Optimizing Traffic Signal Control for Emergency Evacuation

Matthew Durak and Erik Goodman

Michigan State University, East Lansing, MI 48823

Introduction

Agent-Based Modeling (ABM) is an effective strategy for simulating the emergent properties of complex systems. This form of simulation models the behaviors of individual agents interacting with each other and with their world. ABM is well suited for problems such as modeling behaviors during emergencies, or critical incidents. These situations can range from getting people out of a burning building to routing traffic away from a chemical spill. Many frameworks exist for creating ABMs—one such program, used here, is called NetLogo [1].

While ABM works well for modeling individual behaviors to determine the resulting emergent group properties, it is difficult to find optimal strategies for emergency responders to follow in order to save lives, faced with such scenarios. Parameters must be tuned and tested to find optimal strategies for handling various kinds of incidents.

Furthermore, in order to test each possible strategy, one or more simulation runs may take a very long time. It becomes impractical to test the many combinations of parameters for possible strategies in an exhaustive manner, so a more efficient search method must be used. In this case, it is advantageous to use an optimization tool with evolutionary and hybrid computation, such as HEEDS Professional.

While most applications of HEEDS focus on design of physical parts and products, HEEDS can also be applied to many other forms of optimization. This paper discusses the success of HEEDS in the area of traffic control for urban evacuation, using agent-based modeling.

Traffic Model

The agent-based traffic model is written in NetLogo. It was originally based on the "Traffic Grid" model included in NetLogo [2] but has been heavily modified. The premise of the model is that cars drive around in a city, going to destinations determined at random before and during the simulation. After some time, an emergency occurs on one side of the city (for example, a chemical spill) and drivers are instructed to leave the city in the other direction. The model focuses on only a portion of the city, represented as a grid of eight "streets" by eight "avenues." Cars can enter and leave the visible portion of the city freely because the world is normally treated as toroidal (cars leaving one side reappear on the opposite side, when no evacuation is underway). This behavior only changes during an emergency, when cars leaving the edge of the simulated area that is the designated emergency direction actually "stick" at the dangerous edge of the area to be evacuated, and must come back to the visible eight-by-eight world to complete their evacuation.

The eight-by-eight urban area has 1500 cars driving during normal conditions (Figure 1). Each car is an individual agent with variables and behaviors. These cars have random goals, which are represented by going to an intersection or leaving the visible area in a certain direction. The model is simplified so that cars can only drive on the roads, and there are no parking lots for them to enter or leave. Cars each have different target driving speeds, normally distributed around the speed limit.

When the emergency occurs, only half of the cars will immediately begin to obey the commands to leave the city. This models the behavior of people who would ignore orders in such a situation because they need to pick up a family member or friend first, and some people who simply would not immediately become aware of the emergency directive.

Parameters may be set in the model to create additional cars during the emergency which will emerge out of the dangerous side of the city. These cars represent the people who are not on the visible portion of the map but must still evacuate.

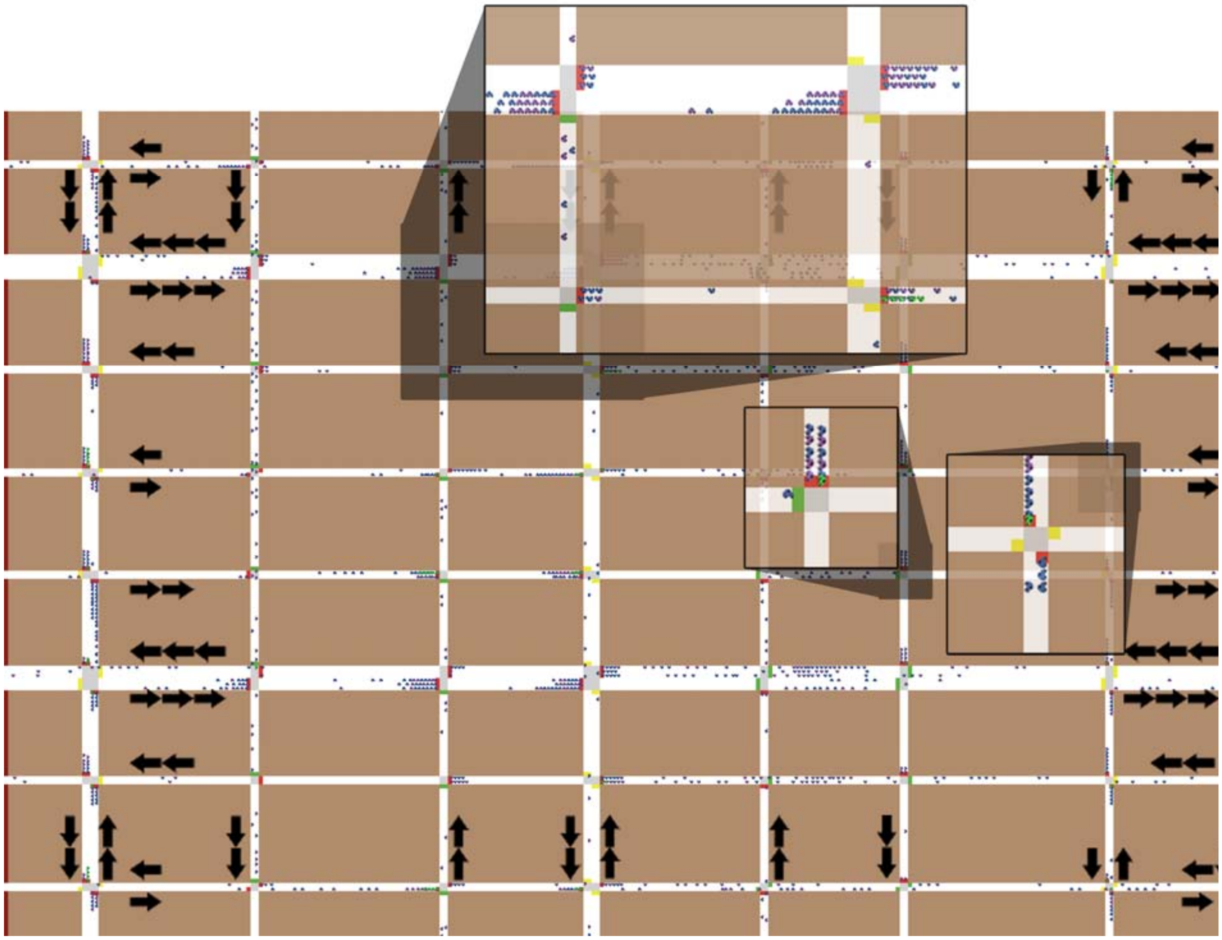

Figure 1: Traffic model in NetLogo with arrows indicating direction of traffic and number of lanes. Some portions of the map are enlarged for better viewing.

Officials can attempt to assist drivers in the evacuation by changing the signal timings of traffic lights when the emergency occurs. Finding an optimal signal control strategy is the task presented to HEEDS.

Chromosome Design and Optimization Goal

The traffic model uses an input file with instructions for how traffic lights should be timed when an emergency occurs. The controls can set the length of a green light, including time for the yellow light, for both the horizontal and vertical direction, and a positive timing offset from the lights to the left of it or below it, so that lights can be synchronized (Figure 2). In HEEDS, this is set up so that each horizontal row of streets has the same time for vertical green lights and a second (different) time for horizontal green lights.

Then, on each of these roads, a "unit offset" of each light's start of green is set, offsetting it from the previous light (to its left) (Figure 3).

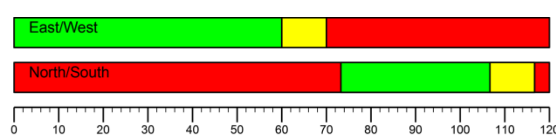

Figure 2: Example of traffic light timing. The top bar represents the color of the light for traffic going east or west while the bottom bar is for the north and south traffic.

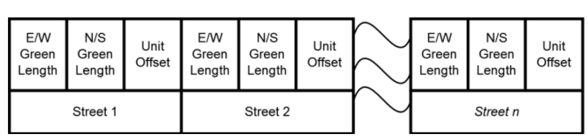

Figure 3: Layout of the chromosome in HEEDS.

This unit offset on the chromosome is multiplied going right to left in the model. For example, a unit offset of 25 would set the offset of each intersection going right to left as 0, 25, 50, 75, ... 25 * n. In this way, lights to the left will begin at a point in time further in the cycle so that, ideally, by the time cars driving to the right reach the next light, it is still green. This left-to-right synchronization is chosen because that is the direction that will be used for evacuation (to the right). This signal timing will synchronize to the desired pattern regardless of the light timings during non-emergency operation. Roads are grouped by row because cars evacuate from left to right so that, ideally, these traffic lights will be timed to let cars hit as few red lights as possible while preserving reasonable flow of the traffic in the other directions.

The goal is to keep traffic smoothly flowing out of the city. The value that HEEDS attempts to minimize is the time it takes to evacuate a specified fraction of the cars plus the number of cars that did not yet evacuate. In the current study, the specified fraction is 93% of the cars, because after that point, there is minimal congestion, and delay in departure is only because some cars have been delayed by intermediate stops they have made. (The "penalty" for cars remaining unevacuated is to allow for some comparisons when changing the fraction from 93% to some other value. Also, if an evacuation runs too long, it will stop, so the count of cars remaining provides some comparison of designs that went over the time limit.)

HEEDS Multi-Agent Design

This ABM optimization problem is run on the High Performance Computing Cluster at Michigan State University. HEEDS is set up to use the SHERPA method for parameter optimization. Evaluation of each potential solution requires running of a significant ABM simulation, so effort must be made to minimize the number of total evaluations needed. SHERPA can automatically distribute the evaluation of individual evacuation strategies over the machines in the cluster to significantly reduce the time to solution. In order to assist HEEDS even further in finding optimized designs as quickly as possible, the problem is set up to use multiple search agents (Collaborative Independent Agents). The idea is to have one HEEDS agent use a simpler (hence faster) simulation model, which then passes on its best designs as they are found to the next agent group, which uses a higher-resolution specification of the solutions and runs each simulation for a longer time.

In the simpler agent, one cannot argue that an agent that starts off evacuating quickly will necessarily end up evacuating well, but one *can* argue that one which starts off evaluating poorly will *not* do well overall. Therefore, the best strategies identified appear to perform very well, when really they are poor in other conditions. NetLogo allows simulations to have a seed set for the random number generator so that these runs can be repeated with the same parameters. In order to find optimal solutions, the designs must work well when tested numerous times with different random number seeds. This is achieved in HEEDS by using a stochastic parameter value for every evaluation. This way, HEEDS will take the better designs, and reevaluate them to determine if they work well under all conditions. This favors designs that are more robust because they work regardless of random variations.

The current setup for this traffic model uses one "coarse" agent in which the simulation continues only until 75% of the cars have evacuated. This is enough to show if a design is performing well, while allowing a shorter execution time so that more evaluations can be done in a given total execution time. The variables have a resolution that allows multiples of ten changes in the values of the light-timing parameters. These numbers are in simulation time units, in which one step is approximately one-half second.

This agent passes its three best designs in every optimization cycle to a more complex agent. This "intermediate" agent runs each design until 93% of the cars have evacuated. This is most of the cars, ignoring some cars that would otherwise still require significant simulation time to leave, while relatively unobstructed by traffic. Also, this agent uses more refined resolution of variables, stepping through multiples of two, which represents one second of real time. This agent then passes its three best designs to a longer running agent.

This final "extended" agent has a longer evacuation time, because in addition to the 1500 cars on the roads, 1500 more cars are introduced to the grid during the emergency from the side of the grid in which the emergency has occurred. This represents additional cars that would need to transit the emergency area in order to evacuate. It is a more realistic measure because the flow of traffic remains steady for a longer time before concentrating near the exit side.

This agent group takes the most computer time per evaluation, but also provides results representing the most taxing traffic conditions. Each of these agents is set up in a different agent group in HEEDS so that they can work in parallel. The HEEDS agents also generate multiple designs to run in NetLogo in parallel (Figures 4 and 5), so long as multiple processors per HEEDS agent are available.

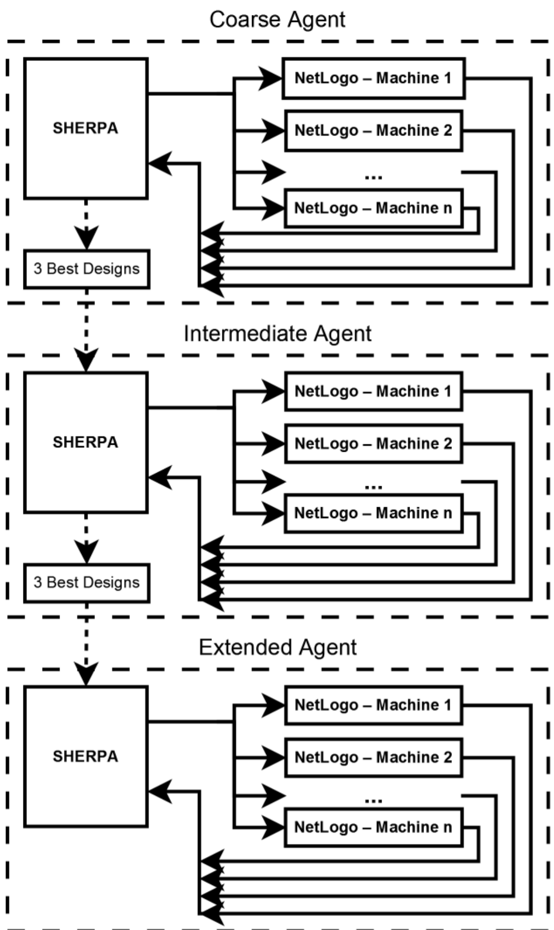

Figure 4: Process flowchart of HEEDS using multiple agent groups and simulating multiple models in parallel.

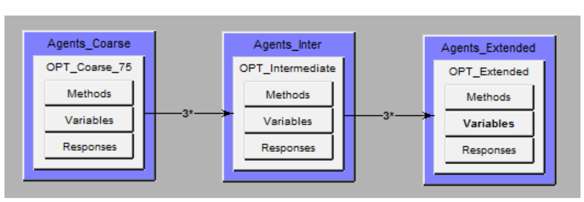

Figure 5: Agent set up from the Assembly page in HEEDS Modeler. This shows the three agents, each in a separate agent group, along with the links they share.

Preliminary Results

The graphs in Figure 6 show the average fitness over time for each agent for a typical run. The numbers are not directly comparable between agents because they are run under different conditions.

The final plot (Figure 6) of the fitness of the extended agent shows both the average fitness when running this three-agent setup, and the average fitness when running that agent alone. When running the extended agent along with the two others, it showed a slight improvement in search. The best design for the agent running with the other two assisting was about 2020 and occurred at design 1421. The best design for the agent running by itself was also about 2020 but did not occur until design 1645. More study is currently underway to assess the effectiveness of using three HEEDS agents rather than one for this problem.

Next Steps and Summary

This problem definition is still being refined as part of an ongoing study. The setup of the agent groups is still being explored. An attempt will be made to develop a hierarchy of fitness functions that will allow rapid exploration of signal control strategies at the coarsest level, and to feed those strategies into more refined agents to speed the overall search. The ABM run to evaluate one evacuation strategy can take anywhere from ten to forty minutes, so even though many simulations are done in parallel on the cluster, an optimization run can still take a long time.

Future versions of the strategy optimization will need to look at more factors affecting traffic flow and analyze the distribution of states of knowledge and goals of the drivers. There are important behaviors that have been left out of the model in order to obtain initial results. For instance, many drivers might not obey the law and sit at a red light for two minutes if they can see some immediate danger. It is important to explore how people behave under emergency situations and deviate from their normal patterns.

It could be beneficial to choose traffic timing strategies that are reasonably fast, while also avoiding driver panic by optimizing for values such as average wait time. It may also be important to look at specific ways that officials can notify drivers of an emergency and direct traffic to different roads.

4572 South Hagadorn Rd., Suite 3A ▲ East Lansing, MI 48823 ▲ TEL 517.664.1137 ▲ FAX 517.664.1175 ▲ www.redcedartech.com

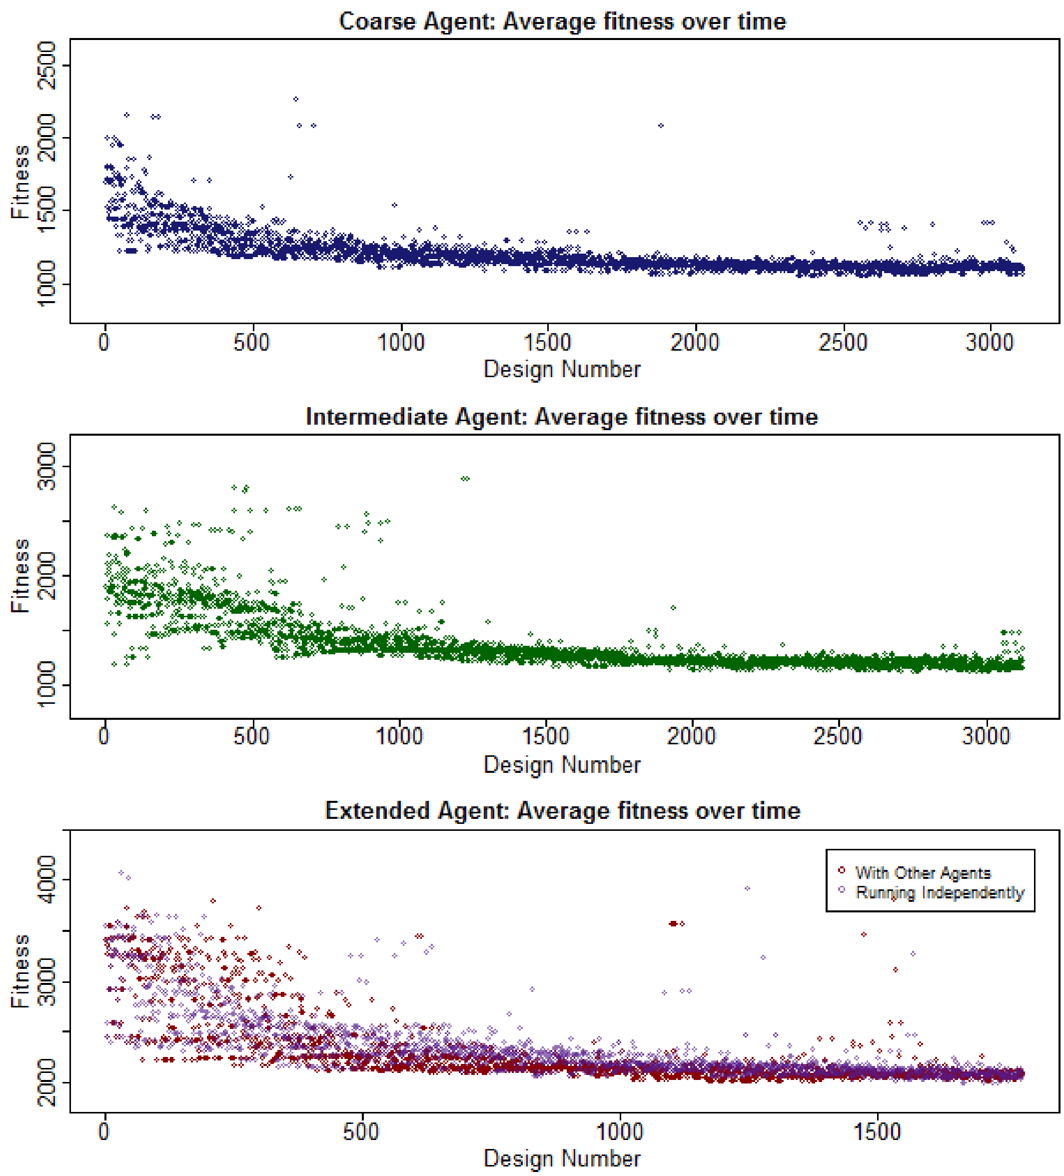

Figure 6. Plot of fitnesses of all solutions evaluated in the coarse agent (top), intermediate agent (middle) and extended agent (bottom, red). Conditions of simulation and fitness calculations are different in each agent, so numerical values are not directly comparable. Bottom plot also superimposes another agent's results (blue) running without benefit of input from top and middle agents; predominance of red along bottom of plot indicates more rapid progress with input from the other two agents, but more detailed analysis and additional runs are needed before any strong assertions can be made.

The process of optimizing the strategy using the ABM is an ongoing effort. But even the current stage of the development illustrates that HEEDS is a powerful tool for strategy optimization using agent-based models.

References

1. Wilensky, U. (1999). NetLogo. http://ccl.northwestern.edu/netlogo. Center for Connected Learning and Computer-Based Modeling. Northwestern University, Evanston, IL.

2. Wilensky, U. (2003). NetLogo Traffic Grid model. http://ccl.northwestern.edu/netlogo/models/TrafficGrid. Center for Connected Learning and Computer-Based Modeling, Northwestern University, Evanston, IL.

4572 South Hagadorn Rd., Suite 3A ▲ East Lansing, MI 48823 ▲ TEL 517.664.1137 ▲ FAX 517.664.1175 ▲ www.redcedartech.com