is still a complex set of linked objects, rotated in three dimensions. A lot of work has been done trying to understand the geometry of the world and how second life presents primitives, but not enough time was given to allow a full map to be generated in any real virtual environment. The limited map does work in an entirely controlled environment created using OpenSim (OpenSim Project Infrastructure Fund, 2011), a second life-like simulator hosted locally. But this then does not allow for gathering trail information from other users.

A different approach to generating the metric map was considered using the trail information given. As this information contained where the free space was, an occupancy grid map could be created using only this information. This would then give bias to the trail based approach and render the experiment useless. By including additional free space to the trail map we were able to get a working probabilistic road map from the solution, however not with enough time to submit this to the conference.

The trail information was gathered over a twenty-four hour period by observing changes in the data reported by the library. Each avatar within a given region was identified and a constantly updating set of coordinates and timestamps was given. These could then be grouped to show routes through the world by linking points together from a single user. Time gaps of longer than five seconds were used to identify where one route started and another begun. Another filter was applied to remove any routes where the user may have been flying, this was done purely from the Z coordinate based on the average level of the land in the world. Finally, if the distance travelled was more than 30m (somewhat more than can be travelled on foot in a 5 second period) then the routes were assumed to be not linked, as it is likely the user teleported from one location to another.

The probabilistic road map algorithm was taken from the original paper Kavraki et al. (1996). This worked well for the occupancy grid map. No tests have yet been done on the proposed changes to this algorithm as we feel it is better to get a true map of the world, rather than the approximation. For the route planner we chose to use A*, as this is a simple planner that should always arrive at a good solution.

3.3 Evaluation

We can evaluate the success of the trail based PRM in several ways. The first method is to generate a full PRM for the world under three conditions:

1. Only the standard PRM algorithm is used
2. Only the trail information is used
3. A combination of the two is considered.

With this we can compare standard PRM in the environment (option one) with the information gathered only using trails (option two). If only the trail information is used then a road map should be generated in a faster time than standard PRM, but will only be able to discover routes where others have been observed. In a large scale world this will not be appropriate as there may be many areas that are only occasionally visited. Option three combines the previous two by allowing the algorithm to select a point either from the trail map or randomly generate one. When the algorithm selects a point from the trail map it is able to look up any routes and connect it to the road map allowing for faster generation still. By allowing the algorithm to also generate random points it will also be able to find new routes through the environment.

We propose to test this by selecting several start and end points, covering both short and long journeys through the world. The first measure would be the time taken to generate a road map for an ever increasing number of nodes. Option three should perform better than the other two in most cases, but especially when there are a large number of nodes to process. We can also change the number of points taken from the trails, rather than a unique point for the PRM algorithm. Another measurement would be how long the agent takes to plan the route from the start to the end point, and how successful these routes are. Measuring the success of the route would be done through visual inspection first, to check that the route is feasible, and then by length.

3.4 Problems and Future Work

The problems we faced in implementing this experiment all related to the complexity of the Second Life environment. After some time we came to the conclusion that it is highly unusual for any user created viewer to attempt to render the world itself, rather than relying on the standard client for this window. This leads to no in depth documentation and understanding of how objects are structured. We hope to solve the map generation problems, conclude this experiment and continue to write it up for publication in the near future.

4 Additions to the literature review

The focus of my reading in recent weeks has been related to probabilistic road maps. I started with the original paper that outlined the algorithm (Kavraki et al., 1996). The motivation behind probabilistic road maps is that they can be generated in any type of environment where a map of blocked regions is known. Later work has been done to try and negate this requirement such as the work presented by Kneebone and Dearden (2009) regarding uncertain environments. By modelling the PRM as a partially observable Markov decision problem the agent is able to decide which route is best to take through an environment despite not having full information at the start. This method also allows the PRM to be updated as more information is gathered, which will be useful for mapping dynamic environments.

Another interesting idea related to trails is presented by Dijoux (2010). They capture human movement through an environment using sensors attached to the human body and use this to generate a map based on current location and actions, such as opening a door. The resulting map looks very similar to the trail maps I have generated previously. Using trails for map building and navigation is an idea which is becoming more apparent in the literature.

As well as this I have been reading around robot navigation and mapping further. Jarvis (2010) uses an occupancy grid representation to model the world and path planning based on distance transforms from this point. The result is something rather like a vector field map (Schmidt and Azarm, 1992), in which the agent moves from one square to another based on the steepest descent trajectory to the goal point. Vu et al. (2011) also use a mapping system based on occupancy grids to hold maps about dynamic environments. Moving objects are detected by noticing changes (especially repeated changes) to individual objects.

5 Outline of any changes to the plan of research

Writing the paper linked to probabilistic road maps has been pushed back to later in the process, when the object map can be fully formed. Otherwise the plan for research is unchanged.

6 Poster presentations and publications

I took part in the Graduate School poster conference at the university in the summer presenting some of the work on trails. This was a good experience in talking to other people about my research and potential directions that it could take. I have made some contacts in the school of Psychology who are interested in how trails can be merged and discovered potential uses for trail use in other fields. After the successes of this day I was invited to present my poster again at the British Science Festival research showcase. This event gave me chance to talk to members of the public as well as academics from other departments about the work I hope to complete.

7 Next steps

Between now and the next RSMG meeting I hope to finish the experiment and paper regarding trails and their uses. The framework will be finished and tested to ensure that we gather all the information available and make this available for the client. The primary focus over the next few weeks will be on creating an accurate object map that is capable of functioning in complex environments. After this point we can implement the baseline system as described in the thesis proposal, adapting and using existing code where possible. Once this is in place we can compare the occupancy grid representation used by the base line system with the object map created and include the improved trail maps to this.

References

Frederic Dijoux. Mapping Indoor Environments Based on Human Activity. *Computer*, pages 476–481, 2010.

Ray Jarvis. Terrain-Aware Path Guided Mobile Robot Teleoperation in Virtual and Real Space. *2010 Third International Conference on Advances in Computer-Human Interactions*, pages 56–65, February 2010. doi: 10.1109/ACHI.2010.23.

L.E. Kavraki, P. Svestka, J.-C. Latombe, and M.H. Overmars. Probabilistic roadmaps for path planning in high-dimensional configuration spaces. *IEEE Transactions on Robotics and Automation*, 12(4):566–580, 1996. ISSN 1042296X. doi: 10.1109/70.508439.

Michael Kneebone and Richard Dearden. Uncertain Probabilistic Roadmaps with Observations. *Proceedings of the 19th International Conference on Automated Planning and Scheduling (ICAPS09)*, 2009.

Open Metaverse Foundation. Libopenmetaverse — open metaverse foundation, October 2011. URL `http://www.libsecondlife.org/projects/libopenmetaverse`.

OpenSim Project Infrastructure Fund. Main page - opensim, October 2011. URL `http://opensimulator.org/wiki/Main_Page`.

Thomas Pedley. Whisper - a light-weight yet feature rich second life chat client, October 2011. URL `http://whisper.slx.cc/`.

Gunther K Schmidt and Kianoush Azarm. Mobile Robot Navigation in a Dynamic World using an unstead Diffusion Equation Strategy. *Proceedings of the lEEE/RSJ International Conference on Intelligent Robots and Systems*, 1:642–647, 1992. doi: 10.1109/IROS.1992.587412.

Trung-Dung Vu, Julien Burlet, and Olivier Aycard. Grid-based localization and local mapping with moving object detection and tracking. *Information Fusion*, 12(1): 58–69, January 2011. ISSN 15662535. doi: 10.1016/j.inffus.2010.01.004.

ZeroC. Welcome to zeroc, the home of ice, October 2011. URL `http://www.zeroc.com/`.