

Steering Traffic Networks[†]

Yoshitaka Kuwata[‡], David M. Hart & Paul R. Cohen

Computer Science Technical Report 94-43

Experimental Knowledge Systems Laboratory
Computer Science Dept. Box 34610
Lederle Graduate Research Center
University of Massachusetts
Amherst, MA 01003-4610

Abstract

When a traffic management system involves many thousands of vehicles using hundreds of streets and highways, it can be difficult or impossible to tell whether the network is flowing smoothly and to predict how modifications to dynamic control parameters will affect the system. For large-scale traffic management problems it is both necessary and difficult to assess the current state of the system and to predict the effects of modifications. In the absence of informed intervention, a system can evolve into a pathological state or process, in which vehicle progress slows or stops completely. We describe an interactive control system whose purpose is to help human controllers steer the evolving state of the network away from a pathology – avoiding the pathology if it has not yet materialized, and disabling it quickly if it has.

[†] This research was supported by NTT Data Communications Systems Corporation and by ARPA/Rome Lab contract #F30602-91-C-0076.

[‡] Yoshitaka Kuwata is a Visiting Researcher in EKSL from NTT Data Communications Systems Corporation.

1 Introduction

"Traffic congestion, and its associated social and economical detriments, is one of the most significant problems facing our nation today. Rapid population growth since the 1950's, general reliance on the automobile as the primary source of transportation, and the declining incidence of construction of new transportation facilities due to their high costs and public concern for energy and environmental issues, are some of the main causes for the alarming congestion levels. Recent studies indicate that American drivers waste about two billion hours a year in traffic jams causing an average annual loss exceeding \$73 billion to the U.S. economy. If this situation does not improve, the number of hours spent in traffic will increase by over 400 percent by the year 2005. Advance Traffic Management Systems (ATMS) is the backbone of Intelligent Vehicle-Highway Systems (IVHS) and will provide the tools and technology necessary to properly address this important problem."

Alberto J. Santiago,
ATMS Program Manager
High Priority National Program Area

When a traffic management system involves many thousands of vehicles using hundreds of streets and highways, it can be difficult or impossible to tell whether the network is flowing smoothly and to predict how modifications to dynamic control parameters will affect the system. Neither task is particularly challenging for small-scale systems: The entire state of the system can be stored, and control modifications can be evaluated by search. However, for the large-scale traffic management problems faced by an ATMS, it is both necessary and difficult to assess the current state of the system and to predict the effects of modifications. In the absence of informed intervention, a system can evolve into a pathological state or process, in which vehicle progress slows or stops completely.

Pathological states are well recognized in the sciences of dynamic processes: We speak formally of "deadlock" and "thrashing" in computer systems, and we speak more colloquially of vicious circles, diminishing returns, bottlenecks, and gridlock. Each of these terms describes processes running amok or slowing down in a characteristic manner. The prediction and detection of pathological states in traffic systems, caused either by specific traffic incidents or by local deteriorations in traffic network flow, can be used as the basis of a visual, interactive network control system. Models of how pathologies arise in traffic networks can be used to develop visualizations – graphical, model-based displays of pathologies. The purpose of such a control system is to help human operators *steer the evolving state of the network away from a pathology* – avoiding the pathology if it has not yet materialized, and disabling it quickly if it has. Detecting a pathology involves successfully matching a model of a pathology to a history of some of the events in a traffic network; a graphic visualization of this history is produced as a side-effect of detection. The model of the pathology is consulted by the system to suggest alternative steering actions to the human controller.

The visualization of the pathology provides an interface between the user and the model of the pathology, allowing the user to explore (through simulation) in a "what if" mode how steering actions – as suggested by the system and of his own design – are predicted to affect the course of the pathology. Such a system can be used both offline as a training tool, allowing users to see problems and ask "what could I have done to avoid this problem," and online to ask "what can I do now to avoid or disable the problem." The intent of the offline systems is to help traffic controllers identify and prevent pathologies as traffic management strategies are being developed and tested in simulations, whereas the online version is intended to identify and correct pathologies in traffic management strategies as they are being executed.

2 Background

A number of researchers have proposed the application of existing AI technologies for controlling traffic networks (among them [3,9,10]). It is not clear, however, that these technologies will scale

up to the demands of large networks. We are developing new AI techniques to detect pathologies as early as possible and to leverage the perceptual and problem-solving skills of human controllers in an interactive search for appropriate steering action(s). This approach has evolved from our research in plan execution monitoring and adaptive plan modification in several complex, dynamic domains: forest fire-fighting and transportation planning.

The Phoenix system is a real-time simulation of forest fire-fighting in Yellowstone National Park [2]. Numerous autonomous agents execute fire-fighting plans under the direction of a coordinating agent called the Fireboss. Complexity and uncertainty in the simulation environment cause circumstances under which a plan is executed to change often (wind characteristics change, new fires might start, etc.). Under these conditions it is important to represent and reason about plan progress [4], since lack of progress can eventually lead to plan failure. The Phoenix planner is designed to use a combination of reactivity, plan repair and replanning to steer the executing plan around obstacles to plan success [5,6].

We applied the idea of plan steering in a transportation planning problem in which the goal is to move thousands of pieces of cargo through a large shipping network with good resource utilization and close adherence to a target schedule. In this domain we are developing an architecture for steering activity in a transportation networks in order to keep trim to the schedule. This architecture is based on models of the pathologies found in these networks, and applies refinements of techniques developed in Phoenix [7]. We employ a network of demons which monitor activity in the transportation network, watching for the development of pathological situations and warning the controller when they arise.

3 Steering Traffic Networks

We are currently applying our plan steering architecture to the problem of managing flow in large traffic networks. As with the previous domains, we are using a simulator (this one called Sapporo) to identify and model pathologies and also (eventually) to implement an interactive steering system. The domain of traffic network management differs from both fire-fighting and transportation in that there is no "plan" or "schedule" to execute. There is most often a small set of goals, such as maximize network throughput while maintaining safe driving conditions. Since most traffic networks have stable topologies, this can become a matter of tuning a default network signal plan to adjust to daily fluctuations in traffic flow. However, there often arise what can be characterized as pathological situations, such as accidents, roadwork on high volume arteries, or unexpected bottlenecks. When these situations arise, network traffic is steered by changing the topology and signal plans for the affected part of the network. Such changes can be modeled in terms of new expected flow patterns and monitored (much as Phoenix and transportation plans are monitored) to assess how effectively they address the pathologies.

3.1 Sapporo

Sapporo is an integrated traffic control system developed at the Forschungs-zentrum Informatik in Karlsruhe, Germany.¹ One goal of the system is to control traffic lights in a large urban road network in order to maximize traffic flow and minimize delay. As the actual traffic system is dynamic and has many uncertain factors, it is impractical to apply global optimization techniques to the entire network.

Sapporo's approach to traffic management has three important features:

1. *Qualitative Simulation Model.* To simplify the problem, traffic flow is handled qualitatively. Traffic flow characteristics are represented as density zones (e.g., free flow, partial free flow, max flow

¹ Sapporo was developed by Bernd Wild and colleagues in the Department of Technical Expert Systems and Robotics at FZI in a joint research project with (and sponsored by) NTT Data Communications Systems Corporation.

dense traffic, traffic jam, and so on). The movement of individual cars is not represented; rather the movement of a road section's density zone is calculated through simulation. The calculations are based on continuum theory; for each intersection the flow of outgoing roads is calculated from both the flow of incoming roads and fixed probability distributions of heading. (Note that a slightly different approach is found in Sugimoto, et al. [8] which relies on qualitative process theory for the basis of simulation model.)

2. *Distributed control and local negotiation.* Each intersection's traffic signals are controlled by an independent agent. Agents either choose an existing (pre-defined) signal plan, modify an existing signal plan, or build a new signal plan from scratch. Actor theory [1] is used as the local negotiation mechanism, in which each agent proposes a local zone plan based on its own perspective, then through a negotiation process each agent's plan is modified to reflect the concerns of its neighbors' plans.

3. *Simulation.* An event-driven, object-oriented simulation is used to check the behavior of the model and of agents.

3.2 Steering around ATMS pathologies

Another goal of an ATMS is to manage traffic lights according to evolving circumstances so that traffic flow becomes as smooth as possible. For example, traffic accidents can happen anywhere within a road network. If a road cannot be used because of an accident, the control system must manage to use other routes while trying to avoid bottlenecks. To achieve this goal, we can apply the steering architecture in the traffic control system. In this context steering consists of 1) predicting or detecting bottlenecks (e.g., low traffic flow or a traffic jam), 2) visualizing these bottlenecks for traffic controllers, 3) exploring alternative traffic-flow plans by using interactive simulation, and 4) changing intersections' signal plans to reflect the new traffic-flow plan.

3.3 An Example

An example scenario of plan steering in an ATMS will be shown in this section. This example assumes the existence of pre-defined signal control plans (by "plan" we mean coordinated sequencing of intersection signal plans in subparts of the network to achieve specific a goal). Two kinds are pre-defined. One is for normal driving conditions in the global traffic network. This kind of plan can be optimized using offline optimization techniques and knowledge about the normal operating conditions. The other kind is a pre-defined parameterizable plan designed to respond to specific recurring situations, such as accident management, detours, or peak




Figure 1. Plan Steering Example. There is a stadium in the center of this map. The lightly shaded area around the stadium is controlled by a stadium demon, which has a default plan for the departing traffic shown by the black arrows. The medium shaded area around section 3 is controlled by an accident demon, which is assigned temporarily to divert traffic from section 3. Gray arrows show the initial plan of the accident demon.

usage. Some of these can be parameterized in advance, such as those for peak usage; others, such as those for redirecting traffic around accidents, involve dynamically reconfiguring flows in the network.

In this scenario, a pre-defined plan for controlling traffic departing from a stadium is used whenever an event held there ends. This plan directs each intersection to assign long signal durations to lanes directed away from the stadium and toward major highways. A *demon* is assigned to monitor the stadium plan, watching to see that departing traffic flows as expected. Figure 1 shows the traffic network around the stadium as an athletic event is ending. Unfortunately, a big accident has occurred just prior to this time in the center of section 3, forcing it to be completely closed and detouring traffic to other roads around the stadium. As usual, a pre-defined plan for accident management has been instantiated, and an "accident demon" is assigned to the area.

Because there are two goals in this situation which will obviously interact, and independent demons monitoring each goal, some form of coordination is required to mediate the interaction. As it is not practical to model all such cases, and it is the responsibility of human controllers to decide what action(s) to take in such a complex case, the demons report the failure of each of the goals to controllers, and assist the controllers in exploring alternative routings and flow control regimes to alleviate congestion in the overlapping areas. The four steps involved in this process are summarized below:

- 1) *Prediction and Detection of bottleneck.* Using traffic models, demons predict the development of a bottleneck near the stadium. Because the goals of the stadium demon and accident demon are different, they detect pathologies in different places. The stadium demon might detect a pathological waiting queue around section 3. The accident demon detects traffic jams around sections 2 and 4.
- 2) *Visualization.* Each demon is trying to visualize the problem with respect to its goal. For example, the stadium demon might show the predicted density profile of section 4. On the other hand, the accident demon is likely to display the same graph of sections 2 and 1 to the controller and point out the traffic jam in these sections.
- 3) *Explore Alternative Plans.* As the result of negotiation between the two demons, several plans are proposed to the controller. Using simulation, the controller can explore the alternatives and compare the results.

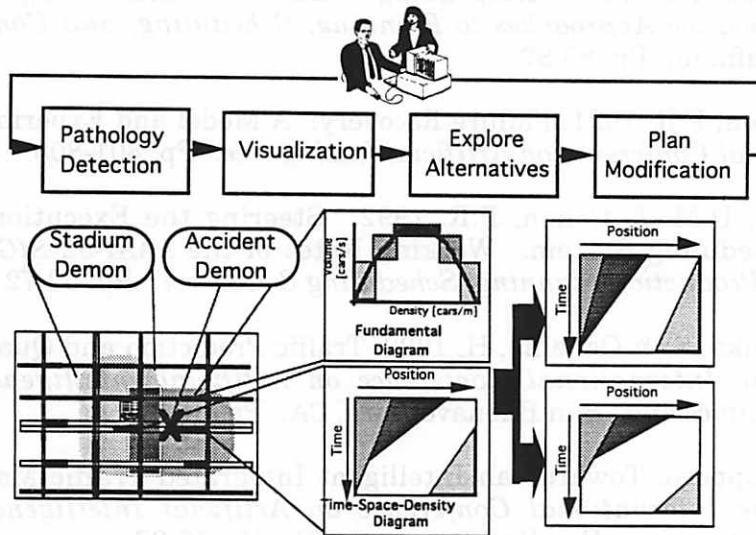


Figure 2. Plan Steering in an ATMS. The *Ideal Fundamental Diagram* represents the relationship between the density of a section and the traffic volume of the road. This diagram is derived from continuum theory. This figure also shows density value (qualitatively) as free flow, partial free flow, max flow dense traffic, and traffic jam. The *Time-Space-Density Diagram* represents the expected status of the road over time. Each gray level represents a different density zone. Darker gray means denser. In this figure, the densest zone, which is traffic jam, shrinks as it moves from right to left and eventually disappears. The *border zone speed* is calculated from the density and volume of the neighboring two zones.

4) *Plan Modification*. After several trials of the simulation of that area the controller can make a decision. After the chosen steering action is taken, demons continue monitoring the progress of the modified plan.

Acknowledgments

This research was supported by NTT Data Communications Systems Corporation and by ARPA-AFOSR contract F30602-91-C-0076. The United States Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright notation hereon.

References

- (1) Agha, G. 1986. ACTORS: A Model of Concurrent Computation in Distributed Systems. MIT Press.
- (2) Cohen, P.R., Greenberg, M.L., Hart, D.M. & Howe, A.E. 1989. Trial by Fire: Understanding the Design Requirements for Agents in Complex Environments. *AI Magazine*, 10(3): 32-48.
- (3) Cuenca, J., Ambrosino, G. & Boero, M. 1992. A General Knowledge-based Architecture for Traffic Control: The KITS Approach. *Proceedings of the International Conference on Artificial Intelligence Applications in Transportation Engineering*. San Buenaventura, CA. Pp. 153-171.
- (4) Hart, D.M., Anderson, S.D., & Cohen, P.R. 1990. Envelopes as a Vehicle for Improving the Efficiency of Plan Execution. *Proceedings of the Workshop on Innovative Approaches to Planning, Scheduling, and Control*, Katia Sycara (Ed.). Morgan-Kaufman. Pp. 71-76.

- (5) Howe, A.E. & Cohen, P.R. 1990. Responding to Environmental Change. *Proceedings of the Workshop on Innovative Approaches to Planning, Scheduling, and Control*, Katia Sycara (Ed.). Morgan-Kaufman. Pp. 85-92.
- (6) Howe, A.E. & Cohen, P.R. 1991. Failure Recovery: A Model and Experiments. *Proceedings of the Ninth National Conference on Artificial Intelligence*. Pp. 801-808.
- (7) Kuwata, Y., Hart, D.M. & Cohen, P.R. 1992. Steering the Execution of a Large-Scale Planning and Scheduling System. Working Notes of the AAAI-92 SIGMAN Workshop on Knowledge-Based Production Planning, Scheduling & Control. Pp. 62-72.
- (8) Sugimoto, T., Nonaka, S. & Oohama, H. 1992. Traffic Prediction and Qualitative Reasoning. *Proceedings of the International Conference on Artificial Intelligence Applications in Transportation Engineering*. San Buenaventura, CA. Pp. 3-18.
- (9) Wild, B. 1992. Sapporo: Towards an Intelligent Integrated Traffic Management System. *Proceedings of the International Conference on Artificial Intelligence Applications in Transportation Engineering*. San Buenaventura, CA. Pp. 19-37.
- (10) Zhang, H. & Ritchie, S.G. 1992. A Real-time Decision-Support System for Freeway Management and Control. *Proceedings of the International Conference on Artificial Intelligence Applications in Transportation Engineering*. San Buenaventura, CA. Pp. 39-58.