Realistic Human Path Planning using Fluid Simulation

Lieutenant Colonel René G. Burgess
U.S. Army War College
650 Wright Avenue
Carlisle Barracks, PA 17013
717-245-3252
rene.burgess@us.army.mil

Christian J. Darken
Naval Postgraduate School
700 Dyer Road
Monterey, CA 93943
831-656-2095
cjdarken@nps.navy.mil

Keywords:
Path Planning, Route Selection, Avenues of Approach, Naturalistic Movement, Gas Diffusion, Fluid Simulation.

ABSTRACT: *This paper describes an approach for obtaining very realistic movement paths through a terrain set by applying the properties of a fluid simulation to produce intuitively human-like results. Similar to the concepts described in the physical world by the Principle of Least Action, realistic paths for human movement generally tend to follow "natural lines of drift." This common military term describes a method of route selection based on least effort expenditure (or highest possible speed) enroute to a goal (or destination). We implemented a fluid simulation (gas diffusion) as a means of determining these routes through a digital terrain set. This fairly simple technique generates what appear to be very realistic avenues of approach for large vehicle formations or for individual vehicles traveling at high rates of speed. This approach seems quite promising for modeling human movement tendencies and appears superior to classic path finding or optimal route selection methods in terms of representing human-like behavior.*

1. Modeling Avenues of Approach

In order to produce models and simulations which will realistically replace the functions of human beings, we must develop methods of encoding human thought patterns or behaviors. One particularly important requirement for improving military simulations is the ability to generate realistic avenues of approach, or likely routes to be followed by an attacking enemy formation of multiple vehicles. Individual path-finding or route selection algorithms are insufficient to this task, as they do not realistically represent the paths selected by a moving force.

Multi-vehicle formations exhibit many of the same characteristics as fluids as they move across the terrain. Whether these vehicles are the disciplined troops of a mechanized infantry battalion or the apparently unrelated commuter vehicles traveling into a metropolitan area, it is easy (for a human!) to observe the fluid-like characteristics of these formations as they travel from source to destination. This phenomenon is widely acknowledged by military authors and subject matter experts through their frequent use of fluid-based terms to describe the movement of forces across the terrain (e.g. "natural lines of drift," "flowed through the gap," "poured across the desert," "path of least resistance," etc…).

For our research at the Naval Postgraduate School, we had a requirement to develop multiple, realistic, natural, obstacle-avoiding, high-speed, and successful paths from a single source to a destination in order to create avenues of approach for analysis by a simulated intelligence officer [3]. We needed to find or develop a pathfinding algorithm which would produce these avenues through our digital terrain set, behaving in much the same manner as the flow of water down a dry, rocky stream-bed. With this analogy in mind we began investigating fluid-based simulations before finding a simplified gas diffusion model. With minor modifications, we've implemented this simple model to produce remarkably realistic avenues of approach.

The underlying principles behind using a fluid simulation as a means of pathfinding are not unique. We believe, however, that the application of those principles to the problem of producing highly realistic representations of human route selection tendencies is quite promising and worth consideration by current and future models of human behavior and tactical decision-making.

1.1 Related Work Using Simplest Path Approaches

Duckham's work on simplest paths proposes that humans prefer a simplest path algorithm to optimal path selection[8]. He proposes a simplest path algorithm designed to reduce the number of turns, stops, and general degree of difficulty in navigating an urban environment to create routes which appear more natural and human-like than some of the shortest path algorithms would produce. Duckham states that further cognitive studies are required to verify the algorithm, but concludes that this approach appears to present considerable advantages due to ease of description and navigation. Duckham's assertion, that easier routes are preferred, is in line with military terrain analysis, which emphasizes consideration of planning factors such as natural lines of drift, path of least resistance, and easiest route. Other researchers have noted similar human tendencies to drift, or move along routes with curved behaviors that consider the effects of velocity, turning radius, and speed [2].

1.2 Explanation of the Phenomenon – Principle of Least Action

In 1746 Maupertuis [10] developed the Principle of Least Action to describe the tendency of elements in nature to seek the minimal effort solution. Euler expanded upon this principle in 1748 with his assertion that a system of bodies at rest will seek a state which minimizes total potential energy, or effort [1]. In one of the linked pages from his series of idleness theory essays, Davis [7] describes the global nature of least action as follows:

> …Almost the whole of physics can be represented in geodetic form. Water running downhill seeks the steepest descent, the quickest way down, and water running into a basin, even one with irregular shape and bottom, distributes itself so that its surface is as low as possible… The Second Law of Thermodynamics requires that thermal systems change along a sequence of configurations, each having a higher probability of occurrence than the preceding configuration...

This paper implements the assumption that the Principle of Least Action can also be applied to human route planning and path selection. Without benefit of cognitive studies to confirm (or refute) this assumption, we implement a fluid-based simulation approach to modeling human movement tendencies when selecting avenues of approach through a terrain sample which we believe is superior (in terms of realism) to classic pathfinding approaches.

1.3 Modeling Gas Diffusion in Real Time

The concept of using fluids to represent human movement tendencies arose out of discussions on how best to model a moving, multi-vehicle element. These discussions repeatedly involved the use of fluid-like terms to describe the properties of a moving unit which led to the search for simple fluid models. Rather than delve into the detailed Navier-Stokes equations discussed in the textbooks [4], we turned towards a simpler solution for creating the same results.

Jos Stam [11] provided an efficient method of simulating movement in gas and smoke fields for use in visual simulations. Stam's simplified version of the Navier-Stokes gas exchange equations were instrumental in achieving the fluid simulation approach utilized in this paper.

1.4 Relationship to Potential Field Pathfinding

The idea of using fluid flow to model movement axes has a connection to previous work in using potential fields for robotic pathfinding logic [9]. In the special case of a pure gas diffusion, the density converges to a solution to the Poisson equation. Of course, it is therefore also a solution to the Laplace equation everywhere but at the sources and sinks. Connolly [5] proposed using Laplace equations as method of obtaining robotic pathfinding without the disadvantage of local optima. He defines a potential gradient for each cell based on the Laplace equations for each of its four neighbors. The solution of the Laplace equation is a harmonic function which is able to achieve the goal without getting stuck in local optima. In a second paper [6], Connolly discusses use of harmonic functions as method for robotic control, which, as he says

> …can be used to generate smooth, collision-free paths without the threat of spurious local minima.

Similarly, Svenson [12] discusses use of exact potential fields to generate avenues of approach before presenting his simplified approach using agents modeled on ant behavior to generate roughly comparable results to a potential field calculation of the entire terrain set.

2. Gas Diffusion in a Digital Terrain Set

One of the strengths of a gas-based terrain analysis is that one can assign several layers of trafficability information to a cell by simply adjusting the gas diffusion properties of that cell. This approach of classifying a database using multiple levels appears to be quite beneficial and extensible to multiple use cases.

For our research, we classified digital terrain based on percent change of slope as one of three categories: "Go" terrain suitable for fast vehicle movement, "Slow Go" terrain suitable for vehicle movement at reduced speeds, and "No Go" terrain unsuitable for vehicular movement. We assigned a maximum speed of 25 kilometers per hour (kph) to the Go terrain cells, 10 kph to the Slow Go cells, and less than 1 kph to the No Go cells.

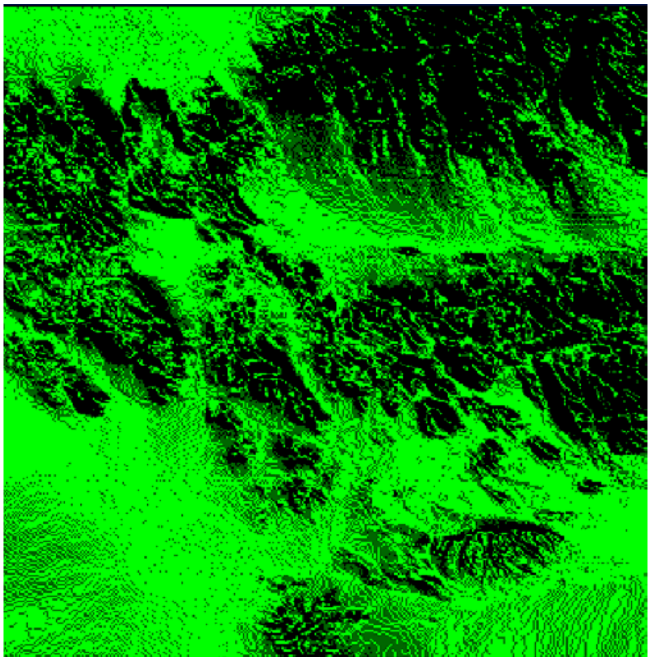

Figure 1. Terrain Map showing color-coded Go, Slow Go, and No Go areas for a mountainous portion of terrain at the National Training Center (NTC) at Fort Irwin, CA.

The mechanics of the gas diffusion model are quite simple. Each point in the data set is treated as a cell in the terrain model which begins with a gas concentration of 1, and with a maximum limit for diffusion to each of its neighbors. For simplicity, we treat each cell as a square, only able to exchange gas with each of its four shared-wall neighbors, thus the maximum exchange (out) for any given cell is 25% of its current volume.

The expected enemy approach direction (or an edge of the map zone) is treated as a source. Each source cell diffuses gas at the maximum rate with its neighbors and is reinstated to full capacity (index value of 1, or 100% full) at the beginning of every turn.

At the other "end" of the terrain set is the sink, or expected enemy objective area. Each sink is able to absorb gas at the maximum rate from its neighbors and immediately empties back to 0 at every turn.

For every turn, the model processes outgoing gas exchange for every cell with each of its neighbors based on the current concentration of gas in that cell times the maximum exchange rate of the cell. Based on the relative assigned velocities (above), we used maximum exchange rates of 25% for the Go terrain cells, 10% for the Slow Go terrain cells, and 0% for the No Go terrain cells. The exchange rate is limited by the minimum exchange rate possible at the two cell edges. A Go cell only exchanges 10% of its volume along the edge that it shares with an adjacent Slow Go cell while it may also exchange 25% of volume on the edge shared with another Go cell.

Outgoing gas changes are stored in a buffer until all outgoing calculations are complete, then the results, both outgoing and incoming gas amounts, are posted to each cell, the sources are regenerated to 1, the sinks are emptied to 0, and the process is repeated. As the process continues, gas concentration in the cells closest to the sink quickly approaches 0 and those cells serve to drain the cells around them. This process continues throughout the duration of the model until equilibrium. Evaluation prior to achieving equilibrium is visually interesting but practically irrelevant, as the model must be at (or close to) equilibrium to tell us where the fluid (or gas, in this case) will truly flow.

The true benefit for evaluating movement occurs when the gas reaches a steady state approaching equilibrium. As the difference in gas exchanged (between any two cells) for each cycle reduces to the point of minutely small changes for the entire database, the field of gradient vectors becomes useful for analysis. Though not true equilibrium, as the amount of gas exchanged approaches zero we have a state of nearly unchanging gradient vectors upon which to base our route estimation.

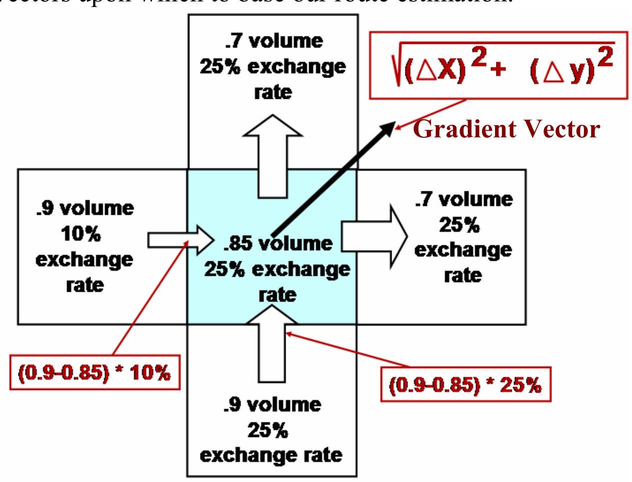

Figure 2. Gas Diffusion calculations for each cell, with the Gradient Vector (and associated Gradient Magnitude) represented by the changes in x and y.

Of course it's possible to compute near equilibrium mathematically, but for our research we implemented an interesting visual method for determining both the degree of diffusion (as it approaches a state of near-equilibrium), as well as providing a means of observing fields of high density flow (such as the movement of water through a narrow section of streambed) by mapping gradient magnitude. Gradient magnitude for each cell is computed by storing the total changes in x and total changes in y for each cell and treat these changes as the components of a force vector at each cell.

The gradient vectors below indicate both the direction of the vector from the center of each cell, but also the magnitude. Gas flowing through the convoluted area on the top of the picture moves much more slowly than that to the bottom of the picture, as can be seen by the size of the gradient arrows. This is both intuitive and informative.

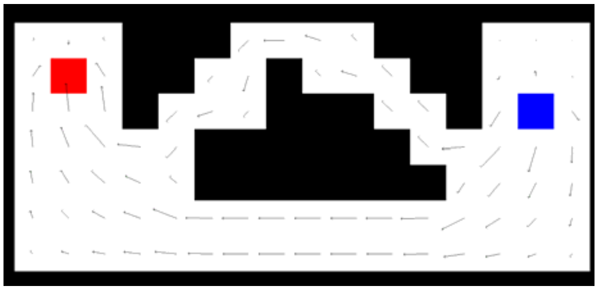

Figure 3. Gradient Vectors from source to sink.

The gradient vectors for each cell form a vector field which will become quite useful in our analysis of routes. Shown below is the gradient magnitude view of our terrain field where each cell's gradient magnitude (as a percentage of the highest magnitude vector in the field) determines the brightness of the red color.

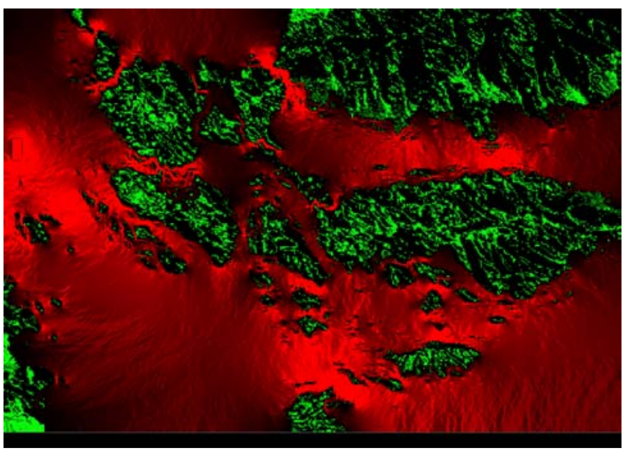

Figure 4. Color Coded Gradient Magnitude Levels for terrain at the NTC. Source is on the right, and the objective (or sink) is on the far left of the picture.

This provides a very intuitive picture to the user of the gas flow over time. It's quite clear by the bright areas on this frame which pieces of the terrain experience the highest gas transfer rates, and where the areas of concentrated high vectors exist. This information can be then used to pinpoint areas of risk to a moving force as they become more physically concentrated, key terrain (as it overlooks these places where the moving force assumes higher risk), and potential anchor points for a defensive deployment.

3. Producing Avenues of Approach

The process of producing avenues of approach from this point (given a gradient field produced as described in the previous section) is similarly simple to implement but produces an excellent representation of avenues of approach through a terrain sector. Beginning at the source cells, we generated routes at each possible adjacent cell. By design, this limited routes to a fixed number for ease of computation and viewing. We adjusted the number of routes by increasing or decreasing the size of the source. Continuing forward towards the sink, we interpolate the next particle position based on the influence of each of the closest gradient vectors proportionate to their strength and distance from the simulated moving particle. This process is displayed in the figure below.

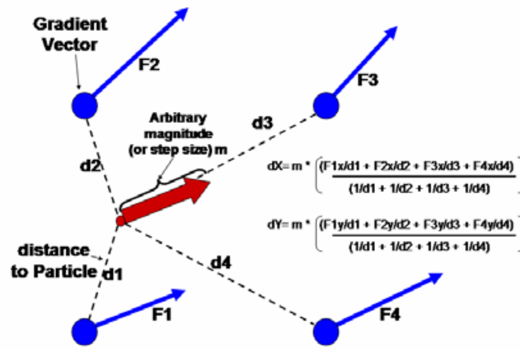

Figure 5. Particle Flow Interpolation.

This process is repeated until each particle has reached the closest sink. Interpolation is a necessary step in order to represent a realistic movement through the gradient field. In the diagram above, step size is listed as an arbitrary magnitude m. We commonly used a magnitude of 0.4 (or 40% of the distance between posts), as this forced at least two interim checks between each set of influencing vectors. The route selector is programmed to consider all vector posts within a distance of 1.5 cells from the particle. This ensures that each particle always received input from a minimum of four vectors as well as considering both past and future vectors when transitioning between one block of four vectors to the next.

One interesting result of this approach is the effect of edges upon movement of particles. Rather than force particle movement to never intrude upon the No Go terrain cells, instead we treated No Go cells as ones with a zero magnitude gradient vector. As a result, then, we experience a form of limited edge effects when particles travel adjacent to No Go terrain as their forward movement (to the sink) is slowed with only half as many gradient vectors to propel them forward. These edge effects also produce occasional visual anomalies for small terrain grids as the routes (especially the slower routes to the sink) appear to travel within the bounds of the first No Go cell adjacent to a Go or Slow Go terrain cell. It is worth noting that this is a purely visual anomaly produced by our representation of point values (the gradient vectors for each data point) as if they were uniformly distributed across an entire grid cell. In fact, the changes between data points are continuous in real life, and allowing routes to drift into the first cell of No Go terrain grid is akin to the physical nature of waves lapping against a rock before returning to the greater stream.

For our research, we chose to model gas diffusion in two dimensions only, treating the changes in terrain elevation as factors influencing speed of travel (or diffusion), but not specifically modeling a 3D gas. Similarly, it is obvious that a precise physical model of fluid movement through a complex grid would include a formal treatment of edge effects, turbulence, coefficients of drag, and fluid density. We believe that using a more complex fluid state or gas diffusion model than the one described here would result in similarly realistic results, albeit at a higher computation cost. In our implementation, the net effect of interpolating particle movement with zero magnitude edge vectors results in slowed particle movement along exposed No Go terrain and faster moving particle routes which curve around obstacles. These effects seem consistent with our observations of human movement patterns. Both the decision to forego a full mathematical treatment of fluid movement properties and our acceptance of de facto edge effects in this simple model appear to be reasonable for our task.

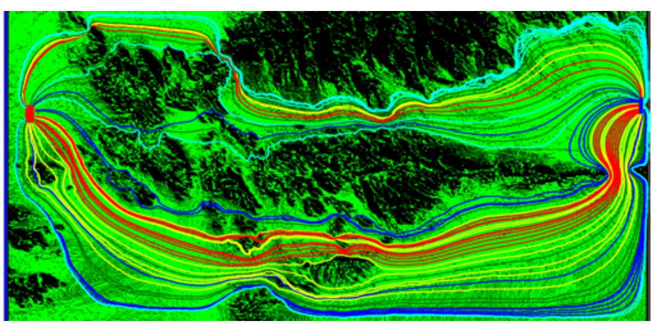

Figure 6. Avenues of Approach sorted by color against Go/Slow Go/No Go background (NTC Terrain).

Avenues of approach through our terrain are evaluated based on speed of particle movement through the grid in total time from source to sink. Shown in Figure 6 are the avenues of approach generated for a portion of the terrain database discussed above. For ease of visual interpretation, all routes depicted are displayed by color ranking based on relative time to complete the movement. The fastest twenty percent of routes are colored in red, followed by yellow, green, blue, and cyan.

Note that this approach describes most likely avenues of approach from a certain point and is used as a means of generating favorable terrain indices for use by a defensive positioning algorithm. As such, the routes are only calculated one time from an estimated source location. Each route from this analysis receives a score relative to its likelihood of use. We use this score as a reward function for related work considering potential defensive deployments.

In an alternative implementation, one could also use gas diffusion to provide a continuous assessment of an attacking enemy force rather than considering routes as a means of predicting the location of a defending enemy force. The initial results would be identical to those presented here, but at each chokepoint or decision point we'd want to reestablish a new gradient field based on the decision point as source and predict further routes into zone from that point. In the U.S. Army, this recursive application of terrain analysis is applied throughout the battle by intelligence officers and would need to be mirrored in software in order to replace some of these functions in a simulation.

4. Implementation Issues

For the large NTC terrain data set prevalent in many of these pictures, the terrain data set consists of 720 x 360 elevation posts – or 259,200 separate DTED Level II elevation posts, each with its own coefficient of fluid transfer. While this volume of data makes the fluid simulation very sensitive to the specific nature of the terrain (note the subtle folds in the terrain from Figures 1 and 4), it also has the net effect of reducing an elegant, simple solution to a tedious computational exercise as we were required to run the fluid exchange program for hours in order to obtain a fluid system which has achieved equilibrium (and can thus provide reasonable avenues of approach).

Coarsening the grid by averaging the values for a block of terrain posts saves considerably on computation time and appears to provide a reasonable estimate for the most likely avenues of approach; the red lines representing the top 20% of fastest routes were remarkably similar to those of the fine grid solutions. Estimating secondary avenues of approach using a coarsened grid is a challenge, however, because the terrain grid will now misrepresent smaller terrain features in potentially misleading ways. This is less of a problem when averaging Go and Slow Go terrain, or at edges between long stretches of Go and No Go terrain, but problems arise at places which were formerly classified as impassable but now appear to be simply Slow Go. Less likely to occur in the high speed avenues of approach, it's almost guaranteed to occur in the slower speed avenues of approach (where the coarsened grid is based on the average of several No Go and Slow Go terrain cells).

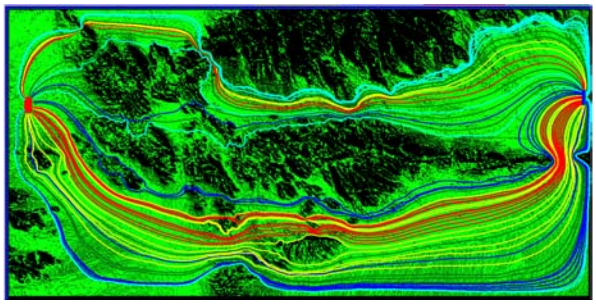
Figure 7. Fine Terrain Grid Avenues of Approach.

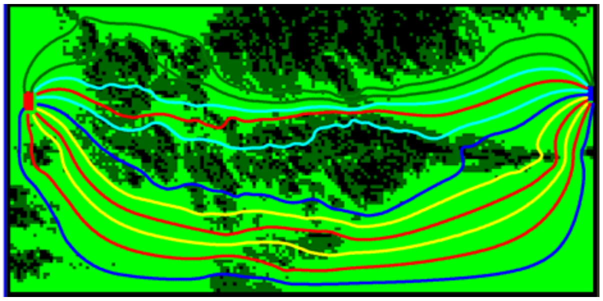
Figure 8. Coarse Terrain Grid Avenues of Approach.

To appreciate this difference, compare the results of a fine terrain grid's avenues of approach above with those obtained by a 1/16th size terrain file (each cell representing 4 x 4 elevation posts from the original data set). Note how the preferred avenues of approach (bottom half of each figure, in red) are very similar, but the secondarily colored lines have some interesting discrepancies in the left-most third of the figure as they traverse the congested terrain at the end of the long valley.

It's reasonable to assume that there are suitable techniques for decreasing the amount of time to achieve a uniform fluid distribution based on the distance of each cell from the source and sink. This would produce a very rough approximation of dispersed fluid volume, but perhaps reduce the time to produce equilibrium with consistent gradient vectors considerably. Using DTED Level 1 data would similarly reduce the data set to a more manageable size, but would be subject to the same loss of subtlety evidenced in the coarse terrain grid. Balancing processing speed with degree of required resolution warrants further study prior to an implemented solution and, though identified here, is not further investigated in this paper. Using the result of the coarse simulation as an initial condition for the fine one is an obvious and easy to implement approach. Because of the equivalence to linear system solution noted above, there is no end of techniques to try here.

5. Evaluating Avenues of Approach

Further study of human path selection is required in order to confirm that the intuitive benefits of the results obtained above are supported by human experimentation. This comparison was not conducted prior to completion of this paper, but is acknowledged to be required before the fluid simulation can be categorically proclaimed an accurate representation of human path selection or wayfinding tendencies. Diffusion-based avenues of approach are compared to an A* search with generally desirable results.

5.1 A* Search Comparison

Comparing the results of fluid simulation optimal routes with those produced by an A* search pattern confirmed that diffusion-based routes are, as suspected, not optimal traversals of the terrain grid in terms of pure time to traverse each cell between source and sink.

We compared a number of small terrain samples using both A* and diffusion method with the cost factor of time to traverse from start to finish. Time was computed based on the distance between nodes divided by the maximum allowable speed at the next node. A* calculations used a cost function of cost to date (along the given set of previous points) plus a heuristic of straight line distance divided by maximum speed. We noted an average of 13% more time to traverse the diffusion based routes in addition to considerable variability between optimal routes (A*) and the most likely (diffusion) routes. The highest variability occurred for those sample sizes which we created with multiple twists and turns along the shortest path and wide paths to the outside. This is quite similar to the route through a major city: technically faster by driving more direct path through the city given no traffic delays, but with a preferred route on the high speed beltway around the city.

Comparing A* search pattern results with diffusion-produced optimal routes would seem to indicate that the A* search produces the "best" path between two points. As discussed earlier, we believe that if the definition of "best" is "most realistic," then clearly the fluid simulation produces better results because of the human tendency to select more flowing, natural routes over mathematically superior shortest path (or even fastest path) routes.

Further analysis of this comparison reveals another flawed assumption: the above comparison assumes that our fastest path calculations included all appropriate costs. On further reflection, however, it's clear that we've ignored the physical costs (in terms of time) for changes in direction, stops, starts, and time to achieve maximum speed. We cannot achieve instantaneous speed changes and similarly do not make changes of direction at constant speeds. A more detailed comparison of the A* routes versus diffusion-based avenues of approach should include cost factors for changes of direction, changes in maximum speed, turns, and stops required. These costs are significant in the real world, and if modeled in our cost comparison would surely lead to a reduced gap between the optimal mechanical solution and the optimal diffusion-produced fluid simulation solutions.

5.2. Considering Military Units versus Single Vehicles

Although we believe that this approach produces realistic routes for both individual vehicles as well as larger military units, it must be noted that the differences in movement behavior between a single vehicle and a larger unit are substantial. The inertial properties of a large group of vehicles (acceleration, turning speed, turning radius, maximum speed, time to decide on route change, and deceleration, etc.) are believed to weigh significantly in favor of a fluid-based simulation as method of computing routes. Single vehicle properties can be accommodated by more frequent estimation of possible routes, as their increased agility makes the likelihood of route changes much greater. This is similar to the method one would use for considering offensive movement through zone as discussed earlier.

5.3.1 Considering Unit Width

Military units travel in a number of different formations based on size of the unit, likelihood of enemy contact, and desired speed. In general, military units will compress to a width narrower than the described "normal" or doctrinally prescribed width for short periods of time in order to achieve a high overall traversal rate. If we assume that the unit *must* travel from the source to the sink, then the diffusion-based most likely avenue of approach is valid *regardless of preferred unit width* because the diffusion-generated most likely route is also the route which supports the highest total flow volume from source to sink.

We believe that this is an important aspect of using diffusion as it more closely models movement of multiple vehicles across a complex terrain set with multiple small no-go terrain sectors interspersed throughout a larger maneuver corridor. Powell [12] describes a model for movement over terrain that makes a strict division of all terrain into go and no-go regions. Large units can have the set of possible routes filtered by minimum width. In effect, the unit is being modeled as a hard disk that can either fit through a given route or not. Slow-go terrain must be forced into one of these two categories. Regions that are mostly good terrain but speckled with no-go terrain must be modeled as entirely go, neglecting the effect of the no-go spots entirely, or the consequence will be that they will appear impassable to large units. This is a stark contrast to the fluid model, which predicts gradual degradation in the attractiveness of a route as the result of slow-go, or speckled no-go terrain.

Some of the most unlikely routes generated by our work (the slowest routes indicated in cyan color from the included figures) would include "impossible" or improbable routes for units of greater than single vehicle size. Excluding these routes from consideration, however, is not recommended as these routes still represent realistic methods for passing from one point to another. Just as a stream will continue to flow through a rocky section if the smooth streambed is too far out of the most efficient path, so too will military units temporarily disperse into smaller movement corridors in order to increase overall movement efficiency.

6. Conclusions and Future Work

Several areas of future research are indicated in this direction. First, future cognitive research is indicated to confirm the specific results obtained by a fluid flow simulation as an approach for modeling military movement through terrain. This research should include investigation of true rates of travel for multiple vehicles over good terrain interspersed with obstacles at fine granularity. Second, speed enhancements in calculating near equilibrium are required to produce the vector field in real time. Finally, further study to compare the results of this simple model with a more detailed treatment of edge and boundary effects in multiple terrain types and sizes. We believe that more complex fluid models would not significantly improve the utility of vehicle movement modeling, but this belief bears further investigation. Each of these questions should be addressed by future work in this field.

Fluid simulations appear to provide an excellent tool for modeling human movement tendencies through terrain. The avenues of approach and routes created by seeding digital terrain sets with simulated particles and tracking their movement from source to sink are natural in appearance, reflect multiple levels of difficulty in an elegantly simple fashion, and seem to provide a great resource for blind computer algorithms to represent the complex functions of the human eye with respect to seeing terrain as an entity rather than as a set of data points. Future work is indicated in refining this tool and reinforcing the cognitive assumptions upon which it is based, but at a minimum the fluid based approach described previously warrants serious consideration as a technique for modeling human movement.

The flexibility of this approach, combined with its fairly simple mathematical foundation, indicate that this model also could be applied to a large set of problems with varying degrees of evaluation or certainty and provide realistic approximations of those environments for minimal computational cost. This model provides an excellent approximation of the effects of a real world physical process. The laws of physics which bound gas diffusion are the same laws which bound movement of any object in the real world, namely the principle that objects (in this case humans) seek a predictable path of least expended effort (or least resistance, or least action, etc..) as they transition from one state to another. In the authors' opinion it is because of the global applicability of these laws that we can apply this simple model to very complex environments and achieve similarly valid predicted outcomes.

7. Acknowledgement

This work was supported by funds from the Naval Postgraduate School, the Navy Modeling and Simulation Management Office, the U.S. Army Training and Doctrine Analysis Center, Monterey, and the Space and Naval Warfare Systems Center, San Diego.

8. References

[1] Beeson, D. Maupertuis: An Intellectual Biography, The Voltaire Foundation, 1992.

[2] Brogan, D.C. and Johnson, N.L., "Realistic Human Walking Paths," In Computer Animation and Social Agents (CASA), p. 94-101, 2003.

[3] Burgess, R.G., Realistic Evaluation of Terrain by Intelligent Natural Agents (RETINA), Masters Thesis, Naval Postgraduate School, Monterey, California, September 2003.

[4] Chorin, A.J., and Marsden, J.E., A Mathematical Introduction to Fluid Mechanics, 3d ed., Springer, 2000.

[5] Connolly, C.I., Burns, J.B., and Weiss, R., "Path Planning Using LaPlace's Equation," In Proceedings of the 1990 IEEE Conference on Robotics and Automation, pages 2102-2106, May 1990.

[6] Connolly, C.I., "Applications of Harmonic Functions to Robotics," In Proceedings of the 1992 International Symposium on Intelligent Control. 26 May 1992.

[7] Davis, S., "Idle Theory, The Principle of Least Action,"[http://ourworld.compuserve.com/homepages/cuius/idle/evolution/ref/leastact.html], September 2003.

[8] Duckham, M. and Kulik, L., "Simplest Paths: Automated route selection for navigation," Accepted for COSIT'03, Lecture Notes in Computer Science, Springer-Verlag, 2003.

[9] Krogh, B. and Thorpe, C. "Integrated Path Planning and Dynamic Steering Control for Autonomous Vehicles." In Proceedings of the IEEE Conference on Robotics and Automation, pages 1664-1669, 1986.

[10] Maupertuis, P.L.M. de. Oeuvres. Dresde, 1752.

[11] Stam, J., "Interacting With Smoke and Fire in Real Time," Communications of the ACM, pages 77-83, July 2000.

[12] Powell, D.R., and others, "Representations to Support Reasoning on Terrain," Los Alamos National Laboratory report, LA-UR-88-3752, 1988.

[13] Svenson, P., and Sidenbladh, H., "Determining possible avenues of approach using ANTS," In 6th International Conference on Information Fusion, pages 1110-1117, Cairns, Australia, 2003.

Author Biographies

RENÉ BURGESS is an Army Lieutenant Colonel currently stationed at the U.S. Army War College in the Center for Strategic Leadership where he serves as Deputy, Strategic Experiential Education Group. While there, he continues his PhD Studies at the Naval Postgraduate School in Modeling, Virtual Environments, and Simulations (MOVES), where he recently obtained his Master of Science. From 1998-2001 he worked at the National Simulation Center in Fort Leavenworth, KS, where he was the first Chief, Digital Battlestaff Sustainment Trainer (DBST), a live-virtual-constructive simulation and stimulation federation currently fielded throughout the U.S. Army.

CHRISTIAN DARKEN is an Associate Professor of Computer Science at the Naval Postgraduate School, where he also collaborates intensively with the MOVES Institute. Previously he was Project Manager of the Decision Support Systems project and Senior Member of Technical Staff at Siemens Corporate Research in Princeton, NJ, where he was variously associated with the Learning Systems, Adaptive Information and Signal Processing, and Software Engineering Departments. He received his Ph.D. in Electrical Engineering from Yale University in 1993, and previously received the M.S. and M.Phil. in Physics from the same institution.