LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

UCRL-TR-202572

NA22 Model Cities Project – LL244T An Intelligent Transportation System-Based Radiation Alert and Detection System

Steven G. Peglow

3/16/2004

Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California, and shall not be used for advertising or product endorsement purposes.

Acknowledgements

This work was sponsored by the US Department of Energy, NA-22 under the technical direction of the PPAC Program at Lawrence Livermore National Laboratory. Significant contributors included Stein Weissenberger, LLNL, Dr. John Lathrop, consultant, and Yonnel Gardes, UC Berkeley. We would like to thank Ed Rowe of Iteris Corporation and Matt Edelman of TRANSCOM for their help and encouragement.

This work was performed under the auspices of the U.S. Department of Energy by the University of California, Lawrence Livermore National Laboratory under Contract No. W-7405-Eng-48.

Table of Contents

Overview **4**

1.0 Concept of Operations **5**

2.0 Logical Architecture **15**

3.0 System Design Issues **24**

4.0 Simulation of Target Tracking/Interdiction **43**

Summary of Accomplishments **80**

Recommendations for Future Work **81**

Overview

The purpose of this project was twofold: first, provide an understanding of the technical foundation and planning required for deployment of Intelligent Transportation System (ITS) -based system architectures for the protection of New York City from a terrorist attack using a vehicle-deployed nuclear device; second, work with stakeholders to develop mutual understanding of the technologiesand tactics required for threat detection/identification and establish guidelines for designing operational systems and procedures. During the course of this project we interviewed and coordinated analysis with people from the New Jersey State Attorney General's office, the New Jersey State Police, the Port Authority of New York/New Jersey, the Counterterrorism Division of the New York City Police Department, the New Jersey Transit Authority, the State of New Jersey Department of Transportation, TRANSCOM and a number of contractors involved with state and federal intelligent transportation development and implementation.

The basic system architecture is shown in the figure below. In an actual system deployment, radiation sensors would be co-located with existing ITS elements and the data will be sent to the Traffic Operations Center. A key element of successful system operation is the integration of vehicle data, such as license plate, EZ pass ID, vehicle type/color and radiation signature. A threat data base can also be implemented and utilized in cases where there is a suspect vehicle identified from other intelligence sources or a mobile detector system. Another key aspect of an operational architecture is the procedures used to verify the threat and plan interdiction. This was a major focus of our work and discussed later in detail. In support of the operational analysis, we developed a detailed traffic simulation model that is described extensively in the body of the report.

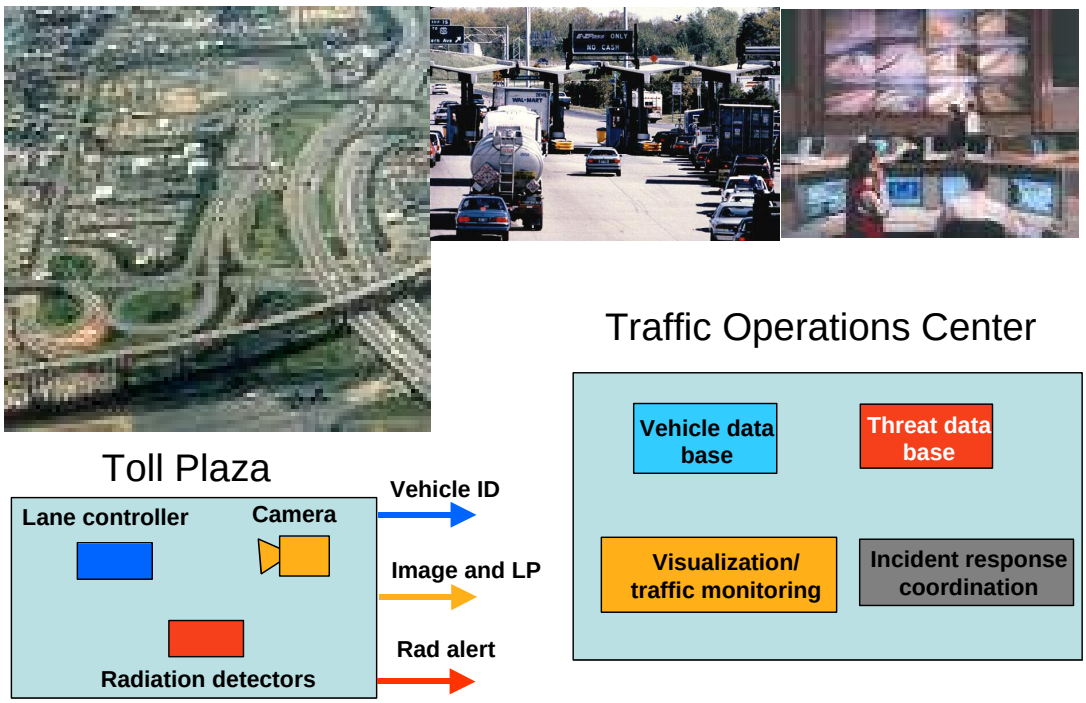

1.0 Concept of Operations

This section describes a Concept of Operations for the Radiation Alert and Detection System (RADS). The ConOps is based in part on information gained from transportation and law-enforcement agencies in New Jersey. RADS and its associated ConOps were developed to explore the feasibility of basing a highway nuclear detection and response system on existing Intelligent Transportation System infrastructure and law enforcement organizations. In this exemplary implementation, RADS is deployed on the New Jersey Turnpike. The overall conclusion of this exercise is that such a deployment is indeed feasible, but that there are certain operational questions that require experiment and testing before the scope and details of a practical system can be defined. Recommendations for future work are made in a later section of this report.

RADS performs the following functions, through a combination of automated processing and human decision making:

• Collects single or multiple radiation readings at stations on the Turnpike and associates these readings with other data such as visual images of the source vehicle
• Compares radiation readings to thresholds and assesses the validity and significance of each data set
• Decides whether to request the interdiction of a suspect vehicle, based on the magnitude and number of radiation readings, vehicle images and other identifiers, associated Intel, specific alerts, and other information bearing on the likely nature of the source of radiation
• Requests a vehicle interdiction following an established procedure
• Provides timely, accurate, and useful information to field officers performing the interdiction
• Supports an interdiction through appropriate SOPs, detectors, and training to the field officers who will perform interdictions and take associated radiation measurements.

<u>ConOps Overview</u>

RADS relies extensively on existing roadway sensors, communication infrastructure, and public operating agencies and response forces. RADS is supported by automated reporting and assessments, but human decision making is also required at key points. Figure 1 shows its high-level communication architecture. In the following, we sketch how this RADS implementation performs detections, makes assessments, and requests and executes vehicle stops.

<u>Detection</u>

RADS detector stations use the control and communication infrastructure of existing NJ Turnpike Authority toll entries and exits, with the addition of radiation detectors mounted on toll station gantries and tied into existing power and communications. The NJ TA's existing control and communication system for registering toll collection violations can

be adapted to RADS purposes by integrating radiation alarm signals with existing imaging, identification, recording, alerting, and messaging functions.

Registering and reporting radiation alarms is very close in function to the routine Turnpike Authority activity of identifying and reporting toll collection violations. Consequently, radiation detection can be added to the Turnpike Authority's responsibilities with relative ease. This view is confirmed by lengthy discussions with the Turnpike Authority and by their expressions of ability and interest in participating in tests of such a system extension.

Detection functions are performed entirely automatically, with alarms and associated data organized and clearly communicated to key human operators. (See "Responsibilities" below.)

Assessment

Given a single or multiple set of over-threshold radiation signals, an assessment must be made on whether to interdict the associated vehicle. To support this decision, RADS organizes and processes a variety information in a structured fashion, including: the number and magnitude of over-threshold radiation readings, vehicle images and other identifiers associated with those radiation readings, Intel and specific alerts that bear on the likelihood of transport of a weapon with a radiation source, and information bearing on the possible presence of a "nuisance" source of radiation (e.g. vehicle type, any "source authorization", radiation spectrum information). The resulting automated recommendation is communicated to human operators (see "Responsibilities below) who make decisions on the basis of this and other information.

Vehicle descriptors and interdiction aids

To facilitate vehicle stops, RADS supplies a variety of information to field officers, including: forecasts of vehicle arrival times at upstream locations, based on current traffic and estimated vehicle speed; color images of the vehicle; front and rear license plate images; and classification of the vehicle into one of eight categories.

Response

Troop D of the New Jersey State Police has sole responsibility for law enforcement on the Turnpike. Among their normal functions are vehicle stops, of which they make in the range of 50-200/day.

The NJ Turnpike Authority is already integrated with the unit of the NJ State Police (Troop D) that has sole responsibility for law enforcement on the Turnpike. Hence, coupling the two agencies together in the ConOps is a natural addition to their current set of normal activities.

Specific responsibilities

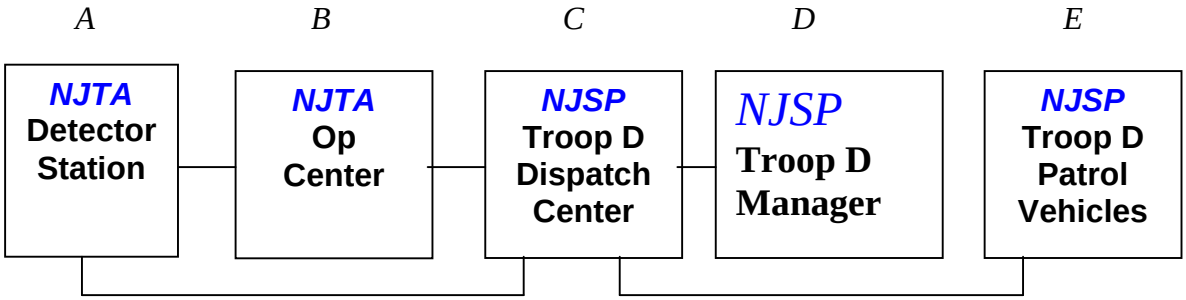

The diagram above illustrates the normal chain of communications, from left to right, as positive radiation detections lead to a vehicle interdiction:

A. One or more detector stations, located at NJTA sites, produce radiation alarms and corollary sensor data about the vehicle that generated the alarms. These signals are transmitted simultaneously to the NJTA Operations Center and the New Jersey State Police Dispatch Center for Troop D.

B. The NJTA Operations Center has the primary responsibility for the following primary functions: communicating verified radiation alerts to the New Jersey State Police and providing the NJSP with data that will assist them in any eventual vehicle interdiction (e.g. vehicle and license plate images, vehicle ETAs). They are also responsible for: managing, operating, and maintaining the RADS sensor system, ensuring the validity and timeliness of RADS reports, interfacing with other transportation agencies (e.g. TRANSCOM and transportation units of NY/NJ Port Authority), advising the NJSP regarding favorable times and places for interdictions (from the viewpoint of traffic and highway conditions), and, in coordination with the NJSP, executing highway control actions such as posting communications on variable message signs, changing lane speed notifications, and effecting lane closures.

C. The NJ State Police Troop D Dispatch Center receives radiation alarm reports from the NJTA. The Assistant Duty Officer (ADO) (present 24/7 at the Dispatch Center) will have the primary responsibility for assessing the validity and significance of RADS alarm reports and deciding on actions to take. These actions could include, e.g., commanding an interdiction, obtaining more information on the vehicle, generating alerts to patrol or other law enforcement units, or determining that the vehicle is a “nuisance” source of radiation and taking no action. The ADO will be assisted in this decision by clearly organized and processed alarm information presented in an effective graphical user interface (GUI), and by an SOP written specifically for this function. The information presented to the ADO will include the radiation history of the vehicle, its time record within the turnpike system, and any vehicle/operator information that has been obtained. As spelled out in the SOP, the ADO will have the option of making decisions on his or her own, or first consulting with the Duty Officer. Both officers will take into account any Intel available, as well as the overall level of security alert.

When the ADO requests a vehicle stop, the ADO communicates vehicle information and ETAs to trooper(s) in designated patrol vehicle(s).

D. When the ADO requests advice, the NJSP Troop D Duty Officer will respond, again according to procedures spelled out in the SOP. The Duty Officer will have available the same information via the same GUI that the ADO has.

E. When the ADO requests a vehicle stop, the designated trooper on patrol performs the interdiction according to an SOP written for this purpose. The SOP spells out procedures for making the stop, for making field radiation measurements, and for alerting other organizations as required. Most desirably, patrol vehicles will be provided with a system that permits the receipt and display of vehicle images. The trooper will be provided with appropriate radiation detectors and trained in their use.

1.1 RADS exemplary design:

A simplified RADS deployment was constructed to illustrate how a part of a regional RADS system might be configured. The design consists of four stations on the New Jersey Turnpike, and could be the basis for a simulation or tabletop exercise involving a vehicle with a radiation source traveling north on the turnpike to NYC via the Holland Tunnel.

The four are located as follows (working "upstream" (west and south) from NYC):

1. At the tollgate entrance to the Holland Tunnel at the Toll Plaza.

2. At the tollgate exit onto I-78 from the NJT at Interchange 14, 5.9 miles upstream from Station 1.

3. On the NJT north at the overpass at Interchange 10 (at the junction with I-287), 22.5 miles from Station 1. (Note: there may be a better location just upstream from the overpass, that would permit capture of newly entering vehicles at this intersection, that would not otherwise be captured by sensors at the overpass itself.)

4. At the NJT tollgate entrance at interchange 7A (at the junction with I-195), 50.6 miles from Station 1.

Notes:

A. Stations 2, 3, and 4 belong to the New Jersey Turnpike Authority: Station 4 is a NJTA toll *entrance*, 2, is a NJTA toll *exit*, while Station 3 is neither. Station 1 is a toll entrance on a NY/NJ Port Authority facility. Thus the set of stations spans a range of realistic location types and operating agencies.
B. Each of these stations as currently constituted (or as projected for near-term upgrade) possesses different combinations of sensors and infrastructure. Consequently, the

RADS deployment at each station may differ, reflecting the different opportunities and constraints at each location. E.g. at the two extremes, Station 2, as an NJTA exit gate, currently possesses the maximum pre-existing complement of sensors and infrastructure, while Station 3 has the minimum (none), as it consists of a presently un-instrumented overpass. Station 4 currently has most of what Station 2 has, minus only camera coverage (and light gate/treadle presumably), as it is an entrance gate. (But note that the computer control and communication is identical at entrance and exit gates, so it should be straightforward to add cameras and integrate them in an identical fashion at entrance gates as well.) Finally, Station 1 is an installation similar to 4, except that it is operated by the NY/NJ Port Authority, whose sensor and communication architecture was not investigated.

C. There will be E-ZPass tag readers at each of the stations (1,2 and 4 existing, 3 assumed as part of a special build). In addition, we assume a TRANSMIT deployment of E-Z Pass Readers at approximately 1.5 mi. intervals on the NJTA, between Interchange 7 and 8A. There will be no readers between 8A and 14 (except at Interchange 10 as part of the RADS installation there), but there will be readers installed from Interchange 14 to the Holland Tunnel approaches, including some on the Pulaski Skyway and two near the Holland Tunnel Toll Plaza. We assume that RADS has been permitted access to every E-ZPass tag number associated with an over-threshold radiation signal, and to every tag called out by specific Intel from authorized law-enforcement agencies. In the case of existing tag readers, this access will be obtained at the roadside at the NJTA server. "Access" is taken to mean the ability to archive and communicate the tag number to authorized agencies, to obtain information about the identity of the owner of the tag, and to track the tag automatically on the roadway through successive tag readings.

A RADS station consists of the following elements:

1. Radiation detectors. A design will specify type and output (counts relative to threshold plus spectrum, etc.) and location (normally overhead on a toll booth gantry, or, in the absence of a toll booth, overhead on the overpass structure for Station 3, or possibly on a sign gantry or TRANSMIT reader support structure).
2. E-ZPass reader. These are existing or planned for implementation at each tollgate, except for Station 3, where there are no tollgates and the reader will be mounted under the overpass structure.
3. Video camera. These are presently located at tollbooth exit gates (Station 2). We assume that for this design they will also be installed on appropriate structures at Stations 1 and 4, and integrated with other sensor readings from those locations. It will also be highly desirable to similarly instrument Station 3.
4. Light curtain. Together with treadles, these produce a vehicle profile.
5. Treadles. These, with the light-curtain, give a vehicle axle count. Readings from light curtains and treadles together place vehicles into one of eight classes.
6. Roadside computer and communication equipment. For Stations 2 and 4, this consists of the existing system of Lane Equipment Controllers (LECs), Lane Control Computers (LCs), Plaza Servers, and fiber optic communication lines. Station 3 should be implemented to be as nearly identical as possible with this system, as

should Station 1—however, the degree to which this non-NJTA station (#1) can be so integrated is TBD.

7. Digital Video Transaction Data Multiplexing Installation. This projected system will permit a digital video image of the vehicle to be captured for any defined "alarm" condition, e.g. an over-threshold radiation signal, and indexed time and location and with other vehicle information, e.g. vehicle type, E-Zpass tag, license plate image, and speed. The alarm is immediately communicated to the NJTA operations center, and the information archived.
8. (Optional) Vehicle presence loop detector, together with associated communication, computer, and signal processing algorithms. This gives the possibility of generating a vehicle signature associated with an over-threshold radiation signal, thus assisting in the tracking of a targeted vehicle.
9. (Optional) Automatic weigh-in-motion sensor.

Sensor information from NJTA properties at Stations 2, 3 and 4 will be collected locally at each station, multiplexed with video and fed to the New Jersey Turnpike Authority Operations Center (TOC). For now, assume that similar data from Station 1 can similarly be collected and communicated to the TOC. Consistent with the existing NJTA procedure for flagging and recording violations, any data associated with an over-threshold radiation signal will be treated as an "alarm", and called out. This alarm information will consist of the following:

1. what happened (the occurrence of an above-threshold radiation signal, with the specific associated "radiation signature" together with "recent above-threshold radiation signal histories", for this or any vehicles)
2. when it happened
3. where it happened
4. "vehicle signature", which will include:
    - vehicle image
    - vehicle class
    - vehicle speed
    - E-ZPas tag (if available)
    - license plate image (if available)
    - loop detector signature (if available)
    - weight (if available)

In addition, the system will record and track the following data:

1. Current vehicle tracking status
    - Time vehicle entered system
    - List and times of stations registering vehicle passage, with radiation readings at each
    - ETAs at upstream locations
    - A measure of the likelihood the vehicle contains a radiation source
2. "Source authorization" data if a licensing system is used to authorize the shipment of benign sources of radiation.
3. Any information obtained on the ownership and criminal status of each vehicle.

4. A record of all data accumulated on this vehicle while on the RADS system.
5. Actions requested, pending, and completed regarding interdiction or further measurements on each vehicle.
6. A complete historical database on all radiation (and associated other) sensor readings at all stations, and all actions taken as a result. This database will be used to periodically review performance and upgrade the system.

1.2 Test Scenarios

These scenarios are appropriate for testing the exemplary design through simulation or tabletop discovery exercises.

Design threats and ground truth

The following basic threats apply to each of the scenarios, with certain variations.

The design threat consists of two variants:

a.) an improvised device
b.) an RDD

Each device is of a size that permits it to be transported in one of several vehicle types: a medium size sedan and van. The vehicles have unobscured license plates. For some threat scenarios, the vehicle possesses a valid E-ZPass tag, for others it does not. There may be a passenger in the cab in addition to the driver. The devices are not armed under transport, and must be manually armed when desired. The type and strength of each source, and the associated shielding, are assumed to be such as to produce the stipulated readings by the radiation sensors in the various scenarios.

The vehicle enters the New Jersey Turnpike (NJT, US 95) at Interchange 7A, from I195, traveling north, at 1300 on Wednesday July 16, approximately 50 miles from New York City, the vehicle's destination.

The vehicle's nominal planned route is as follows: continuing north up the NJ Turnpike, it turns off at Interchange 14 onto the Pulaski Skyway, then into the Holland Tunnel. The vehicle plans to exit the tunnel continue on to its target at One Police Plaza. For most scenarios the vehicle's average speed on the turnpike is approximately 50 mph.

An alternate route is as above, except that the vehicle exits the turnpike at exit 14C.

The scenarios are arranged in order of increasing challenge to the system. Times for Scenario 1are notional, but are based on the information obtained from interviews with various model agencies (see *Response Options Meeting Notes*).

Scenario 1

1200
• A Condition Yellow national alert level was raised yesterday to Condition Red for the New Jersey-New York City area. NJSP organizations correspondingly raise their internal alert levels and activate corresponding standard procedures.
• Specific Intel was received this morning by all RADS management and operators, from NJSP HQ: It is possible that in the next two days a terrorist driving a silver Buick Century, mid-nineties vintage, will be carrying a radiation dispersal device through New Jersey, with a possible destination of NYC.

1230
RADS thresholds are adjusted based on this high level of alert; RADS operators are requested to be on the alert for associations of the specified vehicle type with radiation signals; and all elements of the response force are put on a heightened state of alert.

1300
A strong radiation reading, far above the threshold, is obtained at station 4, from a manual toll entrance, simultaneously registering an alarm at the NJTA Operations Center (TAOC) and at the NJSP Dispatch Center (SPDC). Simultaneously, a black and white image of the vehicle from the rear is obtained at both locations (from the license-plate camera), as well as a color frame of the vehicle (from the front) taken from the toll station video. A partial license plate number is immediately apparent. The video image shows two occupants in a silver sedan. Forecasts are made, with ranges, of vehicle locations if it stays on the turnpike.

1302
The Assistant Duty Officer (ADO) at the SPDC telephones the operator at the TAOC and verifies that the equipment was functioning properly and that the signals are valid ones. The ADO reviews the vehicle images and finds that the pictured vehicle is a good match to the Intel.

1307
The ADO pushes all data to the Duty Officer and requests approval for an interdiction of the subject vehicle.

1312
The Duty Officer at Troop D assesses the request and all associated data, approves the interdiction, and communicates this back to the ADO.

1313
The ADO, having determined the position of patrol vehicles in the best position to interdict the target vehicle, communicates appropriate information to these vehicles with a request to interdict, control the vehicle and its occupants, and as appropriate make radiation measurements.

1328
A patrol car brings the car to the roadside two miles south of Interchange 10, approximately twenty miles from Station 4.

Comments:
This is the "best plausible" case, with the physical RADS system working optimally on signals uncorrupted with noise and produced by an actual device source with a bright radiation signature. Under a heightened state of alert, strongly above-threshold radiation sensor measurements are collected in association with clear visual images (vehicle, occupants, and license plate). The vehicle passes through manual tollgates, giving optimum radiation sensitivity.
By providing the high alert level and the advance Intel, this case also avoids much of the problem of false positives, which under these emergency conditions are unlikely to be viewed as costly. False positive effects will be explored in variant scenarios, under lower levels of alert and absent specific Intel.

Scenario 1 Variations:

1a. If an interdiction were not commanded after Station 4 readings, and no further reading were obtained until Station 2 (Exit 14), how might the interdiction process differ? (Assume that an *alert* had been issued after the first reading at Station 4.) Is there enough time for an interdiction before the Holland Tunnel has been reached? Are there ways to use the traffic control system to assist in this interdiction? Stop or slow traffic at the Station 1 tollbooths? How would these answers differ if the interdiction were commanded after Station 3? In these cases, what would be the value of readings taken at Station 1 itself? How could they be best used?
1b. If vehicle, occupant and license plate images were corrupted (due, e.g. to bad weather) so that only a very imperfect vehicle description was obtained, e.g. a light-colored four-door sedan, how would the process be effected?
1c. If the vehicle possessed an E-ZPass tag and used automatic tollgates, how might the identification, decision and interdiction processes differ? Assume that all signals/images are equally good in spite of the greater speed of the vehicle through the tollgates. Also assume that there are some TRANSMIT tag reading stations along the turnpike between station 3 and 2.
1d. If first detection occurred at Station 3, how might the scenario evolve? Assume that interdiction is commanded after readings at Station 2.
1e. Same as 1d, but with first detection at Station 2, and interdiction commanded either then or after detection at Station 1.
1f. If first detection occurs at Station 1, what system arrangements are required to obtain interdiction in the adjacent toll plaza?
1g. How would this scenario evolve if no prior condition of alert existed, and no Intel were obtained on driver? E.g. if RADS were initially in a low state of alert, and if there were possibly a less extensive response force available?
1h. The vehicle eludes interdiction at the toll plaza, traffic is stopped at the Holland Tunnel exit, and interdiction occurs inside the tunnel.

1i. Following a complete series of detections as in Scenario 1, the vehicle escapes into NYC, and is searched for there.
1j. The vehicle "disappears" from the system following the detection at Station 3, and is not heard from again.

Scenario 2

This is similar to Scenario 1, still involving an actual weapon source (now with more shielding material), but with weaker and more ambiguous radiation signals (now only slightly over threshold at each station). No prior level of alert exists. Non-radiation signals remain clear. The declared status of the vehicle reaches only to "presumptive", following station 2.

Scenario 2 variations:

2a. Weakened or absent other information/signals (e.g. degraded vehicle images due to bad weather).
2b. Vehicle eludes interdiction or is otherwise "lost" to the system.

Scenario 3.

Radiation readings are generated from a nuisance source. All signals strong, as in Scenario 1. Vehicle is interdicted at the Station 2 toll plaza. Source found to be a medical one.

Scenario 3 variations:

3a. Vehicle not interdicted at the Station 2 toll plaza, and proceeds to NYC (not to One Police Plaza, however, as vehicle is not carrying a weapon).
3b Vehicle is approved after matching E-ZPass tag with authorization database, and vehicle is not interdicted.

2.0 Logical Architecture

2.1 Overview

Figure 1 (next page) presents an overview of the logical architecture of the Radiological Alert & Detection System (RADS). On the left, the vertical arrow pointing upward marks the notional route of the suspect vehicle, traveling north on I-95 New Jersey Turnpike. As it does so, it passes by four RADS radiation detection stations. Those are numbered 1 through 4, on the left. The New Jersey Turnpike interchange numbers are marked in the ellipses. Each ellipse is labeled "First Look," "Second Look,' etc. That reflects the concept that the system provides several opportunities to detect a radiation signature from a given vehicle. As mentioned on the figure, as the system gets more looks at a vehicle, it gets "smarter," in that with more information, there may be fewer nuisance alarms (though that is subject to discussion later in this section). On the other hand, additional looks involve less warning before the vehicle reaches the Holland Tunnel.

At each detector station ("look"), a process takes place represented by the rest of Figure 1. Each station has a radiation detector or detector suite. One option is to accompany a detector alarm with immediate recording of images of the vehicle, including license plate, and possibly other measures. Those might include a special imaging set triggered by the radiation detector alarm, as well as possibly toll tag, any other readable tag such as a license to carry radioactive materials, loop detector signature, treadle (axle) count, weight, and any other classification, e.g. the New Jersey Turnpike identifies eight vehicle classifications.

As indicated on the figure, all parts of the figure to the left of the "Green Line" (thick vertical line in B&W), are in the jurisdictional purview of the New Jersey Tollway Authority. The parts of the figure to the right of the "Green Line" are in the jurisdictional purview of the New Jersey State Police -- the agency we would charge with the interdiction function in the New Jersey case.

Upon a radiation alarm, the detector and visual imagery (and possibly other information) is fed to an operation called in the figure, "Threat Assessment." That operation combines the information from the RADS Detector Station with other intelligence and any background alert levels, to reach a decision to interdict or not. A key question is who conducts that operation, and where. That raises the issue of how to get an adequately trained person to monitor the information 24/7, or be available for consultation 24/7, which in turn depends on how heavily processed the system output can be to provide a decision-supporting display.

If the decision is to interdict the vehicle, a request to interdict goes out to an active State Police patrol car in the area. That could be any currently cruising car, or one of a set of specially trained officers possibly in a specially-equipped car. That officer executes an interdiction, which could consist of:

- possibly a drive-by scan and visual inspection with a cruiser equipped with radiation detectors and cameras.
- pulling the vehicle over, either as soon as is safe, or following some guidelines, or to one of a few predetermined sites. Then:
- an interrogation and visual scan;
- possibly a scan with a hand-held radiation detector;
- a decision whether to release the vehicle or request a federal nuclear weapon specialist team, in which case the officer would detain the vehicle until the arrival of that team.

Section 3.1 includes a more complete listing of the options as part of a list of "Response Options."

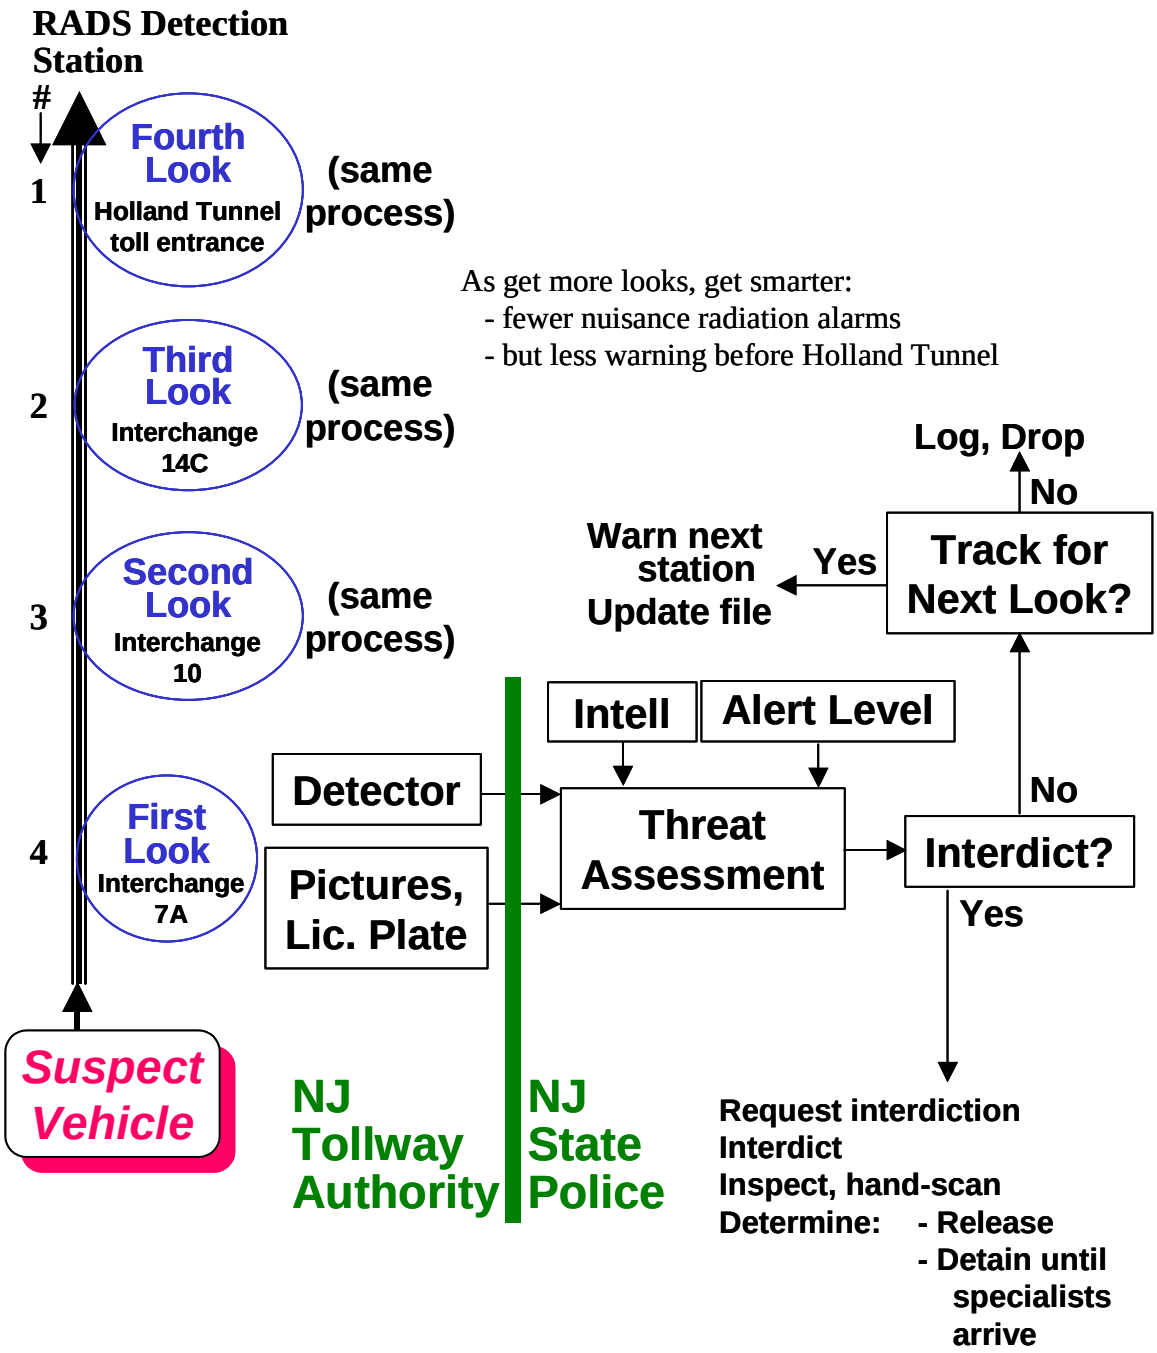

Figure 1. RADS Architecture Overview

If the Threat Assessment Officer decides not to interdict, he or she can decide either to track the suspect vehicle for another look at the next radiation detection station, or simply drop the suspect vehicle, logging it into the records database. If he or she decides to track the vehicle, the system:

- initiates or updates an inference file on the vehicle, which can involve Bayesian statistical updating;

- sets up any vehicle-matching data, which could include all the information listed earlier that may be collected at the detection station and expected time window of arrival, so that the next station can know when it is measuring parameters on the same vehicle that triggered the first detection; and
- warns the next station, if there are procedures that need to be set up, such as a manual visual check;
- sets any special alert level for the State Police cruisers that may be involved, so that they may be more ready to respond to a subsequent request for interdiction;
- sets any special alert level for the nuclear weapon specialist team, again so that they may be more ready to respond to a subsequent request to go to an interdiction site.

The process of data collection, threat assessment and decisions is repeated at each of the subsequent radiation detection stations.

2.2 Key Features of the Logical Architecture

2.2.1 Threat Assessment

Threat assessment is a central function, combining data from the all the detection stations at one site. No inference is local to any one detection station. If we have solid vehicle-matching data linking a vehicle to itself as it passed through one or more previous detection stations, then the system would combine data from the previous hits to build a composite inference about the vehicle, based on the multiple readings. If that vehicle-matching is not solid, then the previous-hit data would be discounted accordingly.

The threat assessment process involves the combination of several types of data:

- radiation detector readings on the suspect vehicle;
- any imagery, including imagery of the license plate and any attempt to lift characters off that image;
- any other data that could be collected, such as any of the following:
    - toll tag;
    - the fact that the vehicle has or lacks a medallion, i.e., a tag establishing that it is licensed to carry radioactive material (and a serial number associated with that medallion);
    - loop detector signature;
    - treadle (axle) count;
    - weight;
    - automatic vehicle classification, e.g., the New Jersey system classifies each vehicle into one of eight categories.
- any country-wide threat level information, including anything from very general levels of alert, to very specific information, such as reason to expect the transport of a nuclear device;
- any threat information specific to the region, route or time of day, which could vary over the same range from general to specific just listed;
- any intelligence that may be able to be gathered from efforts launched by the initial detection, e.g.,

that the license plate matches that of a rental vehicle, which could affect the level of suspicion;
- more static data, such as a list of suspect vehicles.

Considerable effort would be called for to process this long list of information, using decision-aiding algorithms and displays, into as clear a decision-support framework as possible.

Vehicle matching

Two basically different vehicle matching operations are proposed for the RAD System. First is the matching of a vehicle passing through a second radiation detection station with the itself as it previously passed through another detection station on the same trip. That matching can be based on any of the information listed in the previous section, and the time window estimated for the vehicle to arrive at the second detection station. Imagery could include whatever information could be collected on the license plate characters. While license plate characters, toll-tag number, medallion number, loop detector signature, treadle (axle) count, weight, vehicle classification and arrival-time window could be automated, automated imagery matching could be a challenge, which leaves us with the need to consider manual visual matching at subsequent radiation detection stations, which could be prohibitively burdensome with a system that could have a large number of nuisance alarms each day, though it could be instituted during times of high alert.

The second type of matching involves matching the observed vehicle against as many as three static databases:
- suspect vehicles;
- nuisance-alarm vehicles, i.e., vehicles that have created hits before, but then found to be legitimate;
- vehicles that are licensed to carry radioactive materials, but do not have automatically readable tags.

The last two of those databases could be large, suggesting that matching should be automated, which would probably eliminate imagery matching except possibly license plate imagery. That would leave automated matching to be based on the automated data listed in the previous paragraph. The third type of matching listed could be at least partially replaced if a medallion system is set up to mark vehicles licensed to carry radioactive materials. Such a system is described in more detail in Section 5.3.2. If that system involves directly readable tags, such as a toll-tag type of system, then medallioned vehicles could be immediately eliminated as suspects, and would not even show up for the database matching operation suggested here. Note that medallions could specify radiation signature type, so that the system would only eliminate vehicles with medallions consistent with the observed radiation signature.

Interdiction Force: Roles and Interfaces

As indicated in Figure 1, the interdiction force in this application is the New Jersey State Police. Other implementations will have another agency as an interdiction force, for example perhaps the California Highway Patrol in California.

Our appraisal of the situation, and our interviews, led us to the conclusion that the system could not recommend interdiction automatically -- there would always have to be a "human-in-the-loop" to take the "advice" offered by the automated system, combine that with his or her own judgment and other data that might be available, and then make an informed decision as to whether or not to request an interdiction.

So a key feature of the RAD System is its user interface -- its display to the Threat Assessment Officer. The demands of reaction time dictate that the human-in-the-loop decision maker must be able to conduct the threat assessment in just a few minutes. That is, he or she must be on duty or on short-notice call 24/7, and must be able to read the system outputs, combine those with whatever other data is available, and come to an informed decision as to whether to interdict, track, or drop a suspect vehicle. That, combined with operational realities, means that the system must present the data in a way that supports the decision for a person who can be trained up to some level, but who is not an intelligence analyst.

So notionally, we can characterize the system output as if it were a series of colored lights. Of course all the basic system information that could be useful to the Threat Assessment Officer would be made available, but we would also want the system to run through some algorithms to recommend actions at two or more levels. We can refer to those in terms of colored lights:

- Green Light: System operating normally, no suspect vehicles.
- Red Light: System deduces that there is enough evidence to recommend an interdiction. The TAO can countermand that if he or she has other data that suggests that the vehicle is less suspicious than the RAD System could judge based on its internal data, but in the absence of such data, the system lighting "Red" is recommending that the TAO request an interdiction. We could consider that the system could also provide some indication of the likelihood that the vehicle in fact has a weapon on board. We would confer with the interdiction force as to the clearest way to communicate that information.

The above two "lights" comprise the minimal level of advice to the TAO. But the system could go beyond two display levels with:

- Yellow Light - 1: Suspicious readings, but not suspicious enough to recommend interdiction unless the TAO has other data, or may obtain other data, that would increase the likelihood that the vehicle should be interdicted. That is, the system lighting a "Yellow Light" is indicating either of two things:
  - If the TAO has other evidence raising the level of suspicion, then he or she should recommend interdiction; or
  - Later evidence that might be gathered at downstream radiation detection stations could lift the level of suspicion enough to recommend interdiction.

- Yellow Light - 2: There could be different levels of "Yellow Light," indicating different levels of suspicion, though all of them intermediate between recommending interdiction and recommending no action. We would confer with the interdiction force as to the clearest way to communicate those different levels of suspicion.

As should be clear from the previous sections, the New Jersey State Police, or in general the interdiction force, is a key player in the RAD System. We can define three types of roles:

1.) Set up its parts of the system:
- Participate in development of system concept of operations, as explained below in Section 2.3.3;
- Participate in development of the user interface, as discussed in the previous section;
- Deploy the necessary equipment, as determined by the particular ConOps, such as hand-held scanners,
  perhaps cruisers outfitted with detectors for drive-by scans;
- Carry out the necessary training of Threat Assessment Officers, and of the officers who will carry out interdictions;

2.) Carry out threat assessments.

3.) Carry out interdictions.

2.3.3 Concepts of Operation for Interdiction Force

The development of interdiction concept of operations (ConOps) has three major parts:

1.) ConOps for conducting threat assessments. That involves determining who staffs the system display 24/7, if that person is the Threat Assessment Officer, or if that person calls in the TAO. Aside from participating in the development of the user interface, as discussed above, this also involves developing the guidance to TAOs regarding how to carry out the decisions involved, what data to seek out and use other than the RADS data, and how to use that data. With the New Jersey State Police, we would suppose that the Assistant Duty Officer (ADO) would monitor the RADS outputs at all times (direct feed, not through call-takers). Any non-routine decision to interdict would involve consultation with the Duty Officer, who would be on call with less than ten minutes response time. Note that the decisions include:
- the setting of different alert levels, if that is an option. That could include alert levels for the interdiction officer and the nuclear weapon specialist team.
- whether or not to request a drive-by scan, if that is an option.
- whether or not to request an interdiction.
- if not an interdiction, whether to track the vehicle or drop it.

2.) ConOps upon decision to interdict:

How to use the information provided (e.g., imagery, license plate information, the time the vehicle left the last radiation detection station) to intercept and identify the vehicle;

How to pull over the vehicle, including any limitations on where the vehicle can be pulled over;

If the interdiction data includes the likelihood that the vehicle in fact carries a nuclear weapon,
then how that information is to be used;

How to institute any traffic control measures called for;

How to use any criteria to determine whether to visually inspect, conduct a scan with a hand-held detector, or detain and stand off;

How to conduct an interrogation and visual inspection;

How to conduct a scan with a hand-held detector;

How to detain and stand off, while awaiting arrival of the nuclear weapon specialist team;

How to report back to headquarters, or to the RAD System;

How to use criteria to determine whether to release the vehicle or continue to detain it;

All the reporting requirements, not only those between the interdicting officer and others,
but also among all involved agencies, for every eventuality.

3.) ConOps upon decision to conduct a drive-by scan, if that is part of the overall ConOps:

How to use the information provided (e.g., imagery, license plate information, the time the vehicle left the last radiation detection station) to intercept and identify the vehicle;

How to conduct a drive-by radiation scan, including any measurements to determine if it was a satisfactory scan;

Note that the drive-by offers the opportunity for a visual scan also, so: How to conduct a drive-by visual scan;

How to conduct a second scan, if called for;

How to report back the results of the radiological and visual scans;

If ConOps calls for an autonomous decision in the field whether or not to interdict: How to decide whether or not to interdict;

If ConOps calls for a Threat Assessment Officer decision whether or not to interdict: How to maintain surveillance of the vehicle while waiting to see if there is a request for interdiction;

All the reporting requirements, not only those between the drive-by scanning officer and others,
but also among all involved agencies, for every eventuality;

If the decision is made to interdict, then the ConOps for interdiction, discussed above, becomes operative.

2.3.4 Key Questions for the Interdiction Force

Sections 2.3.1 though 2.3.3 cover the user interface, roles, and concepts of operations for the interdiction force. But there is a set of questions to be asked to characterize the level of performance that would be provided by the interdiction force, and are more basic to designing the RAD System than are the interface, role and ConOps issues discussed in the previous two sections. Those questions, and the answers we collected in our New Jersey meetings:

Q1: What warning, and what visual ID information, does it take to intercept, drive-by scan, and interdict?

A1: Just a few minutes warning would be called for. A more specific answer would require a more detailed description of what would be involved. Simple visual ID information would be adequate. While that was the answer, again, we got the impression that a more detailed, careful examination of the question might have revealed some concerns. Casual observation of the New Jersey Turnpike suggests that unless the visual information is somehow quite specific and unique, such as a very unusual vehicle with distinctive markings, or license plate information, it would be quite easy to intercept the wrong vehicle.

Q2: What would be the cost/manpower impacts of different system configurations, e.g., N interdictions per day, M alerts per day?

A2: Quite hard to get a specific answer to these questions. After much discussion, we concluded that the perspective was: How many interdictions per day, all of which would be nuisance alarms, could be tolerated with the current force. While never getting a direct answer, we concluded that we should work with, as a working assumption, the idea that 10 interdictions per day might be the maximum tolerable. That does not seem an unreasonable number, given the State Police estimate of normally from 50 to 200 interdictions per day. With additional funding and staffing, a higher number could be accommodated. That is only a broad guess based on the discussion. In fact, in an actual deployment, this question and answer should be carefully worked out. As far as alerts per day, we would have to work out more completely what would be involved in an "alert," in terms of operations changes, before a reasonable answer could be expected.

Q3: Could the "beats" followed by each cruiser be revised to keep them within range for responding to requests for drive-by scans and interdiction?

A3: A general answer was that they wouldn't have to revise any beats -- the general pattern of driving by cruisers would keep some cruiser generally in range to respond to a request for a drive-by scan or interdiction. In an actual deployment we would want to investigate actual driving patterns for a more careful check on what response times could be expected.

Q4: What alert levels would the interdiction force want to define?

A4: This question would need a quite extensive orientation before it could be answered in an informed way. Our interviews in New Jersey did not involve enough time, or

enough preparation on our part, to lay out the issues and alternatives well enough for informed answers to be provided.

3.0 System Design Issues

This section presents the system design considerations identified and developed in the course of the study.
They are sorted into three categories: Response Options, Deployment Options and Concept of Operations Choices. Though note that there is some overlap between the Response Options List (Section 3.1) and Concept of Operations Choices (Section 3.3) -- they are two different perspectives on system design issues.

3.1 Response Options List

Many of the findings of the CY03 work can be expressed in terms of a listing of the "Response Options" to be considered in any design of a future system:

3.1.1 Data Collected on Suspect Vehicle

Radiation data. In background data-collection mode that could be more extensive, in terms of radiation types/energies, than in operational mode.
Take pictures of vehicle.
Take extra pictures of vehicle, prompted by radiation alarm (this suggested by a New Jersey Turnpike Authority person).
Experiment with license plate pictures and character inference from that image.
Consider pictures that allow at least a counting of the occupants, e.g., from the back.
Record toll-tag number (in the New Jersey case: E-ZPass).
Record medallion (license to carry radioactive materials) presence, type and number, if there is one.
Record loop detector signature.
Record treadle (axle) count.
Record speed, to be combined with treadle signature to infer wheelbase, and it may be used to process radiation data.
Record the weight of the vehicle.
Cross-reference the weight of the vehicle against its visual characteristics.
Cross-reference the weight of the vehicle against its estimated curb weight (inferred from imagery) and observed number of occupants.
Record vehicle classification, e.g., the New Jersey Turnpike classifies vehicles into eight categories.
For a vehicle that is suspect but allowed to pass through to subsequent radiation detection stations:
use time and speed to calculate time window of arrival at the next station.
For a vehicle arriving at the station from a previous station where it recorded a hit, possibly:
- set up for manual matching of the observed vehicle to imagery from previous

station(s);
- set up for automatic matching on all digital data collected on that vehicle from previous station(s).

3.1.2 Staffing

Who staffs the threat assessment function. Two possible positions, though one option is for the monitoring person could make interdiction decisions without consultation:

- monitors RADS outputs (current concept for NJ State Police: Assistant Duty Officer);
- consults on some or all interdiction decisions (current concept for NJ State Police: Duty Officer).
-

3.1.3 Equipping/Training

Outfitting all cruisers that could be dispatched for interdiction, or only some of those cruisers, with drive-by detectors.

Equipping and training all officers, or only some officers, in RADS interdiction, including drive-by scanning, visual scanning and scanning with a hand-held detector. (Equipment could be assigned to particular cruisers.) One option: In the New Jersey case, rather than train and equip all State Police officers involved, the system could train and equip specialists in the New Jersey Dept of Environmental Protection, who could be called out for each interdiction. Then the system would only have to train and equip 24/7 coverage in that department.

3.1.4 Supporting The Interdiction Decision

The criteria for that decision.

Whether those criteria should be varied by background alert levels.

The system display to the Threat Assessment Officer. For example, it could be simply "Green Light" (no suspicious vehicle) and "Red Light" (request interdiction), or it could include "Yellow Light" (suspicious vehicle -- request interdiction if the TAO has other evidence raising his or her level of suspicion), to several levels of "Yellow Light" indicating different levels of suspicion. In addition, the "Red Light," or even a "Yellow Light," could include some measure of the likelihood that a weapon is in fact on board.

3.1.5 Response Force Readiness

Whether or not there should be alert levels for the interdiction force. That is, if the Threat Assessment Officer determines not to interdict a vehicle, but to continue tracking it, should there be one or more heightened alert level(s) in the interdiction force, and what should be the operational changes associated with each alert level?

Whether or not there should be alert levels of the nuclear weapon specialist team. The levels could correspond with, e.g., a vehicle being tracked but not interdicted, then a vehicle being interdicted but before it is known if the team will be called.

3.1.6 Interdiction, and Additional Scanning, Procedures

Drive-by radiation scan.

Arranging for at least two officers to conduct the interdiction.

Limiting where vehicle is pulled over.
Considering traffic control around the interdiction point.
Procedures once the officer pulls the vehicle over:
- approach and visually inspect;
- approach and scan with hand-scanner;
- do not approach, but call for nuclear weapon specialist team.

What nuclear weapon specialist team is to be called on, based on what evidence.
How rapidly that team is to be dispatched and sent to the site.

3.2 Deployment Options List

Similar to Section 3.1, some of the findings of the CY03 work can be expressed in terms of a listing of the "Deployment Options" to be considered in any design of a future system:

Note: There is some pairing between deployment options and response options. That is, a given deployment option will have associated with it a different set of more effective response options than there would be with another deployment option.

Which detectors, and suites of detectors. In particular: More costly detection to collect more information regarding radiation type (alpha/beta/gamma, energy levels). More costly detection to detect smaller signals. Traded off against a larger number of detection stations with the same budget.

Siting the array of detectors:

First Issue: If in fact it is found that multiple detector passes improves the system's ability to detect a weapon: For a fixed detector budget: Trading off how thoroughly a given route is covered vs a larger number of routes covered.

Second Issue: For a fixed detector budget: Trading off number of routes covered (increased by placing detectors closer in to the targets, in the FY03 case the Holland Tunnel) vs warning time.

While this may seem a short list, other options are best framed as response options or concept of operations choices.

3.3 Key Concept of Operations Choices

Section 2.3.3 above laid out ConOps issues to be handled by the interdiction force. Several of those issues call for more technical work to identify what ConOps makes the most technical sense. Those ConOps issues are best addressed from a systems design point of view, before discussing them with the interdiction force. This section describes the key issues that fall into that category:

Note that there is some overlap between this section and Section 3.1. These two sections represent two different perspectives on system design issues.

3.3.1 General List of Key ConOps Choices

Criteria to request an interdiction:

This would involve consideration of the probability of missing a terrorist attack, weighed against the cost of maintaining a system with a high nuisance alarm rate. In Sections 5.3 and 6.2 we present more complete discussions of considerations in setting those criteria.

Criteria to request a drive-by scan:

This would involve considerations similar to those for requesting an interdiction. It would also involve assessing the benefits of such a system, in terms of reducing the number of nuisance alarms involving a full interdiction, at the cost of less costly, less disruptive drive-by scans.

Guidance to the interdicting officer regarding the likelihood there is an actual weapon in the vehicle:

This entire issue needs to be thought through. First, how much could the system deliver, in terms of any actual perspective on the likelihood of an actual weapon, and how different would that information be among interdictions? Second, how would the interdicting officer use that information? Note that we suggest different ways that information could be useful in the next subsection.

Where to pull over the vehicle:

This question would involve technical considerations, such as how much population risk could be mitigated by restrictions placed on where to pull over the vehicle, balanced against the feasibility of controlling where to pull over the vehicle, against any added risks of delaying the pull-over past the earliest possible time, and finally, the added burden on the interdiction force for the additional time involved for each interdiction, given that almost all interdictions will be nuisance alarms. Clearly, interdiction officers should be included in this ConOps development, since they are the experts and users of this particular ConOps. Then political considerations could come in to play, and could become dominant. This issue might best be addressed by conducting carefully structured panel sessions with expert panels, to settle the technical considerations, then with that preparation, if still necessary, convene panels of political representatives to address political issues. Frankly, the fact that almost all interdictions will be nuisance alarms suggests that, after all tradeoffs are considered, especially burden on the interdiction force for each interdiction, the vehicle be pulled over as soon as is safe, unless strong evidence exists that the vehicle is carrying a nuclear weapon. This issue illustrates one case where the likelihood that the vehicle is carrying a nuclear weapon could be useful.

Traffic control measures to accompany any interdiction:

This question would involve technical considerations, such as how much population risk could be mitigated by traffic control measures, balanced against the associated traffic congestion burdens, and societal peace-of-mind considerations. Again, this issue may be dominated by the fact that almost all interdictions will be nuisance alarms. Again, this may only become a consideration for special cases where strong

evidence exists that the vehicle is carrying a nuclear weapon, and so another example of the usefulness of likelihood-of-weapon data.

Criteria for visual inspection versus scan with hand-held scanner versus detain and stand off:

This question would involve technical considerations regarding the effectiveness of visual inspection versus scanning with hand-held scanner, balanced against risk to the officer and the neighborhood around the interdiction. This would involve laying out the sequence of operations between possible early actions by the intervening officer and later actions by the nuclear weapon specialist team, all assessed against the background of likelihood of an actual weapon. Again, the criteria would vary with likelihood of an actual weapon.

Hand-held scanner scanning procedures, including safety measures:

These procedures could be developed using trial scanners, vehicles and simulated weapon sources. Officers could be consulted regarding existing procedures for approaching suspect vehicles. The risks involved and safety measures and could be of paramount importance.

Visual search procedures, including safety measures:

These procedures could be adapted from what are probably existing interdiction force procedures for visual search. While a visual search could be more effective than a scanner search, considerations of civil liberties and intrusiveness could be important. That suggests that these procedures, and the choice between visual and scanner search, should be informed by consultation with legal authorities.

Stand off procedures:

Again, there are probably existing interdiction force procedures that should be reviewed.

Nuclear weapon specialist team:

This is an area that has not yet been examined, and was considered beyond the scope of the current project. However, there are four key issues concerning that team that should be directly integrated with the interdiction force ConOps:

- Criteria determining when the team would be called in;
- What would be their operational response time, i.e., would they be on call for immediate roll-out, and helicopter, versus surface transportation with sirens and lights, versus surface without sirens and lights. That response time would have important impacts for the effectiveness of the system.
- What team, exactly, would be called in;
- What would be their concept of operations, i.e., what would they actually do with the vehicle and its occupants.

The last two issues are key in that they in part determine what functions are left up to the interdiction force.

3.3.2 Special Issue: Uses of likelihood-of-weapon data.

The system could display to the Threat Assessment Officer different levels of likelihood that the suspect vehicle is carrying a weapon, yet still recommend interdiction. If the signal is low enough to be plausibly within the range of legitimate sources, but still worth

interdiction, then the likelihood that the vehicle is actually carrying a weapon would be less than in the case where the signal is higher than could be expected from any legitimate source. Actual differences in procedures between the low-likelihood and high-likelihood cases would have to be developed. But here, for example purposes only, are some possible differences in procedures that could be considered:

Low likelihood the vehicle is carrying a weapon:
- No drive-by scan.
- Vehicle pulled over as soon as is safe.
- No traffic control measures.
- Officer conducts visual inspection and interrogation, then scans with hand-held scanner as necessary.
- Nuclear weapon specialist team not notified unless the officer determines it should be.

High likelihood the vehicle is carrying a weapon:
- Nuclear weapon specialist team notified upon threat assessment decision and launched toward likely pull-over site. Very high likelihood could warrant the use of a helicopter to reduce response time.
- Drive-by scan, as much for a visual scan as for a radiological one.
- Vehicle pulled over in an area intended to minimize impacts as much as possible.
- Traffic control measures are instituted.
- Officer conducts visual inspection, interrogation, attempts to maintain that the interdiction is due to a speeding violation, tail light out, etc. Detains vehicle, while standing off, until nuclear weapon specialist team arrives.

If nothing else, these example procedures should make clear that there are a number of issues to be thought through for system implementation.

3.4 Overall System Performance

3.4.1 Measure of Overall System Performance: p(interdiction | weapon)

With the previous sections as background, it is now possible to characterize overall RAD System performance. That performance can be measured in terms of probability of interdiction given a weapon in the vehicle, p(interdiction | weapon). Note that that is carefully chosen to match the scope of the RAD System, and what it can and cannot control. It can only call for and carry out an interdiction until the nuclear weapon specialist team arrives-- We can treat system performance issues, for now, as if the system can't control the success or failure of the incident after the nuclear weapon specialist team arrival. That is not completely true, since aspects of the system, such as where the vehicle is pulled over and what the interdicting officer does in his or her phase of the interdiction can impact the overall success of the incident, but that involves details of procedures that were not pursued in the FY03 effort, and so will be set aside for now.

An alternative measure would be p(interdiction | weapon signal strength after shielding). That measure would "give the system a break" regarding what it can and cannot control. That is, all the system can "see" is the weapon signal after shielding. But as will be explained in Section 5, important parts of the system can be designed to reduce the interdiction threshold on the signal the system sees, while maintaining a particular nuisance alarm rate. So we need to go to a "higher level" of system performance measure, p(interdiction | weapon), which is, after all, what we ultimately care about. Note that the RAD System leaves room for incorporation of other intelligence and data, but that can only work to the system's benefit (probably), and is beyond the control of the system, so can be left outof considerations of system evaluation.

3.4.2 Areas of system design/performance that affect that overall performance

There are four general areas of system design and performance that affect overall system performance. Each of those areas involves very different aspects of system design:

Area 1: Detection At Each Station:
This involves two quite different areas of development:
- technology choice, among types of detectors, detector size and location.
- possibly setting up a medallion system, as discussed in Section 5.3.2.

Area 2: Allocation of Detection Along Route and Among Routes:
This is a matter of system assessment. It depends on the results of investigation of the benefits of multiple looks, discussed in Section 6.2. It involves system-wide tradeoffs between different dimensions of coverage, and coverage versus warning time. For example, if the goal is protection of the Holland Tunnel with a fixed number of detectors, those detectors may be arrayed with several detectors along each of a few routes, or one detector per route, covering many more routes. Also, the closer to the entrance to the

Holland Tunnel, the more thorough the coverage of routes with a fixed number of detectors, but at the cost of shorter warning time.

Area 3: Threat Assessment:
This involves several different system design features:

- the processing of raw system outputs into decision aids / decision-aiding displays;
- the assembling of other data that could be of use to the Threat Assessment Officer;
- the training of the Threat Assessment Officer;
- the organizational system of staffing the Threat Assessment Officer function, which would involve a combination of direct monitoring of RADS output displays and other related information, and possibly calling in a more trained person on short notice, to offer a more trained opinion.

Area 4: Interdiction:
This is largely an organizational process with a number of features:

- Investigate how much the interdiction force can be persuaded to conduct N interdictions per day, where almost all of them are going to be nuisance alarms.
- Consider additional funding and staffing, so that the interdiction force could have the resources to support a larger number of interdictions per day.
- Consider combining the interdiction function with other functions, such as vehicle inspections, and certifying that the vehicle has the appropriate permits for carrying radioactive material. Costs of the system could be defrayed by fines levied for infractions discovered.
- Consider multiple alert levels, where the interdiction thresholds are adjusted such that under low alert only very strong signals cause an interdiction to be requested, perhaps such that only about one interdiction per day is requested, while at higher alert levels the thresholds are lowered so that perhaps 50 interdictions are requested per day, but only on those high alert days. The system as notionally characterized here would be easily capable of exactly those adjustments to thresholds. As experience is gained, system operators will know very accurately (but with a known uncertainty) how many interdiction requests they can expect for any given threshold setting. The next section discusses how in a system with different detection capabilities for different radiation types/energies, interdiction thresholds can be adjusted separately for each radiation type/energy.

3.4.3 Intrinsic Advantage of RADS vs a Portal Monitoring System

A key system cost is maintaining an interdiction force, typically highway law enforcement, ready with a very short response time. With RADS, that is efficiently accomplished by being set up such that interdiction force officers on their standard patrols can be in position to respond quickly to a request to interdict. While a portal system, on the other hand, must involve an interdiction force that must be kept in close proximity to the portal, and so either be dedicated to the system, or be limited in what other duties they can perform. Either system must maintain preparedness for a very rare event, so that a system that “ties up” the response force with little or no capability to

perform other functions is quite costly relative to a system, such as RADS, that allows the interdiction force to perform their normal duties at all times.

3.5.0 Key System Challenge: Signal Detection Out of Legitimate-Source Background

3.5.1 The Challenge

The purpose of the RAD System is to detect a nuclear weapon. For clear reasons, system detection performance will be a function of the weapon signal strength. A given system will have a high probability of detecting a weapon with a very large radiation signal, but then that probability must decline for weapons with smaller radiation signals, due to size or shielding. That's the detection performance side. Now on the cost/burden side: Aside from the system capital and operating costs, the most important "cost" of operating the system is the burden it places on the interdiction force to interdict suspect vehicles. That is a burden because almost all, hopefully all, of those interdictions will be nuisance alarms. That is, they will be interdictions of vehicles that have a radiation signature, but the radiation source is legitimate.

Tying those concepts together, we can state the key systems challenge of RADS succinctly: Achieve the best possible system performance, as measured by p(interdiction | weapon), with an acceptable nuisance alarm rate.

3.5.2 Signal Detection Theory

Starting with the basics: Typical signal detection theory (SDT) discussions begin with two bell curves, offset but overlapping. The one on the left is "Noise," the one on the right is "Signal + Noise." Then you draw in an detection threshold where the bell curves overlap, then measure the probability a system with that threshold will detect a true signal given a signal, and the probability it will sound a "false alarm" (detect a "signal" which in fact is generated by the "Noise" distribution). But in the RADS case, the situation is importantly different.

What is the "Noise" distribution for RADS? For clarity, in the rest of this report we call that "legitimate-source background," but in this SDT discussion we'll refer to it as "noise." For the proposed detector locations on the New Jersey Turnpike, the average daily volume for the detector seeing the most traffic (Station 3) is about 109,000 vehicles per day (for 2002, data reported directly from New Jersey Turnpike Authority via an Open Public Records Act request). But we can assume that almost all of those vehicles have no radiation signature, and so will be easily rejected by the detector system once it is appropriately calibrated. The "Noise" distribution, in the SDT sense, for RADS is the population of vehicles that carry legitimate radiation sources. Those would be vehicles carrying medical sources, well-logging equipment, nuclear density testers (aka soil density testers), food irradiation sources, persons undergoing certain medical treatments, and perhaps other nuclear sources we simply don't know about. What is that population?

That is, how many vehicles with what sort of legitimate radiation signatures would pass by a RADS detector in a day? We don't know. In the same Open Public Records Request referred to above, we asked the NJTA for the number of permits to carry radioactive sources it issues in a year. The answer for 2002: 15. We can assume that that number is not a good start at estimating the number of vehicles per day that would pass by a RADS detector with a legitimate radiation source.

What is the "Signal + Noise" (S+N) distribution for RADS? One might be tempted to assume that there is no S+N, since the system can be considered to be phase-locked vehicle by vehicle (i.e., it views the world as discrete potential sources, one per vehicle). But in fact, there might be adversaries smart enough to carry a legitimate source along with the weapon, and so there could actually be an S+N vehicle. But more likely, the adversary vehicle would only have a weapon source, with a radiation signal of some magnitude, and so we should consider an "S" distribution as well. In either case (S+N or S), the RADS roadside detector will report any vehicle with a radiation signal above some threshold, and leave it up to the system response (interdiction, etc.) to use other means to discriminate legitimate sources from a weapon. (Though Section 5.3.3 below discusses the possibility of using radiation type and energy to aid discrimination.) In any case, the S+N or S distribution for RADS has not yet been developed. We could generate a list of possible weapons and then deduce the possible signals, accounting for any of a number of shielding strategies, but we will now suggest a simpler approach.

Figure 2 presents a completely hypothetical "Noise" (legitimate-source background) distribution. We can simply run the RAD System in normal operation for some long time, and collect background information such as that indicated in Figure 2. The numbers are notional only. For example, we suppose an average of one vehicle per day will pass the sensor with a signal strength of "19," while 20 vehicles per day will pass the

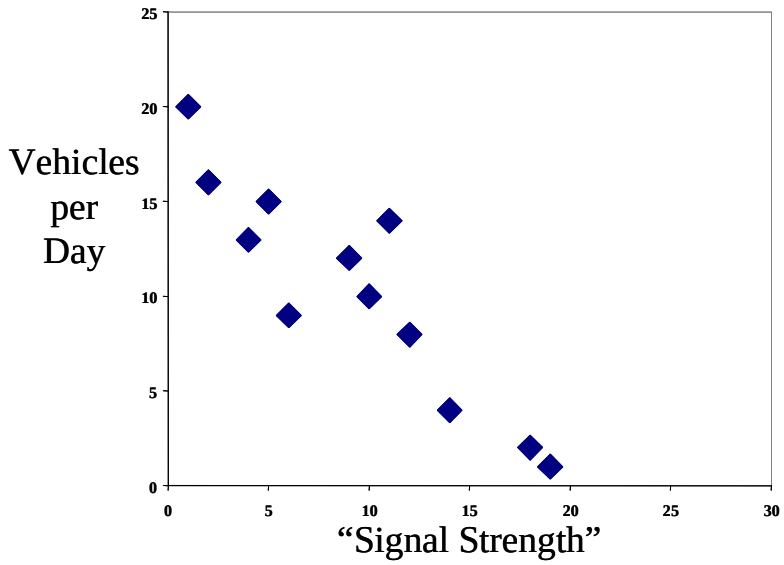

sensor with a signal strength of "1."

Figure 2. Hypothetical "Noise" (legitimate-source background) distribution.

We won't present a "Signal" or "Signal+Noise" distribution, because we know so little about it, but more importantly, because we can characterize ways to address the "challenge" presented in Section 3.5.1 above without needing to know that distribution, as will be explained in Section 3.5.3.

3.5.3 Three approaches to Address the Challenge

Approach 1: Collection and use of data on legitimate-source background.

We assume we are gaming against an intelligent adversary who knows about shielding. So we can assume he or she will employ some level of shielding. Given that, we can assume that whatever the S or S+N distribution is, the system will have a higher p(interdiction | weapon) if we can set the threshold for requesting an interdiction lower. But lowering the threshold, even if it is technically feasible, comes at a cost, as it always does in SDT: a cost in higher "false alarms," or in this case, nuisance alarms. For example, with the background presented in Figure 2, if we set a threshold at signal strength = 20, we would get no nuisance alarms, but we would only request interdictions for relatively "bright" (over "20") sources. We could set the threshold at signal strength 13, and detect less-bright sources, but would then have to tolerate an average of 7 nuisance alarms per day. A threshold of signal strength 8 would detect even better-shielded weapons, but at a cost of 51 nuisance alarms per day.

But that is one way to address the challenge: Simply run the RAD System for a period of time, to collect the legitimate-source background data depicted in Figure 2, then combine that with information as to what rate of nuisance alarms the interdiction force would tolerate, to set the threshold. If we suppose a tolerance level of nuisance alarms of 10 alarms per day, then according to the example data in Figure 2 we could set the threshold (given no special intelligence) at signal strength = 13. Note that we can exploit the rare-event character of this detection challenge to assume all detector hits in the baseline-collection phase are nuisance alarms. Strictly speaking, we will be pretty confident that they are all in fact nuisance alarms if no adversary nuclear weapons are discovered for some period of time after the data collection period.

Figure 3 presents the relationships between the legitimate-source background, the acceptable nuisance alarm rate, and system performance (notionally = p(interdict | weapon)). Reading the figure from bottom to top:
1.) Legitimate-Source Background Spectrum: Is, at first, a given. It just has to be measured.
2.) Acceptable Nuisance-Alarm Rate: Is a matter of what the interdiction force will tolerate.
3.) Interdiction Threshold: Follows from the first two elements.
4.) Threat Spectrum: We don't know what it is, but we know its rough outlines.
5.) p(interdiction | weapon): Without the Threat Spectrum, we can't know its value, but we know it gets higher as the Interdiction Threshold gets lower.

Note the logical relationships revealed in Figure 3: The legitimate-source background, at least initially, is a given. Given that, there is a direct relationship between acceptable nuisance alarm rate and system performance. We know that even though we don't know the threat spectrum. That is, even in the absence of knowledge of the threat spectrum, so we can't actually estimate p(interdiction | weapon), we can still see that the lower the interdiction threshold, the better the system performs as it would be measured by p(interdiction | weapon), if we could measure that. In Section 5.3.2 we will see how we can shift the legitimate-source background "down and to the left," and so improve system performance even at a fixed acceptable nuisance-alarm rate, in a way made clear by Figure 3. In Section 5.3.3 we will see how more sophisticated radiation measurements have the potential to effectively reduce the interdiction threshold against the threat spectrum, but in a way not reflected in Figure 3.

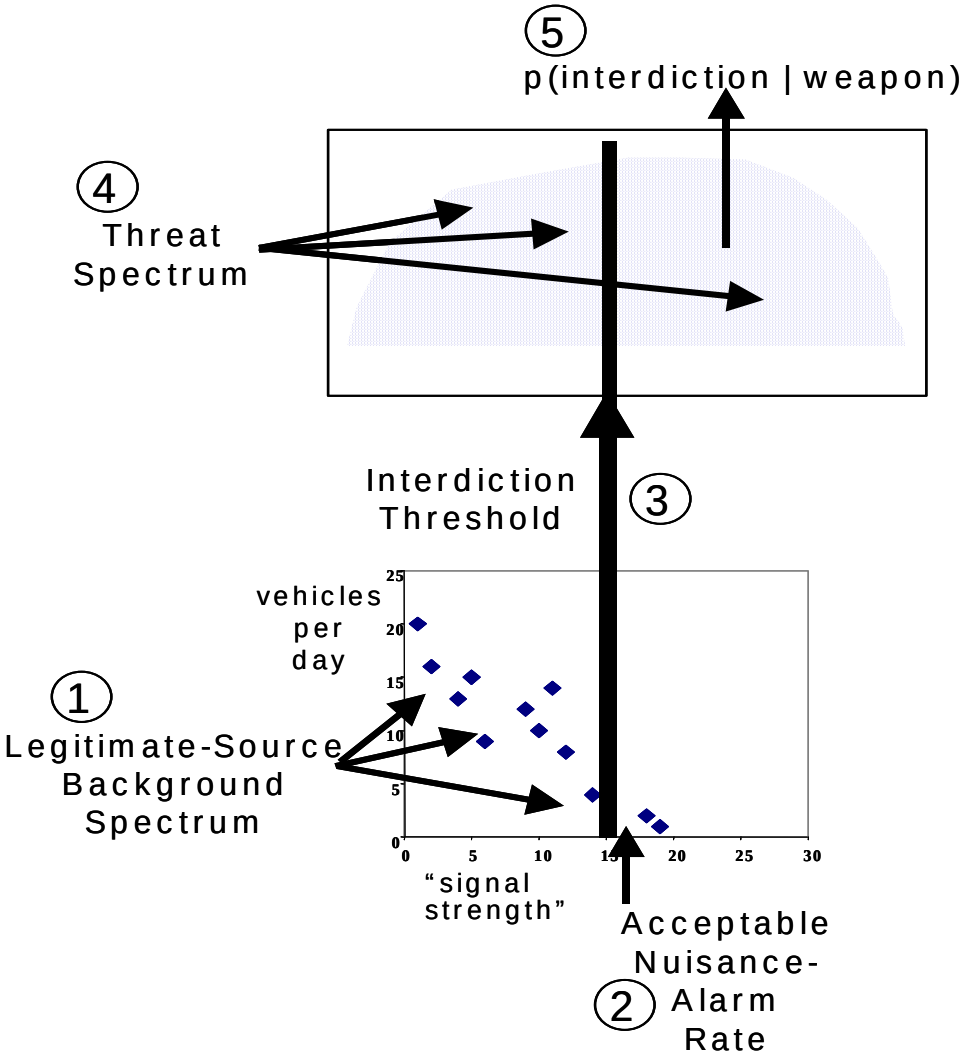

Figure 3. Key Signal Detection Theory Relationships

The initial period of system operation can involve only baseline data collection, with no interventions requested. That data can be used for detector design and optimization.

Then the system can begin requesting interventions, while baseline information collection can continue throughout system operation. As time goes on, more detailed information about the legitimate-source background can be collected, and so more sophisticated detection systems (detectors and algorithms) investigated.

Note that, using the empirical background information, the system, or the Threat Assessment Officer, can shift the threshold for best effect. For example, during times of high alert the interdiction force could be persuaded to tolerate 50 interdictions per day, essentially all of them nuisance alarms, and so the threshold could be set to signal strength = 8, given the hypothetical data of Figure 2. Alternatively, if the TAO has particular intelligence concerning white vans, he or she can set the threshold perhaps at as low as 8, then, using the imagery collected with each detector hit, screen those detections with signal strength between 8 and 13 to only request interdictions for white vans. Turning this thought around into a system specification, a recommendation for RADS implementations is to enable the TAO to shift thresholds easily from the control console.

Thresholds could be adjusted for special vehicle-types. For example, suppose trucks carrying large loads of granite are found to produce a large signal relative to a small well-shielded weapon. A policy decision could be made to reject those detections based on accompanying imagery or weight information, though note that that would leave a strategy open to the adversary, if such a policy became known. That in turn raises the issue of the need for security in descriptions of the algorithms for requesting interdictions, and the desirability of "concealing" those algorithms within decision-aiding software.

Approach 2: The medallion concept.

Empirical study could find that the legitimate-source background makes necessary a quite high interdiction threshold necessary to keep nuisance alarms at an acceptably low rate. In that case, system performance, in terms of p(interdiction | weapon) could be enhanced by launching a program of requiring automatically-readable "medallions" on vehicles carrying legitimate radiation sources. That term is classically associated with taxi cabs, as a tag indicating that the vehicle is licensed to participate in a particular, regulated activity. In this case, the medallion could be an electronically readable tag, like a toll tag, attached to a vehicle as a certification that that vehicle is licensed to carry radioactive materials. It could be specialized in terms of radiation type, energy, and signal strength. It could be attached in such a way that it could not be removed to another vehicle and still work. It could have an expiration date built in to its returning signal.

A medallion system would then have the advantage that medallioned vehicles could be automatically removed from the detection process. The net effect would be to "lower" the spectrum of Figure 2, making possible a lower interdiction threshold for a given nuisance alarm rate. As medallions would be introduced, the system could automatically observe the lowered legitimate-source background (after medallion deletion), as medallioned vehicles become automatically deleted from the detection process. In fact, the RAD System could be operated in such a way that nuisance alarms can always be

kept under, say, ten per day. As medallions penetrate the population, the interdiction threshold could be lowered accordingly.

Thinking carefully through the SDT logic presented in Sections 5.2 and 5.3.1, it would make the most sense to medallion the largest sources first. Further, we would want to try to implement as universal a system of medallions as possible for the larger sources before imposing medallions on smaller sources. The logic there: Medallions are of the most help when they allow a lowering of the no-alert/no-special-intell interdiction threshold while keeping the nuisance-alarm rate at an acceptable level. So if that level is, say, ten interdictions per day, as long as there are ten non-medallioned vehicles per day above a certain threshold, medallions on sources below that threshold would not help lower the threshold, except in the special cases described before of elevated alert, or intelligence about a specific vehicle type.

While for clarity and motivation we began this section with the case of the legitimate-source background requiring a "high" interdiction threshold, in fact for any legitimate-source background at all, medallions would enable a lower interdiction threshold, and as discussed before, in a game against a shielding adversary, the lower the threshold the better the system effectiveness, i.e., p(interdiction | weapon).

Note that the medallion concept introduces an essential change in the nuisance alarm issue: It would change each interdiction from being simply a "nuisance," to being a medallion-enforcement stop, including perhaps the levying of a fine that would help defray the associated interdiction costs. The fine would also be an incentive to purchase a medallion to prevent future fines. The system would naturally have a quite high enforcement rate for vehicles emitting a signature above a certain level, and that level would be reduced as the system "rides the noise spectrum down," keeping the nuisance alarm rate at its target value, as medallions penetrate the population.

The medallion concept raises other issues: We can suppose an appreciable fraction of radiation sources on the New Jersey Turnpike northbound originate outside of New Jersey. So the medallion system would be larger than New Jersey in scope. Essentially, New Jersey would be saying, "If you want to transport nuclear material in our state, you must apply for a permit and carry the associated medallion, or else you may be pulled over and cited." That is not an unreasonable position for a state to take for any hazardous material. In fact, the National ITS Architecture (Version 5.0) includes features for control of transport of commercial vehicles, e.g., Process Spec 2.1.1.2 - Monitor Commercial Vehicle Route (http://itsarch.iteris.com/itsarch/html/pspec/p0120.htm), that are not dissimilar in impact to the medallion concept presented here. More generally, a medallion system would have benefits in terms of hazardous materials regulation that would be in addition to the benefit of increased RADS detection performance at a given nuisance alarm rate.

The medallion concept raises an essential question: Is it appropriate for a society to impose a system of medallions for regulating the transport of radioactive materials, given concerns for adversary actions, even though those actions would be quite rare?

Addressing that must combine both technology and policy. That is, technologists must assess how much can be gained in detection effectiveness by such a medallion system, then the political process must decide if the additional security is worth the societal and private-sector costs. This is the classic engineering economics cost effectiveness question, but in the broader scope of societal benefits and burdens.

Note that there would be pronounced advantages to having a serial number unique to each medallion that would be included in its report-back signal, so that, e.g., if a medallioned vehicle is reported stolen, its medallion could be flagged by the system for no-discretion interdiction, independent of the radiation reading from the vehicle.

Note that if a medallion system succeeds in reducing the signal strength at which the system would trigger a threat assessment, then at some point the system may need to change detectors to detect that smaller signal.

Note that there will probably always be some legitimate sources where it would not be practical to enforce medallions. Example: People undergoing certain medical treatments. Those cases put an upper limit on the number of legitimate sources that can be automatically deleted from the detection system, and may place a lower limit on the interdiction threshold. While even in those cases temporary medallions could be issued with short times to expiration, that would increase the vulnerability of the system to stolen or counterfeit medallions.

Approach 3: The concept of detection of different radiation types/energies.

Another type of information provides a more or less orthogonal approach to improving p(interdiction | weapon): The different radiation types/energies to be found in the legitimate-source background, and to be found in the threat spectrum. With an understanding of the radiation types/energies to be expected from the threat spectrum, i.e., from the range of nuclear weapons and shielding strategies that could be expected, combined with a collected baseline background of legitimate radiation sources, system designers can specify detector parameters that would be most effective at lifting weapon signatures out of the noise of legitimate sources.

Putting those thoughts into a project planning perspective: Initial baseline data collection can involve a pilot set of detectors that measure several radiation types/energies combinations. Information gained in that phase can be used to identify less expensive detectors for wider deployment.

A system designed to detect different radiation types/energies can have different interdiction thresholds set for each radiation type/energy. Those different thresholds can be set with an analysis of the threat spectrum, so that the maximum p(interdiction | weapon) can be achieved for a given number of interdictions per day. Then if intell alerts the system operators to a particular threat with a known radiation type/energy-range, the interdiction threshold for that radiation type/energy-range can be differentially lowered and the others raised, to maintain the same number of interdictions per day, but now

focused on the identified threat. Or of course the other thresholds can be left unchanged, raising the number of interdictions per day, but only raising that number for radiation signatures characteristic of the identified unusual threat.

Analysis of background in different radiation types and energies can be linked to an analysis of the likely costs and penetration of a system of medallions for legitimate radiation sources. In one direction, the analysis just described can result in recommendations for a medallion system focused on those legitimate sources that most look like threat signatures. In the other direction, an investigation of which medallion systems would be most feasible would lead to recommendations for detectors focused on lifting weapon signatures out of the noise of legitimate sources after those sources are partially reduced by the most-feasible medallion system. Finally, as the medallion system penetrates the population of radiation carriers, the detectors could be altered to do the best job of lifting weapon signatures out of the population of legitimate sources that remains after medallioned sources are eliminated.

But it remains an open question if detection of different radiation types/energies would be cost effective. That would have to be investigated in systems studies trying out the ideas presented in this section.

3.6.0 Other Analyses and Issues to Consider

3.6.1 Analysis of the Benefits of Multiple Looks

A key question for system design is whether or not “multiple looks” helps. That is, can the system do a better job of detecting a weapon if the same vehicle is examined by multiple detector stations. If not, then detectors are best allocated to cover the most routes, one detector per route. If multiple looks do help, then more than one detector may be invested in certain important routes.

The key issue for this question: How much noise in the detection process is introduced by the speed and geometry of the car-sensor path variations among different detection-station pass-throughs by a given vehicle. If those factors are identical from one pass-through to the next, then we can expect identical readings and get nothing out of multiple looks. If those factors vary importantly by pass-through, then multiple readings can reduce that noise and lead to a better idea of the radiation actually emitted from the vehicle. Strictly speaking, that question could be answered by careful geometric modeling. But empirical study is probably called for.

Note that a better measurement of the radiation emitted from a vehicle, if that could be gained by multiple looks, still only helps marginally in the basic SDT challenge discussed in Section 5. It simply provides a less noisy measurement of the radiation emitted -- it does not otherwise help lift a weapon-carrying vehicle out of the background of legitimate-source vehicles.

Note there are reasons to have multiple looks other than pure signal detection: It could assist in picking the car out of the crowd, if the imagery/license plate information is insufficient to identify the vehicle for the interdicting officer, who may not be able to begin searching until several minutes after the suspect vehicle has left the first detection station. In that case, an interdicting officer could be stationed at the second detection station, then that station could signal the waiting interdicting officer when the suspect vehicle passes through. If there is toll tag information, that second identification could be done by any toll-tag reader, without radiation detection. If trial runs find that in fact the interdicting officers have difficulty identifying the suspect vehicles based on the information the system can provide after the first detection and the delays involved, then a "second-look" system may have to be set up simply to overcome that difficulty.

3.6.2 Value of Information (VOI) Analysis

There are specific evaluation techniques, under the heading of "Value Of Information" (VOI) analysis, from the field of decision analysis. Those techniques analyze how additional information can be used to reach better decisions, evaluates the incremental value of those better decisions, then assigns that incremental value to the information, hence "Value Of Information." There are some information-collecting features that could be part of RADS, and those could be subjected to VOI analysis. That is, a VOI analysis could be done on the information made available by an information-collection feature. In this case, something less than a full VOI analysis may be called for, since a full one might require a full probabilistic treatment, and there would be severe difficulties in estimating probabilities of some adversary actions. There are many aspects of the RAD System that could be subject to VOI analysis. Here are four examples of information-collection questions that could be addressed with VOI analysis, selected to illustrate four different ways that analysis could be used:

Example 1: Starting from a basic detector, would it be cost-effective to add detection capability to differentiate between radiation types/energies? An analysis of the legitimate-source background and threat spectrum would allow a VOI analysis of the benefits of detector suites that could discriminate particular radiation types/energies, which could then be traded off against the incremental costs of those suites. As discussed in Section 5.3.3, this analysis could be linked with an analysis of the costs and feasibility of particular medallion programs. Then this analysis could be extended to an evaluation of differentiating between radiation types/energies in a medallion environment. Again, a complete, probabilistic VOI analysis would almost certainly not be feasible, but some form of it could be worthwhile.

Example 2: If the analysis of the benefits of multiple looks (Section 6.1) finds that there are benefits, then a VOI could be applied to determine if those benefits are worth investing a second detector station on a single route, compared to covering two routes with one detector each.

Example 3: Would it be burden-effective to conduct a drive-by scan? Again, a full VOI would almost certainly not be feasible, but an orderly look at how much better the system could determine whether or not to conduct a full interdiction by using a drive-by scan could inform the process of developing decision criteria between requesting no interdiction, a drive-by scan, or going directly to requesting a full interdiction.

Example 4: Taking the VOI idea one step further: Combining analysis and some empirical validation, we could find that an algorithm that combines weight, axle spacing (based on treadle signals and measured speed), vehicle type (automatically typed by imagery) and loop detector signature could provide a good index of likelihood of a weapon. The reasoning would rest on the idea that a weapon and shielding would have an unusual pattern of size, weight and loop detector signal (i.e., ferro-magnetic signature). Background data collection could work at a pattern-recognition level to identify profiles that are seldom found in background that would be indicative of a weapon, perhaps simply high weight per ferrous signal other than bulk trucks (sand, liquids, etc.). Thus a vehicle with such an unusual profile would be flagged, and marked for recommended interdiction on the basis of even a very low radiation signature, or even possibly, based on analysis, no radiation signature at all. We could then conduct a VOI analysis on that system: Is the cost of that suite of sensors justified by the improved decisions the algorithm supports?

3.6.3 The need to interact with legal agencies.

One significant discovery made during interviews: Deployment of a RAD system would require interaction with state legal agencies, since there would be a need to establish the appropriate legal framework for interdictions based on the evidence collected by the RAD System.

3.7 Three biggest system challenges

The design of a system to effectively detect terrorist weapons being transported in surface vehicles with roadside radiation detectors is quite challenging. The three biggest challenges are interrelated:

1.) We have to assume that we are gaming against an intelligent adversary, so weapon signatures are apt to be weak after shielding, and such a signature must be lifted out of a ubiquitous background of legitimate radiation sources. The signal detection challenge is such that, without using the system features listed in the next section, we may find only unattractive hit rate / nuisance alarm rate pairs.

2.) The system must rely on an existing interdiction force to actually carry out the interdiction. Because of that, there can be significant resistance to a system that involves high nuisance alarm rates.

3.) The signal being sought is, one would hope, quite rare. So the system in typical operation will result in almost entirely nuisance alarms. So the system in typical operation imposes burdens on the interdiction force in terms of interdictions that

turn out to be nuisance alarms essentially all the time, in exchange for the capability to detect (at a probability less than 1.0) a nuclear weapon which can be expected to be quite rare.

3.8 System Features to Address Those Challenges

The stated set of challenges indicates that we should consider every information, information-processing, institutional and ConOps feature available to us to design an effective system. Those features, discussed in this report, include:

1.) Making other data available to the Threat Assessment Officer, such as imagery, medallion presence, loop detector signature, treadle (axle) count, vehicle weight, vehicle classification, any calculated indices of likelihood of a weapon, and intell, so that he or she can make the most effective decisions possible, to request interdictions focused on the vehicles with the highest likelihood of carrying a weapon.

2.) Working to provide system results to the Threat Assessment Officer in as clear a decision-support format as possible. That could involve decision aids to process the information to make the decision task as clear as possible.

3.) Working closely with the interdiction force to design the system to minimize operational burden.

4.) Developing a database of the legitimate-source background, then combining that with assessments of the threat spectrum, to design detectors / detector suites to most effectively lift weapon signals out of the noise of legitimate sources. (Section 5.3.1)

5.) Investigating the use of non-radiological information, such as weight, wheelbase, vehicle classification and loop detector signature, to assist in discriminating vehicles more likely to be carrying a weapon. That would entail both background measurements and some characterization of weapon-carrying vehicle signatures on those measures.

6.) Considering launching a medallion system to install remotely readable tags on vehicles carrying legitimate sources of radiation, enforced by the RAD System itself, fines, etc.. While seemingly unattractive in administrative burden, such a system may be found to be necessary for acceptable system performance. Though that remains an empirical question.

7.) Considering designing that medallion system such that what would be considered nuisance-alarm interdictions could in fact become medallion-enforcement interdictions. Such an interdiction has two advantages:

   1.) Each interdiction serves a useful purpose even if no weapon is found, and so is less apt to engender resistance by the interdiction force officers.

   2.) Each interdiction could involve the levying of fines, which could encourage compliance and defray the incremental costs of each interdiction.

8.) Including in the detection and (perhaps) medallion processes the consideration of a system that can differentially detect different radiation types/energies, optimizing the system by setting different interdiction thresholds for the different radiation types/energies, so that radiation type/energy information is used to focus

interdictions on the vehicles with the highest likelihood of carrying a weapon. (Section 5.3.3)

9.) Carrying out the interdiction threshold optimization in a way intelligently linked to the process of enforcing a medallion system. (Section 5.3.3)

10.) Intelligently allocating the detectors among routes and stations along each route, to get the best protection from a given number of detectors. This would depend upon an analysis of the benefits of "multiple looks," i.e., having a vehicle pass through multiple detectors. (Section 6.1)

11.) Using value of information analysis to optimize what information is collected and used. (Section 6.2)

12.) Combining all the features listed here in an integrated systems optimization process to achieve the highest probability of interdiction given a weapon, given the willingness of the interdiction force to tolerate a certain number of nuisance-alarm interdictions per day.

4.0 Simulation of Target Vehicle Tracking/Interdiction

The objective of this part of the Model Cities project is to develop and employ a simulation of a specific section of the New Jersey Turnpike and certain adjacent roadways, using the Paramics microscopic traffic simulation package. The simulator will be used to illustrate and demonstrate the feasibility of a new detector system.

The following tasks were identified at the beginning of the project:

- Model road geometry for a 50-mile section of the New Jersey Turnpike, focusing on the northbound direction towards New York City;

- Model four sensor stations at specified locations. Sensor stations will detect certain properties of a specific target vehicle.

- Simulate one (or more) target vehicle and its interaction with the network and sensor stations

- Base the simulation on a fixed path for the target vehicles between specified origin and destination

- Implement plausible traffic conditions on the New Jersey Turnpike for purposes of demonstrating the simulation capabilities. Different types of traffic conditions would be desirable, such as free-flowing and congested conditions.

- Develop and implement graphical displays to facilitate the visualization of target vehicles as they progress through the network, their states (as related to readings from the

sensor stations), the identity of the vehicle and elapsed time. A forecast mode is desirable, with an estimated time for the arrival of target vehicles

- Perform a set of simulations based on a defined typical scenario. Document the work in the report and a set of movie animations to be used for demonstration purposes.

4.1 Methodology

Once the general goals of the project were defined, it was necessary to identify the geographical area to be covered, the precise location of the sensor stations, and the details of the scenario to be simulated.

It was agreed to focus on the northbound direction of the New Jersey Turnpike (I-95), for the section between Interstate 195 and the Holland Tunnel, a distance of about 50 miles.

Along this section of roadway, four sensor stations were to be implemented in the simulation. The locations are as follows:

- Station 1: At the tollgate entrance to the Holland Tunnel at the Toll Plaza
- Station 2: At the tollgate exit east onto I-78 from the NJT at Interchange 14
- Station 3: On the NJT north of the overpass at Interchange 10
- Station 4: At the NJT tollgate entrance at Interchange 7A (at the junction with I-195), 50.6 miles from Station 1.

The scenario to be replicated in the simulation would have the following characteristics: the target vehicle will first enter the NJT at Station 4, driving north. The vehicle will pass through each of the other stations in descending order, activating a series of graphical and analytical events to represent the process of detection, tracking, travel time forecasting and intercepting.

4.2 Overall network description

The study area, shown on Figure 4.1, includes approximately 50 miles of the I-95 New Jersey Turnpike from the junction with Interstate 195 Freeway, to the entrance of the Holland Tunnel. The modeled network contains both directions of the I-95 freeway, all interchanges and a number of adjacent major arterials. The network boundaries were chosen after the scenarios to be studied had been identified. In all scenarios, target vehicles to be detected are traveling northbound on I-95, successively passing by four detector stations on their way to New York City via the Holland Tunnel.

The network was coded in the latest release of Paramics, Version 4.1.

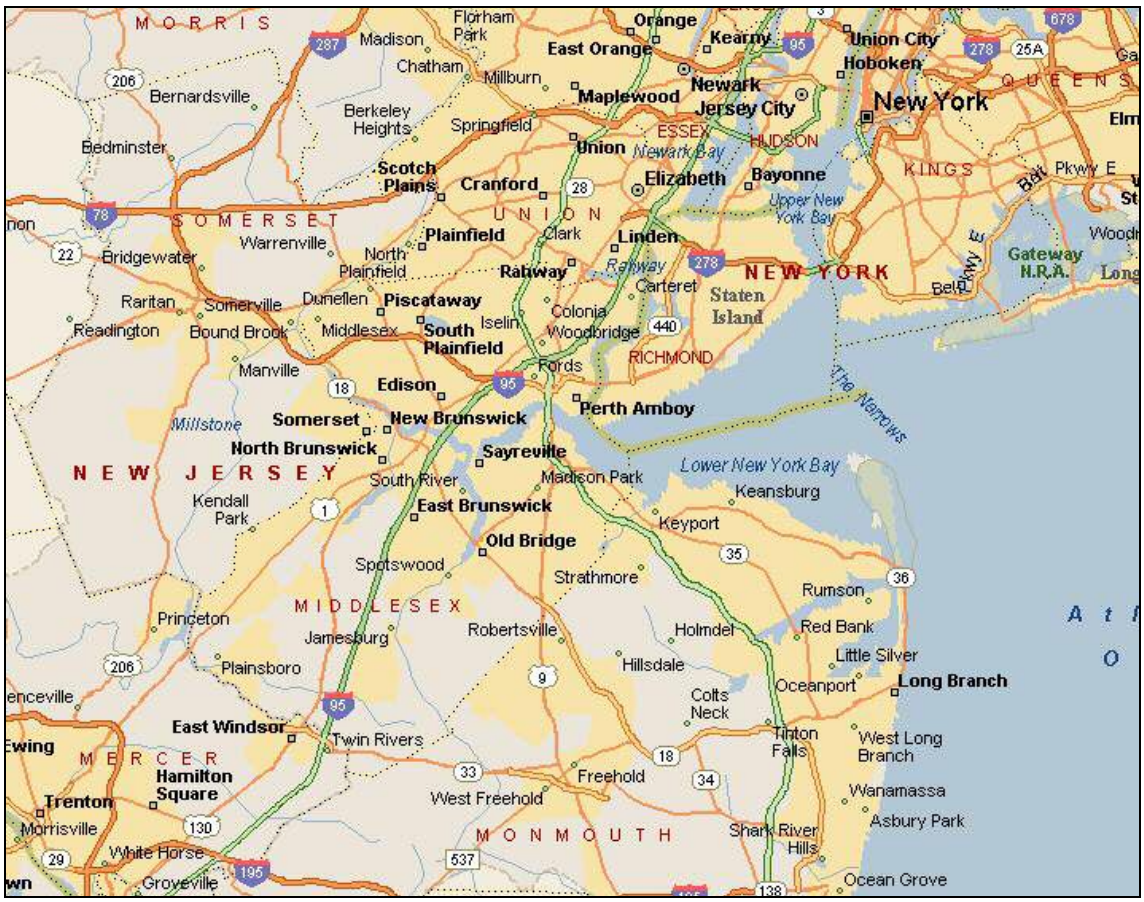

Figure 4.1: Study area

4.3 Use of GIS to generate initial PARAMICS input files

Typically, network coding in Paramics involves the use of overlay files as templates to build the model road geometry. The overlay files are loaded into the Paramics graphical user interface and used as a background layer to manually position the nodes and links in the Paramics network under construction. The files can be aerial photos in Bitmap (bmp) format, drawings in AutoCAD (dxf) of TGA (tga) formats. The overlay function provides valuable support in the network coding process. However, when dealing with large and complex networks, the coding task remains a labor-extensive process. In order to speed up the process, a tool was recently developed to automatically convert GIS files into a Paramics compatible format.

Called S2P (Shape file to Paramics), the conversion program was developed by the University of Santa Barbara under contract for Caltrans. The tool had been formerly tested by the research team, and had been found to perform well. It was therefore decided to use S2P in the process of coding the new I-95 network.

The first step was to acquire the necessary GIS files for the area under investigation. A resource was identified on the Internet, the Environmental Systems Research Institute (ESRI). It is possible to order GIS files online at www.esri.com. The files are available by ZIP code. For this study, 26 ZIP codes were falling within the study boundaries and the corresponding GIS files were ordered. Each GIS file includes 12 types of Arcview shapefiles (*.*shp*) such as *Highway.shp*, *County.shp*, *and so on*. However, for the purpose of this study, only the *Highway.shp* in the each GIS file was used. The *Highway.shp* includes information on highways and major arterials, which is precisely what is needed in coding the road geometry in Paramics.

The process of conversion in S2P includes a decision between two generalization models: Model 1 – Douglas Poiker and Circular Arcs, and Model 2 – Biarcs. The network output produced by S2P is highly sensitive to the choice of the conversion method, and the setting of the various parameters associated with each method. The number of nodes and links in the network output can be quite different, significantly affecting not only the visual aspect of the network geometry but also the traffic performances predicted by the model. A number of tests had been previously carried out as part of an evaluation of the S2P program performed for Caltrans by the research team. These former tests were very useful in the process of selecting appropriate conversion methods and adjusting the conversion parameters in the I-95 project.

It was finally decided to use Method 1 (Polyline Douglas-Poiker) with the default parameters.

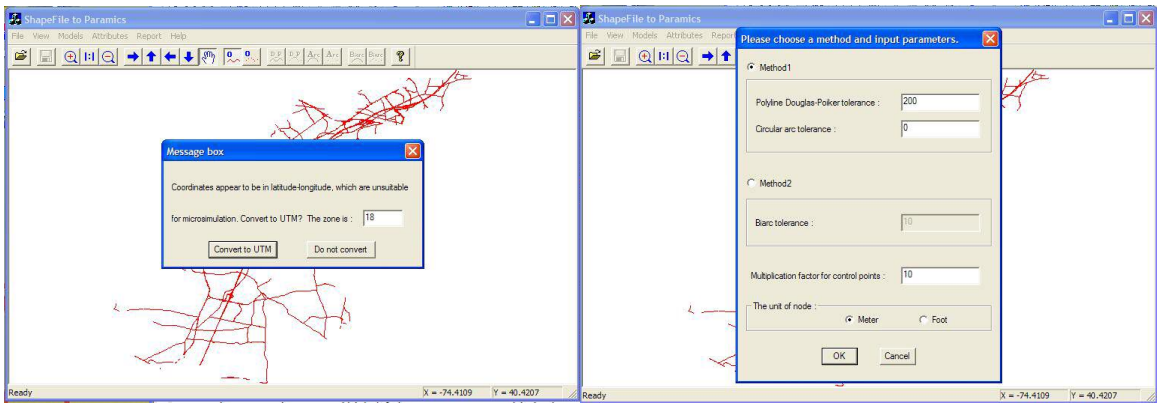

Figure 4.2: S2P conversion software

Because the ZIP files were initially received in three sets, the same conversion process was carried out three times. The three resultingnetworks were finally combined into one network in Paramics by using the *Paste Network* function under *Editor Option* in Modeller (see Figure 4.3).

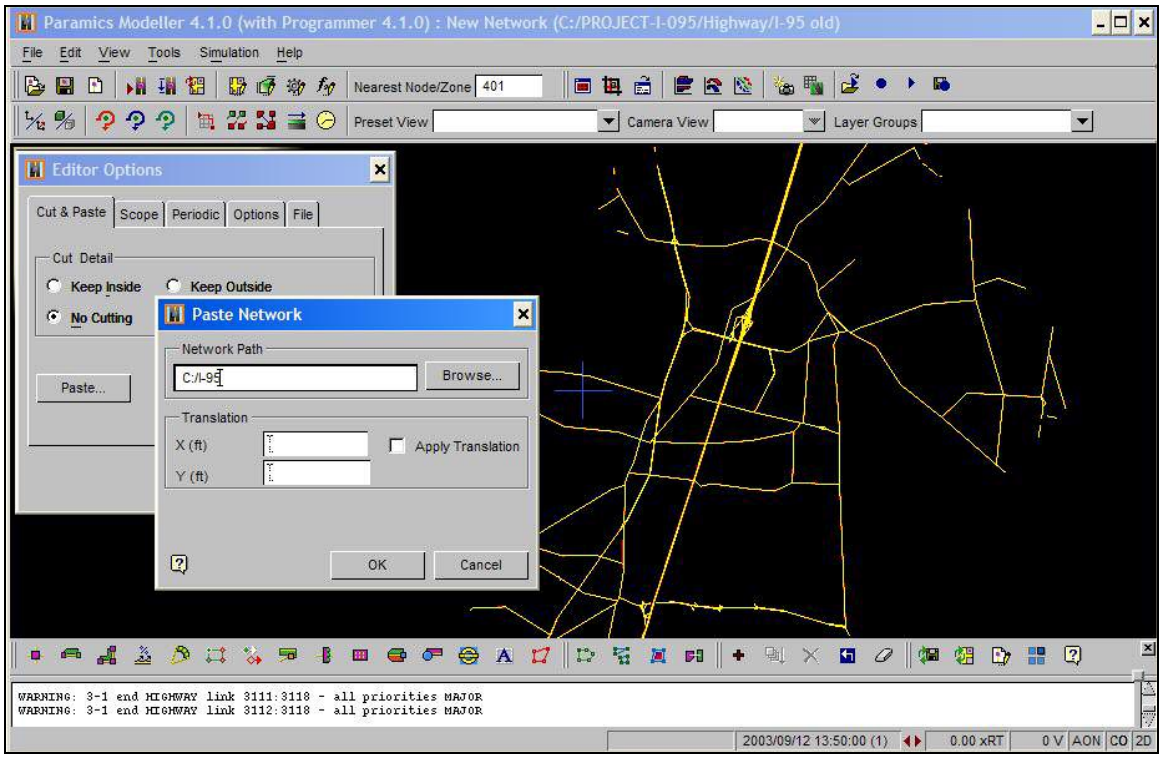

Figure 4.3: Combining the three sub-networks into one

This method of network building provides a fast and cost-effective way of generating a first network structure in Paramics. However, it was recognized that the network generated automatically had to be carefully checked and manually adjusted in order to be

suitable for the purposed of the project. Those necessary network adjustments are presented in the next section.

4.4 Network geometry adjustments

The network created with S2P had a number of deficiencies, including problems with the position of nodes, the network connectivity, and the lack of many link attributes such as speed limits. As a first step, many arterials initially represented in the network generated with S2P were deleted, because they were not directly part of the defined study area. This happened because the original GIS files had many more streets than the ones selected for the model.

The resulting overall network is shown on Figure 4.4. This figure is a screen capture of the graphical user interface of Paramics. It shows the roadway facilities in red.

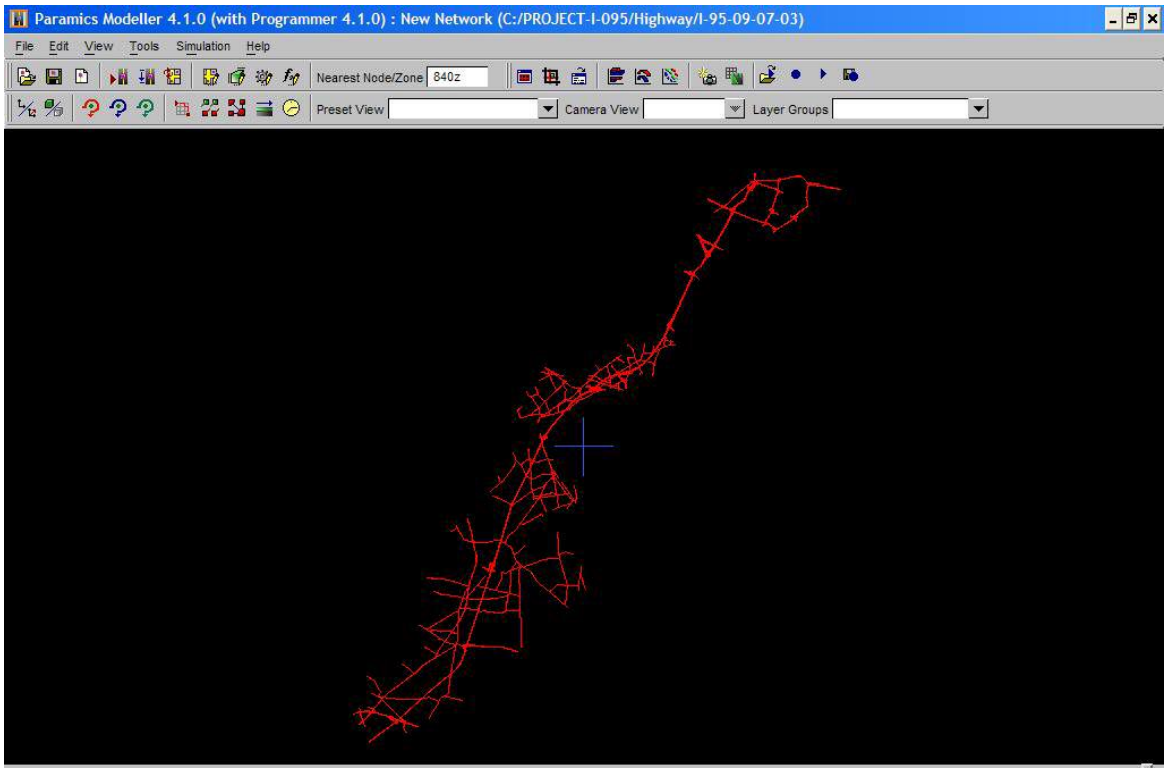

Figure 4.4: Overall Network in Paramics

Among the remaining links, data was required to describe the details of the roadway, such as the number of lanes and posted speed limits. The NJ Turnpike authority provided data on the number of lanes on the mainline freeway and the number of toll lanes at each toll station between exit 7A and the Holland Tunnel (see Table 4.1).

This information was coded into the Paramics model by modifying the link category file. A major section of the New Jersey Turnpike within the study boundary has two roadways for each direction (between Exit 8A and Exit 14). The inner roadway is primarily used by passenger cars whereas trucks should use the outer roadway. The restriction strategy in Paramics Modeller was activated to replicate this situation. Further details on the implementation of traffic restrictions are provided in the next section.

Section	**Number of Mainline Lanes**	
	Northbound	**Southbound**
Exit 7A – South of Exit 8A	3	3
Exit 8A – Exit 9	5	5
Exit 9 –Exit 11	6	6
Exit 11 – Exit 14	7	7
Exit 14 – Exit 14C	2 (Eastbound)	2 (Westbound)

Source: New Jersey Turnpike Authority

Table 4.1: Lane configuration on the NJ Turnpike mainline

Another major issue had to do with curved links: the method and parameters applied in the conversion process did not allow for the creation of curved links. This was done in order to simplify the geometry of the output network, and minimize the number of nodes/links created which enhances the traffic flowing performance of the network. The downside of this choice is that the visual aspect of the network is of lower quality in the absence of curved links. It was decided to manually introduce curvatures on the freeway ramps and connectors. This was done for all interchanges on the I-95 freeway.

Other problems occurred for coding interchanges based on the files produced by the S2P conversion program. The crossover sections between the highway, ramps and arterials were badly coded. It was necessary to manually revisit the coding of each interchange to produce a realistic network representation.

The process involved downloading aerial photos of each interchange, and using these photos as background images to manually adjust the geometry of the Paramics network. The aerial photos formatted as Bitmap (*.bmp) were collected from the Microsoft® TerraServer USA Homepage (Ref. 4). The U.S. Geological Survey (USGS) provides the Microsoft® TerraServer USA site with images and maps of the United States. The

images are available to download and use for free. By using the aerial photos and manually adjusting the coding, the geometry of interchanges and toll stations could be represented with a high level of details.

As an example, the use of a Terraserver aerial photo as an overlay in Paramics is illustrated on Figure 4.5 for a particular interchange along the I-95 freeway. The Paramics network structure appears in red on this screen-capture picture.

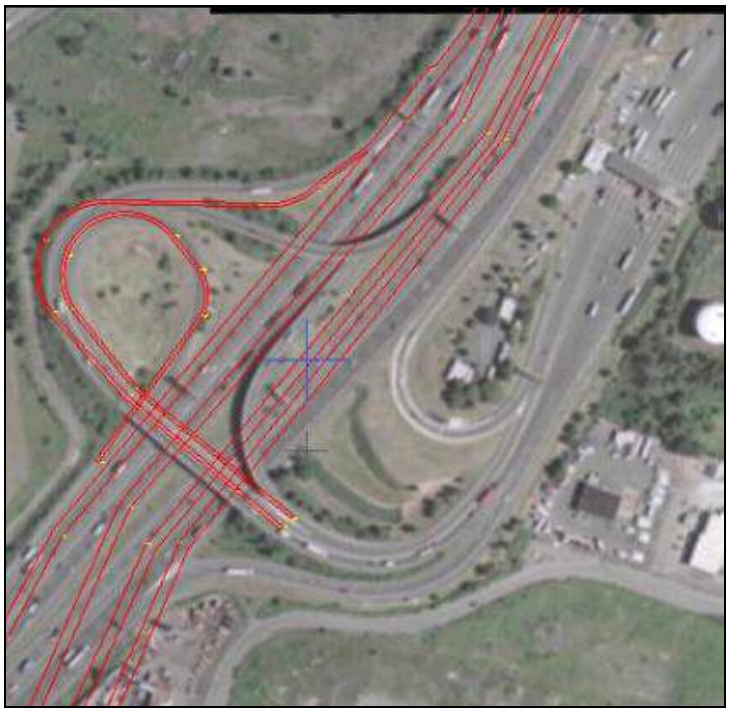

Figure 4.5: Aerial photo used as overlay for interchange coding

The modeled network includes thirteen toll stations along the Turnpike (between Exit 7A and the Holland Tunnel Toll Station). Among these thirteen toll stations, four will be equipped with radiation detection sensors (see Figure 4.6 and Table 4.2) in the simulation exercise. The four toll stations selected for installation of detector stations required special attention in the model coding process, to make sure the geometry and vehicle behavior observed in the field were represented as closely as possible.

In order to increase the realism of the simulation, it was required to differentiate vehicles using Electronic Toll Collection (ETC) systems. Paramics was able to replicate the features of ETC systems by properly coding the toll lane geometry and adjusting lane restrictions. Two or three lanes at each toll station are designated to allow only ETC vehicles: these vehicles are not required to stop for fee collection when they use the special ETC toll lanes.

Restriction strategy in PARAMICS Modeller forces non-ETC vehicles to use the other lanes on the toll station; those non-ETC vehicles do have to stop and pay at a toll booth. In the simulation, non-ETC vehicles stop for 2 seconds at the entrance toll booth and for 3 seconds at the exit toll booth (these values could easily be changed).

More details on the proportion of ETC are provided in section 4.4.

More details on the use of lane restrictions in PARAMICS Modeller are provided in section 4.5.

One example of toll plaza coding is shown on Figure 4.7. The toll booths appear as orange rectangles. The left rectangle is the entrance both while the right rectangle is the exit booth.

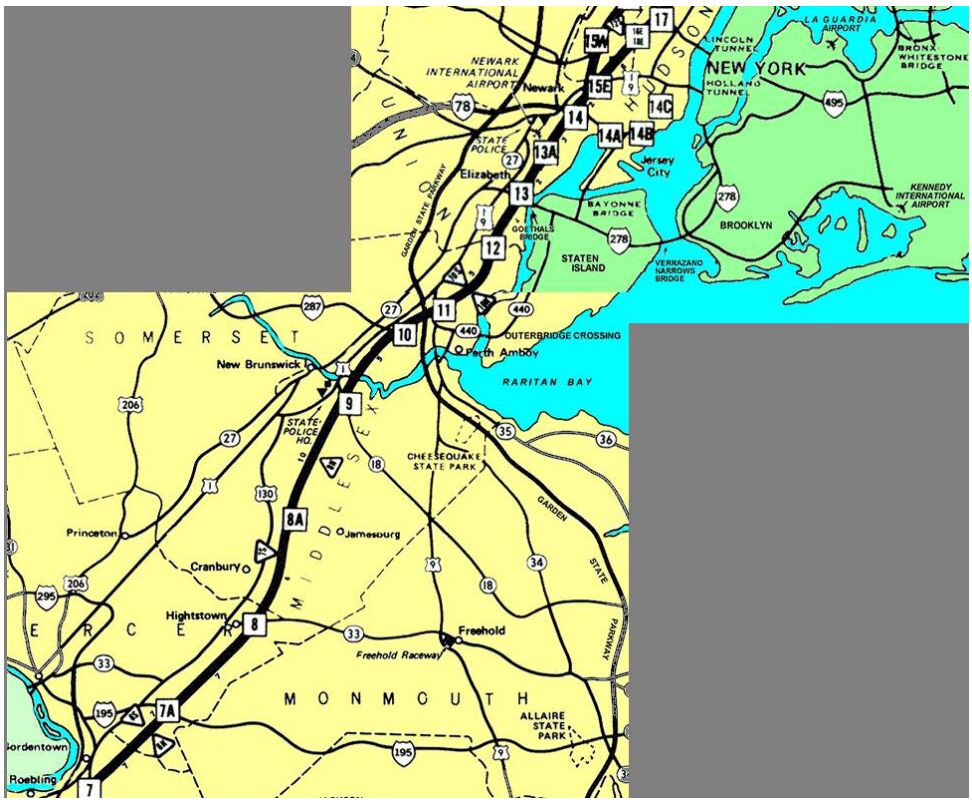

Source: **http://www.state.nj.us/turnpike/nj-vcenter-maps.htm**, Sep, 2003.

Figure 4.6: Location of toll stations in the study area

Interchange Number	Location	Total number of Toll lanes	Milepost
7A*	I-195, Trenton, Hamilton	10	60.0
8	NJ-33, Hightstown, Freehold	5	67.6
8A	Cranbury, Jamesburg	9	73.7
9	NJ-18, New Brunswick	16	83.3
10*	I-287, Metuchen, Perth Amboy	14	88.1
11	Garden State Parkway	26	90.6
12	Carteret, Rahway	7	95.9
13	I-278, Elizabeth, Staten Island	21	99.9
13A	Newark Airport, Eliz. Seaport	21	101.6
14	Newark Airport, I-78, US 1 and 9	27	104.7
14A*	Hudson City, Ext, Bayonne	11	N3.5
14B	Jersey City, Liberty St. Park	5	N5.5
14C*	Holland Tunnel	12	N5.9

Source: NJ Turnpike Authority, Ref.5

Note: *indicates the toll station which has vehicle detection system

Table 4.2: Details on toll stations

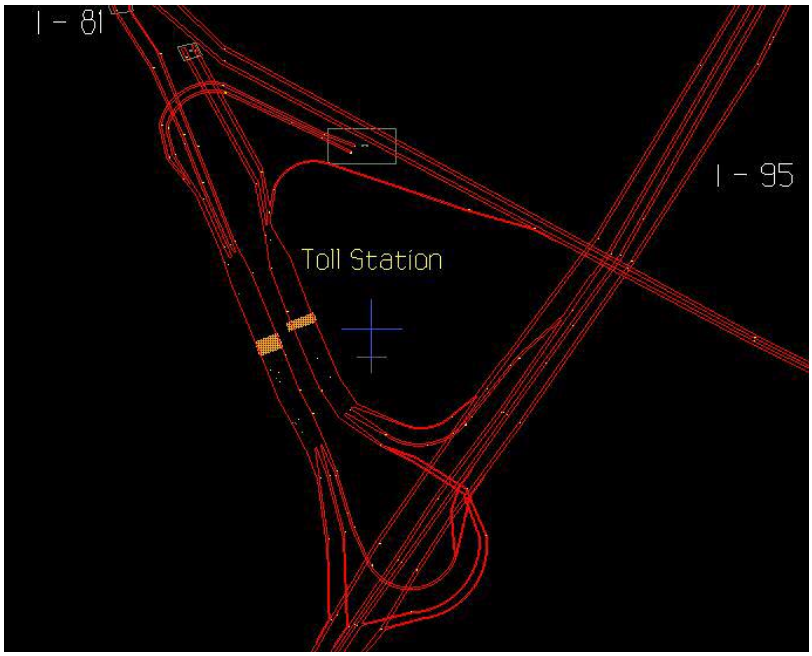

Figure 4.7: Example of toll station coding

4.5 Zone structure and traffic demand

The demand side of the traffic simulation requires the specification of a zone structure which will be used to quantify the traffic demand from zone to zone in the Origin-Destination trip table.

The zone structure consists of the number of zones, zone size, and zone locations. For purposes of this study, where the focus was on analyzing traffic operations on the Turnpike, a simplified zone structure was adopted. Zones were positioned at entrances and exits of the Turnpike in order to generate traffic flow on the freeway.

In this initial study, the traffic impact analysis did not include parallel surface streets or other highways, even if they are sometimes represented in the network due to the method used for initially generating the supply side of the simulation. It is important to note, however, that typical corridor studies in Paramics involve the use of complex zone structures, similar to the ones applied in transportation planning model studies. Transportation planning models, like TRANPLAN or EMME/2, produce zone structures and generate traffic demand based on various socioeconomic input data such as a residential population, number of workers, median house income and so on. This type of zone structure would be recommended to perform a more detailed analysis at the corridor level.

As mentioned above, the zone structure was kept relatively simple since the network mainly considered the main freeway, New Jersey Turnpike.

In Paramics coding, the zone has the role of releasing traffic demands to destinations and attracting traffic demands from origins. The zone specification is based on two rules. First, a zone has to include at least half of a link connected to the zone, otherwise it cannot be recognized as a zone in PARAMICS Modeller. Second, a node located in a zone is defined as a zone connector.

Following these two rules, the I-95 Paramics network was fitted with 28 zones.

After building the zone structure, the traffic demand files could be prepared. The traffic demand file is a table specifying the trip demand from each origin to each destination, in a given time period. Two tables were produced, one for a low traffic scenario (representative of off-peak traffic conditions), and a high traffic scenario (representing peak-period traffic conditions).

Each traffic demand consists of two types of vehicles: passenger cars and trucks. Passenger cars are either ETC vehicles or Non-ETC vehicles, depending on the toll payment method. Trucks are either ordinary goods vehicles or light goods vehicles, with different lengths.

ETC vehicles were distinguished by light blue color and others are shown up in gray color.

The various vehicle characteristics are summarized in Table 4.3.

Vehicle Type	Specification	Color	Size (ft)	Proportion
Passenger Cars	Electronic Toll Collection (ETC) vehicles	Light Blue	13.12	60%
	Non-ETC vehicles	Gray	13.12	20%
Trucks	Light Goods Vehicles	Gray	19.69	10%
	Ordinary Goods Vehicles	Gray	26.25	10%

Table 4.3: Vehicle specifications and characteristics

4.6 Traffic assignment

Traffic assignment is the process of choosing routes to go from an origin to a destination. In Paramics as in real life, different drivers/vehicles can have different ways of making decisions regarding routing. Link costs used to make routing decisions can be a combination of time, distance and toll costs. Link costs can be either fixed, made to vary base on historical perceived patterns, or continuously updated based on real time information as the simulation is running.

Broad assignment techniques available in Paramics fall into three main categories: *all-or-nothing assignment, stochastic assignment, and dynamic feedback assignment*.

All-or-nothing assignment assumes that all drivers traveling between two zones choose the same route (ie. the lowest cost route) and that link costs do not depend on flow levels.

Stochastic assignments methods try to account for variability in travel costs or drivers perception of those costs. These methods assume that the perceived cost of travel on each network link varies randomly, within predefined limits.

Dynamic feedback assignment assumes that drivers who are familiar with the road network will reroute if information on the present state of traffic conditions is fed back to them. This is achieved by taking real time information from the Paramics model and using this data to update the routing calculations.

In the I-95 application, a combination of *All-or-noting Assignment* and *Stochastic assignment* was used. By doing so, most vehicles will travel on the shortest path from their origins and destinations, while a few vehicles will pick other paths.

As mentioned earlier, Paramics offers the option of specifying link restrictions to force a vehicle's routing and lane choice. The definition of restriction affects the routes

calculated for each vehicle type and change the number of routing options available on the network.

Consequently, two kinds of restrictions have been defined to replicate actual traffic situations. Firstly, there are two roadways between Exit 8A to Exit 14 on New Jersey Turnpike for the northbound direction. Passenger cars can use both the inner and the outer roadway, while trucks are only allowed on the outer roadway between Exit 8A to Exit 14. The first restriction has been defined at all the links on the inner roadway and trucks were restricted to use the inner roadway by implementation of this restriction.

Secondly, each toll station has two types of lanes depending on the payment method. The ETC-equipped vehicles are the only ones allowed to use the ETC designated lanes. The Non-ETC vehicles can only used conventional lanes. They are supposed to stop a few seconds to pay toll by hand. This is illustrated on Figure 4.9.

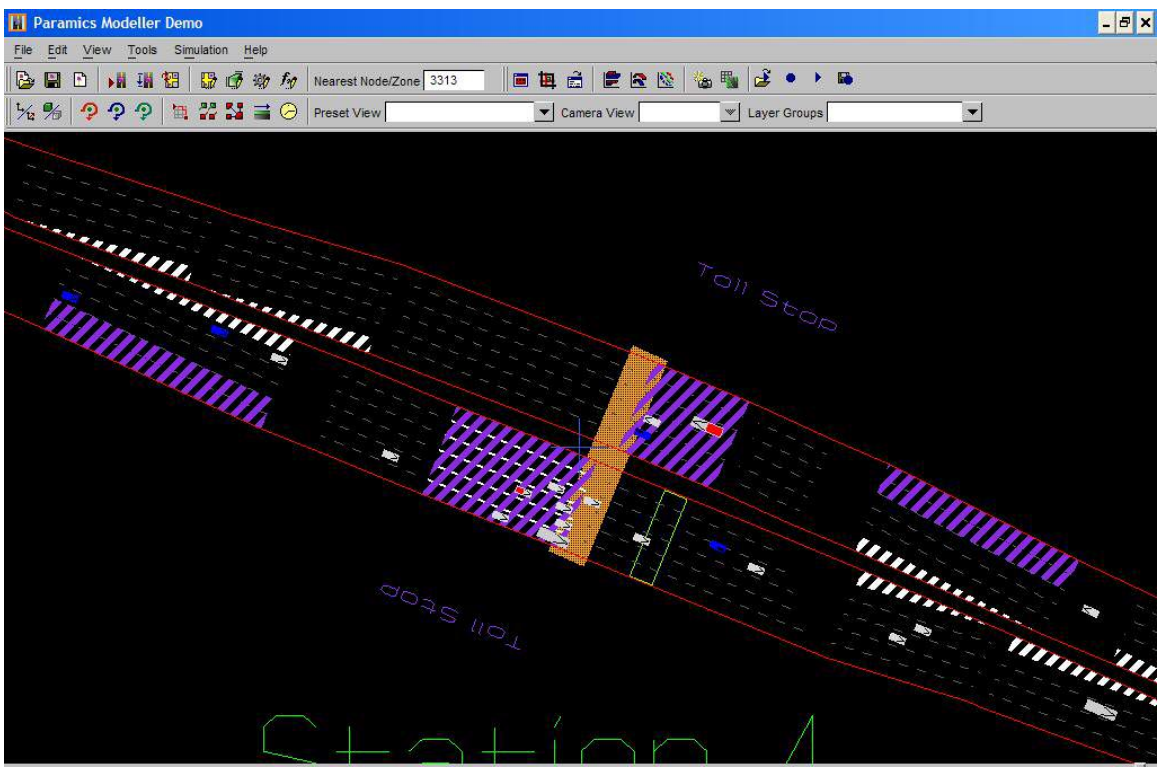

Figure 4.8: Lane restrictions (in purple) at a toll station

4.7 Detector Station Placement

The scenario definition phase identified four locations where detector stations were to be implemented in the model. The detector stations are equipped with special sensors capable of detecting specific properties among traveling vehicles.

For purposes of the simulation exercise, the focus was on vehicles traveling northbound on the NJ Turnpike towards the Holland Tunnel. The vehicles crossing the entire modeled network were to successively encounter four detection stations along the way (see Figure 4.9):

- Station 4: located at the Interchange 7A tollgate entrance (junction with I-195), at the southern edge of the study area
- Station 3: located on the mainline freeway, just north of the overpass at Interchange 10 (junction with I-287)
- Station 2: located at Interchange 14 tollgate exit(towards Route1/9 via the Pulaski Skyway)
- Station1: located at the tollgate entrance to the Holland Tunnel

Figures 4.10 and 4.11 are screen-captures of respectively, Station 3 and 4.

The properties of the detector sections are explained in more details in Section 3 of the report. Each station has a specific role in terms of identifying vehicles, tracking vehicle movements, predicting arrival times at subsequent locations, or triggering other appropriate actions. In the graphical user interface of Paramics Modeller, the detectors are shown as traditional loop detector devices, which are used to collect basic traffic performance data such as vehicle headways, speeds, or flow rates. But the APIs developed as part of the I-95 project considerably extended the capabilities of the conventional loop detectors by adding a number of features directly relevant to this application. This API development work is documented in details in the next section.

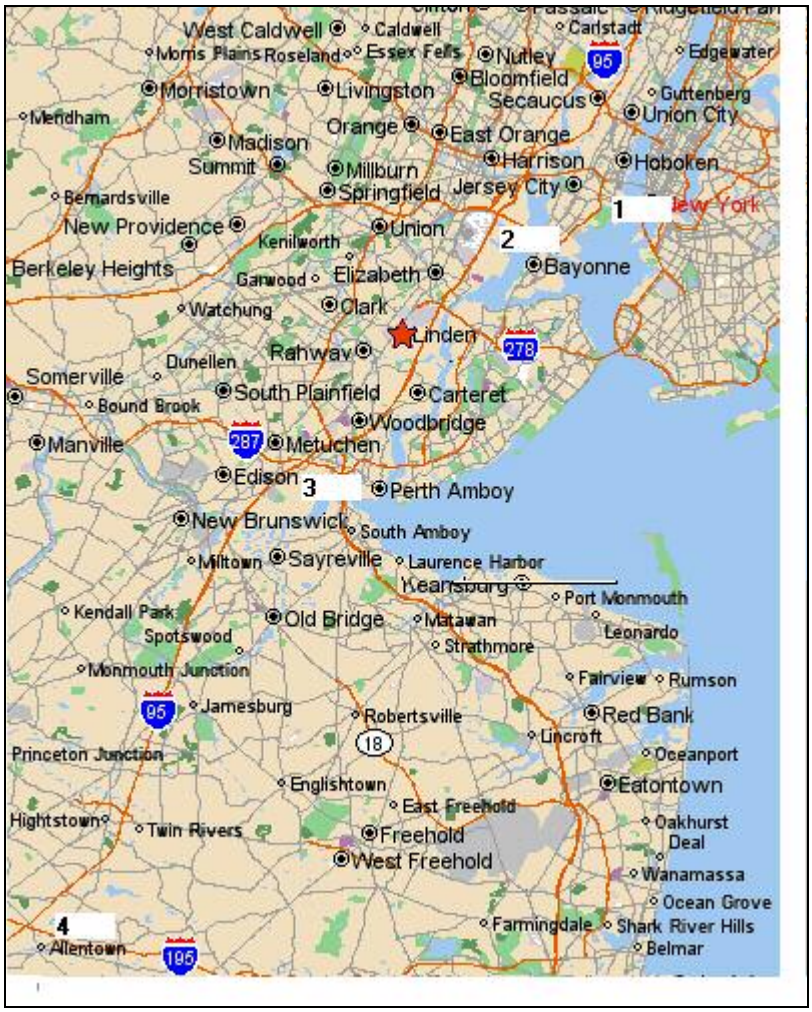

Figure 4.9: Location of four detector stations

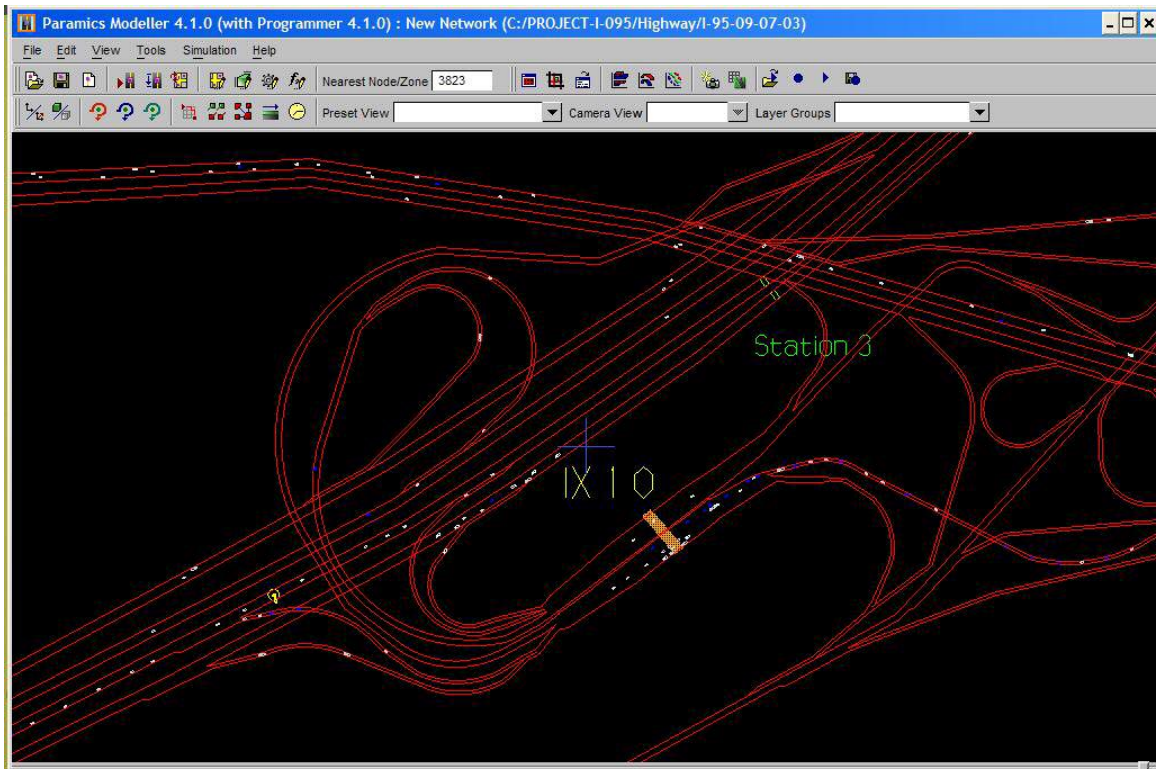

Figure 4.10: Station 3 (at Interchange 10 overpass)

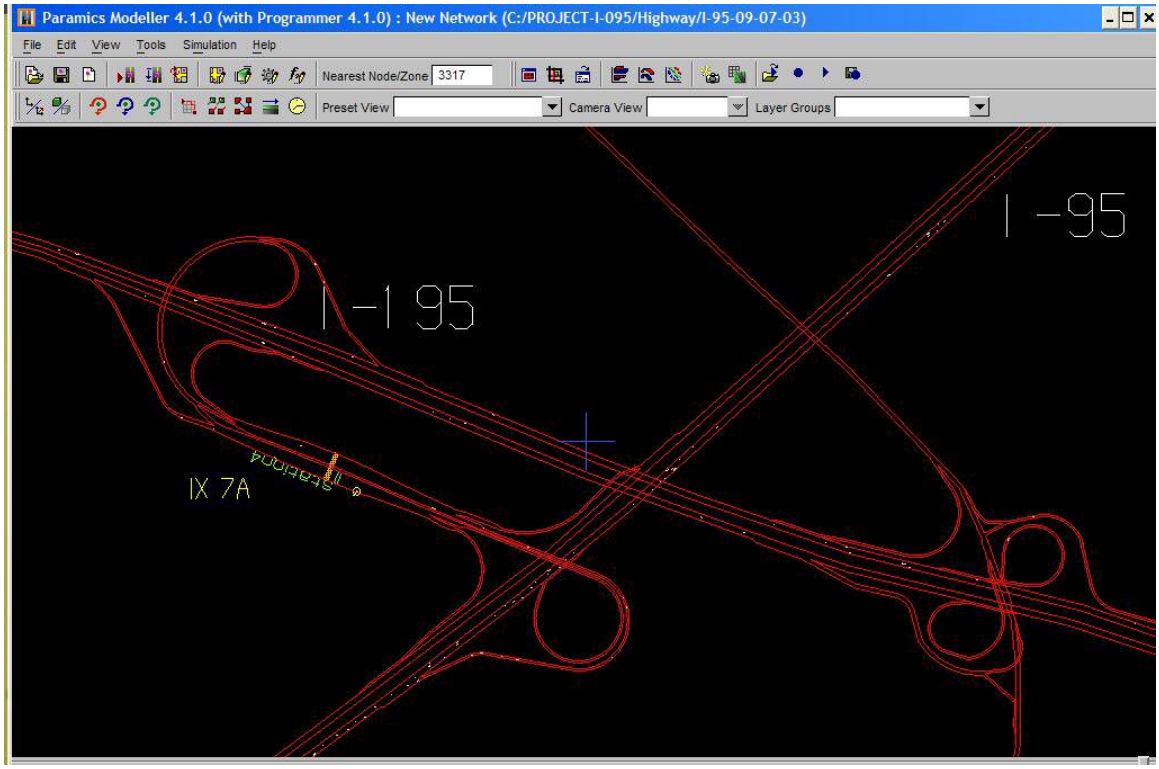

Figure 4.11: Station 4 (at Interchange 7A tollgate entrance)

4.8 Development of New Functions Using API (Application Programming Interface)

In this section, the new functions developed to enhance the Paramics model for purposes of the I-95 project are presented. The chapter is divided into three parts. First, the requirements are defined by identifying the specific desirable features. Then, the capabilities of the standard Paramics model to meet these requirements are examined. Finally, a description of the newly developed functions is provided.

The general goal of the simulation experiment is to build a tool capable of representing the roadway facility, and the specific features of the radiation detection system. With regard to representing the general roadway environment, Paramics Modeller provides an ideal tool, with its high quality graphical capabilities. The network coded as described in Chapter 2 of this report offers a perfect foundation to demonstrate the feasibility of the detection system.

If modeling the supply side did not raise any major problems, the simulation of the detection system was obviously more of a challenge, given that Paramics had not been developed for that particular application. However, the research team could take advantage of the API (Application Programming Interface) that allows the Paramics user to modify logic functions within the core model, or develop new functionalities and plug them into the main model.

It was important, as a starting point, to precisely define what these functionalities should be. In order to do so, a scenario was built, with the aim of replicating it as the simulation runs. The scenario involved a series of actions occurring at various locations along the modeled section of freeway.

More specifically, the scenario can be described as follows:

The traffic flowing on the NJ Turnpike is typical of an average day. Two demand patterns are considered: off-peak and on-peak periods. A vehicle traveling north on the NJ Turnpike enters the freeway at the southernmost entrance. As this vehicle passes detector station 4 (at the toll entrance gate), an alarm is turned on and the status of the vehicle changes to “suspect”. A travel time estimation module is activated to predict the vehicle arrival time at various points along the trip towards the Holland Tunnel. When the vehicle passes by Station 3 (located 28 miles further), the alarm is confirmed, and the travel time estimation is updated. At Station 2, the status changes again, this time from “presumptive” to “confirmed”. Finally, when the vehicle reaches the entrance to the Holland Tunnel, after the alarm has been confirmed one more time at Station 1, the vehicle is intercepted and prevented from entering the tunnel.

In order to simulate this scenario, a number of development tasks needed to be performed. They are classified into five categories:

- Target vehicle releasing
- Vehicle detection

- Vehicle tracking
- Travel time forecasting
- Vehicle interception

Once the requirements were identified, the approach consisted in exploring the features of the standard Paramics program to identify if any of theses requirements could be met without any additional development.

4.9 Exploration of relevant existing functions

The latest release of Paramics, Version 4.1 offers new features that could potentially be used to perform some of the tasks previously highlighted. In particular, the incident modeling tool and the vehicle marking function were found to be worth investigating with regard to their potential use as part of the I-95 project.

The incident modeling function directly available in Paramics Modeller is designed to investigate the impact of certain types of vehicle events on the traffic. For example, if a car is on fire, out of fuel, or any other out-of-orders, it will stop suddenly or pull over gradually, but in either case it may block one lane for a certain period and hold the traffic, creating a bottleneck temporarily.

In Paramics Modeller, this type of event can be simulated. A specific input file, called "incidents" is required. An example of incidents file is shown on Figure 4.12.

```
incident definitions
type 1 "vehicle fire" 0x00ff00ff wait time 00:45:00 in lane 1
passing speed 35 kph opposing speed 30 kph

incident locations
link 96:16 at 08:15:00 at 1050 m type 1
```

Figure 4.12: Example of "incidents" input file

In this example:

- "Type 1" refers to the vehicle type,
- "vehicle fire" indicates the incident type to be used on the tag attached to the incident vehicle.
- The hex number is the internal color for marking that special vehicle.
- "wait time" is the duration of the incident.
- "passing speed" is the speed at which the following vehicles are passing the incident vehicle
- "opposing speed" indicates the speed of traffic on the opposite direction, which is also affected by the incident
- "incident locations" section positions where and when the incident is going to happen in the Paramics network.

In the case of Figure 4.12, the information gathered by the model is that: "a vehicle of type 1 will get on fire at 08:15:00 at lane 1 on the link 96:16 at a distance of 1050 meters from the downstream node of the link. The vehicle on fire will be marked in magenta and the traffic in the same direction will travel at 35 kilometers per hour when passing the site and the opposing traffic will travel at 30 kilometers per hour. The incident will last 45 minutes".

Figure 4.13 shows a screen capture of Paramics under incident conditions. The stopped vehicle appears inside the purple circle, and the text associated "vehicle fire (00:29:43)" indicates the incident type and the remaining duration of the event.

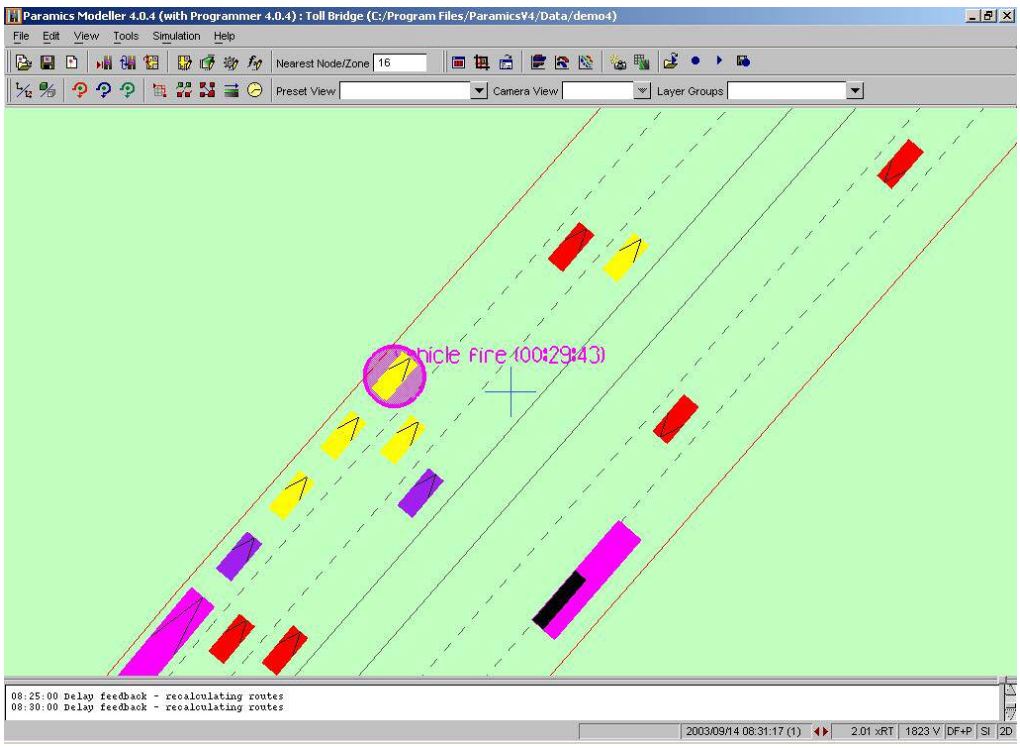

Figure 4.12: Example of incident
(left-hand drive)

The incident modeling tool at first appeared to have several features similar to the ones required in the detection project, including marking and tagging target vehicles, timing of incident/suspect vehicle, and controlling vehicle speed.

However, after further analysis and experimentation, a number of limitations were identified with regard to how much could be done using the standard incident modeling tool provided by Paramics. For instance, there is no direct way to generate a series of incidents occurring to the same vehicle. The process of selecting vehicles for incident occurrence is a random process, which cannot be controlled. In our scenario, on the

contrary, it is assumed that the same vehicle will successively encounter the various stations and see its status updated accordingly.

Another limitation has to do with the type of message that can be associated with a vehicle under incident conditions, limited to a single line. The prescribed scenario calls for much more details regarding vehicle characteristics, and projected arrival times at subsequent locations, which could not fit in one line.

Paramics Modeller also provides a feature called “Marking Rules” which was designed to associate different colors or annotations to vehicles meeting specific conditions defined by the user at the beginning of the simulation. Multiple conditions can be specified.

As an example, Figure 4.14 shows a screen capture of the Marking Rules Manager. It can be seen that marking rules may be specified based on different criteria including:

- Zones: specific trips from certain origin(s) to certain destination(s) can be marked.

- Vehicle types: some or all vehicle types can be selected for marking. Also the familiarity level (familiar vs. unfamiliar) can be used to make further distinction.

- Links: define the location(s) where the marking rules will first apply, if the vehicles traveling on those links meet the other conditions.

These different settings provide a combination of options available for identifying certain vehicles at specific locations, and assign user specified colors and tagging annotations. However, there are a number of limitations when considering the specific requirements listed previously. For instance, no timing criteria is available to specify when the marking rules should begin to apply; a marking rule applies to a group of vehicles and no individual vehicle can be single out for applying specific markings; finally, the annotation size is limited and would not allow for a complete travel performance report to be presented.

Figure 4.13: Vehicle Marking Rules Manager

Incident modeling and vehicle marking functions provide a good starting point to evaluate the capabilities of the existing software, and better assess the directions for future development.

Vehicle marking illustrates that vehicle characteristics can be changed while the simulation is running, when a set of triggering conditions are met. These characteristics include coloring, tagging, highlighting, and blinking. They provide a full set of visual effects to draw on. The marking manager also shows how to apply the rule only after the vehicles passes at specific location, a feature highly relevant in our simulation scenario.

Incident modeling shows the possibility to generate and control the occurrence of specific events to random vehicles, and the slow-down or stopping of these vehicles.

If these existing functions provide a good starting point to work with, it was recognized, however, that they did not have the flexibility to allow for the full range of requirements to be met. The replication of the scenario identified earlier required the development of specific features, which was performed with the Application Programming Interface available in Paramics Programmer. This development effort is being described in detail in the next section of the report.

4.10 New function development using API

Paramics Programmer provides access to the Application Programming Interface (API) component of the Paramics suite. Through API, users can access functions at the core of the simulator, providing a way to modify or replace some of the inner logic within the model. It also provides access to basic information produced by the model as the simulation runs, so that the user can plug-in its own subcomponent and interface it with the rest of the simulation platform.

In the latest release of Paramics Programmer (Ref. 6), the API functions were significantly improved. The upgrades included more than 700 new API functions, an improved structure and format of the functions, and extended core model access. Programmer Version 4 also now supports Java development, in addition to the traditional C language support.

The API plug-ins work as dynamic link library (DLL) file format. When Paramics simulation is running, it will look in the fixed location (ParamicsHome/plugins/windows/) for dll files. If there is one, it will load it into the simulation core engine and the functions incorporated in the dll file will take effect, either by extracting the desired information from the simulator, or overriding parts or all of the traffic modeling logic embedded in Paramics. This type of API provides high flexibility for user defined operations on the basic simulation platform.

In the I-95 project, the API plug-ins were written in C language. The API development effort was organized around the five tasks previously highlighted:

- Target vehicle releasing
- Vehicle detection
- Vehicle tracking
- Travel time forecasting
- Vehicle interception

Each of these five subtasks will be successively examined in the following sections. The entire source code is also available as an appendix to the present report.

This task aims at producing special target vehicles within the simulation. The target vehicle will initially start its trip as any other typical vehicle; its status of target vehicle will appear later while the simulation is running and when specific conditions will be met.
From the modeling standpoint, vehicles loaded onto the network are a series of objects with different attributes, such as height, type, familiar or unfamiliar, etc.

In this application, all vehicles also need to carry additional information with regard to their potential status as target vehicles (such as whether the vehicle is a special vehicle, whether the special vehicle has passed a specific detection station, etc…).

With the Paramics API, this is accomplished by attaching a “user defined data” structure to all released vehicles. The data structure, written in C language is shown below:

```
struct VDATA_s
{
                                                                    Bool
                                                                    special;

                                                                      //whether it
is a special vehicle
                                                                    schar
message[128];

                                                                      // tag
information
                                                                    Bool firstPass;



                                                                      // whether it
passes the first station
                                                                    Bool
                                                                    secondPass;


                                                                      // whether it
passes the second station
                                                                    Bool
                                                                    thirdPass;


                                                                      // whether it
passes the third station
                                                                    Bool
                                                                    fourthPass;


                                                                      // whether it
passes the fourth station
};
typedef struct VDATA_s    VDATA;
```

When certain conditions are met, a specific vehicle will be selected internally by the model to become a potential target vehicle. A number of subsequent actions will then be applied to that particular vehicle.

The target vehicle selection process can be based on two different methods: probabilistic, or fixed time. In the probabilistic approach, a predefined proportion of the vehicles will become special targets (for instance 1% of the overall traffic). Alternatively, in the fixed-time approach, the target vehicle generation will occur at a user specified time in the simulation clock. These two methods reflect options available in the incident modeling tool previously described in paragraph 3.2.1.

Both methods have been successively implemented. It was found that the fixed-time approach provides more control to the user in terms of determining where and when the special target vehicle will be generated. Therefore, it was decided to apply this fixed-time technique when generating target vehicles in all subsequent investigations.

The user specified clock time at which the change of status will occur is defined in a file called "API_UCdetect", which as the following structure:

```
tool "Special Vehicle interception"
API coefficients 1

3.0 "Showing Time ( XX minutes later)"
```

The first line of this file identifies the function that will appear in the "Tool" menu of the Paramics Modeller user interface. The second line indicates how many parameters are in the parameter file. The third line specifies that the special vehicle will be released onto the network from its origin zone after 3 minutes of simulation run.

The Paramics model will call this new file after a line is added in the simulation "configuration" file: the required line is "read parameters file 'API_UCdetect' ".

When selecting the special target vehicle, in view of the scenario previously described, it is necessary to ensure that this vehicle is heading towards the Holland Tunnel, and therefore is bound to drive by the four detection stations located along the path. Special vehicles will only be selected among the vehicles going to zone 2 (Holland Tunnel).

In the scenario to be simulated, only one target vehicle is to be generated. This was the logical first step to be performed in order to illustrate the new simulation features. However, there is no reason why additional target vehicles could not be generated and controlled in a similar manner. It would require only minor refinements to implement this function, for instance using the probabilistic target vehicle selection process previously introduced. That could be performed as a logical continuation of the current simulation effort.

The vehicle detection subtask deals with the series of events triggered when a target vehicle crosses a detection station. Target vehicles are identified, marked and traced each time they pass a station.

As first introduced in Section 4.5, detection stations are modeled using the loop detector function available in the standard Paramics. The loop detector in Paramics, as in real-life, provides a way to collect information as each individual vehicle crosses the device. Traditionally, it is used to compute such traffic statistics as volumes, densities or speeds. The same loop detectors, however, also provide a way to access specific API functions related to vehicle characteristics, the link (road section) where the loop is located, and the status of the loop detectors. In this sense, loop detectors provided the most obvious method of replicating radiation detector stations for this project.

Visually, in the Paramics Modeller graphical interface, loop detectors appear as two parallel lines crossing the roadway, as illustrated on Figure 4.14.

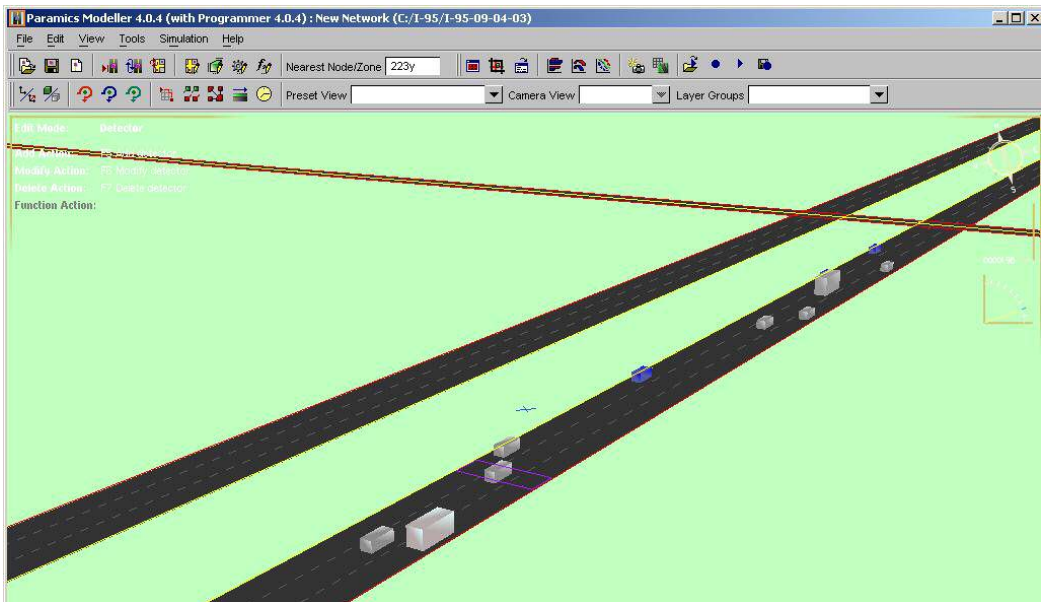

Figure 4.14: Loop detector (at bottom, in purple)

Four detection stations were to be implemented along the modeled section of I-95 (see Figure 4.9). Station 3 has the characteristics of being installed on the mainline freeway at a location where the northbound section of I-95 is actually divided into two parallel roadways. As a consequence, two loop detectors were required to cover both facilities. For the location of Station 2, two options were initially considered, leading to the installation of two loop detectors.

In the plug-in, the four detection stations are therefore named as “station4”, “station31” and “station32”, “station21” and “station22”, and “station1”, respectively.

The API function qpx_VHC_detector(VEHICLE* vehicle, LINK* link, DETECTOR* detector) provides the interaction between vehicle, road and detectors. The detection logic is as follows:

qpx_VHC_detector(VEHICLE* vehicle, LINK* link, DETECTOR* detector)

```
        {
                                                                    if (vehicle is
special)
                                                                    {

                                                                    if (detector is
station XX)

                                                                    {


corresponding actions;

                                                                    }
                                                                    }
        }
```

“corresponding actions” refers to the process of marking and tracking the vehicles that have been detected as target vehicles.

Vehicle tracking is the process of following the target vehicle once it has been identified, and update its status at different locations as it moves along its desired path. This task is the logical extension of the vehicle detection function. When passing different stations, the detection station will determine the status of that vehicle and update the information given to the user.

In the control logic, the “corresponding actions” introduced in section 4.3 is now expanded to the following:

```
if (detector is station 4)
                                                                    {
                                                                      user defined
data firstPass is set “true”;
                                                                      vehicle color
changed to yellow;
```

```
                                                                    vehicle is
marked “SUSPECT”;
                                                                }
if (detector is station 3)
                                                                {
                                                                    user defined
data secondPass is set “true”;
                                                                    vehicle color
changed to orange;
                                                                    vehicle is
marked “PRESUMPTIVE”;
                                                                }
if (detector is station 2)
                                                                {
                                                                    user defined
data thirdPass is set “true”;
                                                                    vehicle color
changed to red;
                                                                    vehicle is
marked “CONFIRMED”;
                                                                }
if (detector is station 1)
                                                                {
                                                                    user defined
data secondfirstPass is set “true”;
                                                                    vehicle
begins blinking;
                                                                    vehicle is
marked “Intercepted”;                                           }
```

The color changes and various annotations associated with the target vehicles through these APIs are the only characteristics that would differentiate those vehicles form any other ones. In other words, the traffic performances of the special vehicles, as well as all interactions with the rest of the traffic are not affected by the new functions.

In addition to the visual marking features, it was desirable to introduce a travel time prediction component. The idea was to be able to predict when a target vehicle will reach the next detection stations along its trip, so as to be able to take appropriate actions in a timely manner. The travel time prediction would take place as the target vehicle passes each detection station, providing a way to validate and update the predictions as the vehicle progresses towards its destination.

Looking at the literature on this topic, it appears that there are two main types of prediction models, reactive and proactive (Ref. 7). Reactive models look backward at the information accumulated before on the vehicle trip, and then try to predict the travel performance in the future. Proactive models are more complex in the sense that they consider individual vehicle route choice behavior each vehicle, and apply an "equilibrium assignment" technique to derive future travel information (Ref. 8).

As part of the I-95, a simplified version of the reactive prediction method was applied. The characteristics of the network are such that no complex route choices are required, and a straightforward "All-or-Nothing" assignment is totally appropriate. This assumption greatly eases the task of predicting travel times, as no dynamic route choices are possible. Vehicles stick to their obvious least-cost route.

In Paramics Modeler, there is a "Vehicle Route" function that can compute travel costs directly. This travel cost can be any combination of time, distance and toll. When the link cost is set to be distance only, the "Vehicle Route" function will result in returning distances between the specified points. This feature was used to calculate the exact distances (as in the modeled network) between the various detection stations. This information is obviously a critical one when predicting travel times from one station to the next. Table 4.3 shows the resulting distances.

Section	**Distance (miles)**
Station 4 to Station 3	28.33
Station 3 to Station 2	17.33
Station 2 to Station 1	9.31

Table 4.3: Effective distances between stations

These distances were then used as a fixed input in the API function developed for travel time prediction.

The simplified approach taken for predicting arrival times assumes that travel speeds falls within the range of 50 to 70 miles per hour. This is obviously a simplification of real-life, but it was thought to be a reasonable assumption to make as part of this initial simulation effort. It would be highly desirable to collect historical information about actual traffic conditions on the New Jersey Turnpike, so as to be able to apply a speed range based on actual measurements.

Considering the 50 to 70 mph speed range, with perfect knowledge of route choice, travel distances and with access to the simulation clock, it becomes straightforward to predict the range of predicted arrival time at each station. As an example, when a target vehicle passes Station 4, the travel time range to reach Station 3 is simply given by: [(Distance 4-3) /70, (Distance 4-3) /50].

Once this information is computed, or updated, it has to be communicated instantaneously to the user, as the simulation runs. The "Reporter" window is used for that purpose. The API function calls for the travel time estimates to be available in the "Reporter" window of Paramics Modeller. Screen shots showing the travel time estimates will be presented in the next section of the report.

Travel time reports are generated at each station. The travel time prediction applies to all remaining stations along the trip. For example, when the target vehicle passes Station 3, the travel time report will predict arrival times at Station 2 and Station 1, based on actual current simulation time.

Vehicle interception is the last phase of the scenario. It deals with the process of stopping the vehicle that has previously been confirmed as a target vehicle, and making it pull over to the side of the roadway. These actions are set to occur just past the last station (Station 1), before the target vehicle enters the Holand Tunnel.

Paramics Programmer provides the ability to control the vehicle speed as well as its location on a certain link (road section). The function qps_VHC_speed(vehicle, float value) is used to set the vehicle speed.

The logic that was implemented to stop the vehicle is the following:

```
                                                    At each
simulation time step
{
 set the speed of the target vehicle to zero;
 set the distance of the target vehicle from the end of the
```

```
    link to a fixed value;
     …
}
```

To provide full control on the behavior of the target vehicle, a mechanism to force the vehicle change lane and pull over to the side is also required.

Two approaches were tested in an effort to implement the pulling-over function.

The first method was to override the lane-changing model, according to the following logic:

```
Lane changing Model
{
if (the vehicle is target and it is passing station 1)
if the vehicle is on the shoulder lane
stop it;
else
move the vehicle to one lane outer than the current one;
repeat until it is on the shoulder lane.
Else
resort to the original model
}
```

The second method that was tested consisted in overriding the route choice model, forcing the vehicle to move to a dummy auxiliary road specially added immediately downstream of Station 1. This second logic is as follows:

```
Route choice model
```

```
{
  if (the vehicle is target and it is passing station 1)
direct the vehicle onto the auxiliary route;
  else
resort to the original model
}
```

Unfortunately, within the limited time frame available, it was not possible to successfully implement any of these two alternatives. The vehicle-stopping feature is available; the vehicle pulling-over remains to be developed.

4.9 RADS Scenario Implementation

Once the network facility was coded in the simulator and all the necessary additional functions had been developed, the next step was to run the simulation with the enhanced functions and check the ability of the tool to reproduce the desired scenarios.

The experiment phase of the project focused on three critical areas:

- The ability of the model to replicate typical traffic conditions on the NJ Turnpike;
- The performance of the added marking and tracking functions;
- The reliability of the travel time prediction tool.

These three areas will be successively examined in this section of the report.

The network geometry resulting from the coding process described in Chapter 2 provides a very good visual representation of the NJ Turnpike section under investigations. Both directions of I-95, as well as all interchanges are coded in fine level of details. The toll stations located at all entrances and exits were given special attention, so that vehicle behavior would be as realistic as possible. As a result of this thorough network coding effort, the simulator provides a high quality graphical representation of the study area.

In addition to the supply side, the other main component of the simulation is the demand data. Because of time and budget constraints, it was not possible to gather or collect all the necessary traffic data required to produce a realistic origin-destination matrix. The quality of the OD matrix is a key requirement for a traffic model to be calibrated. However, the purpose of the project was not to develop a calibrated model for the NJ Turnpike, but instead, to build a network capable of illustrating the various components of the new detection system. For that purposes, a hypothetical origin destination demand table was sufficient, ensuring that the traffic performances of the NJ Turnpike was somewhat realistic.

For demonstration purposes, it was decided to develop two demand scenarios, one representing off-peak conditions, and the other representing peak-period conditions. The two corresponding demand tables were built manually, using engineering knowledge but no actual data. They are obviously a simplification of actual traffic patterns, but were thought to be sufficient at this stage of the project.

Among the features required for illustrating the functions of the detection systems, marking and tracking of vehicles is of critical importance. The simulation scenario defined in section 3.1 calls for specific requirements in that area.

With the help of the API, Paramics was capable of meeting the requirements that were identified. At each station along the route, target vehicles are assigned specific marking codes depending on their status. This provides a useful tool to follow those vehicles as they progress along their route towards their destination. In addition to providing a

visualization tool, the model can be used in a control center to organize the response to the emergency situation.

The series of screen captures presented below (Figures 4.1 to 4.4) illustrates the marking events occurring to a target vehicle as it passes through the various stations,

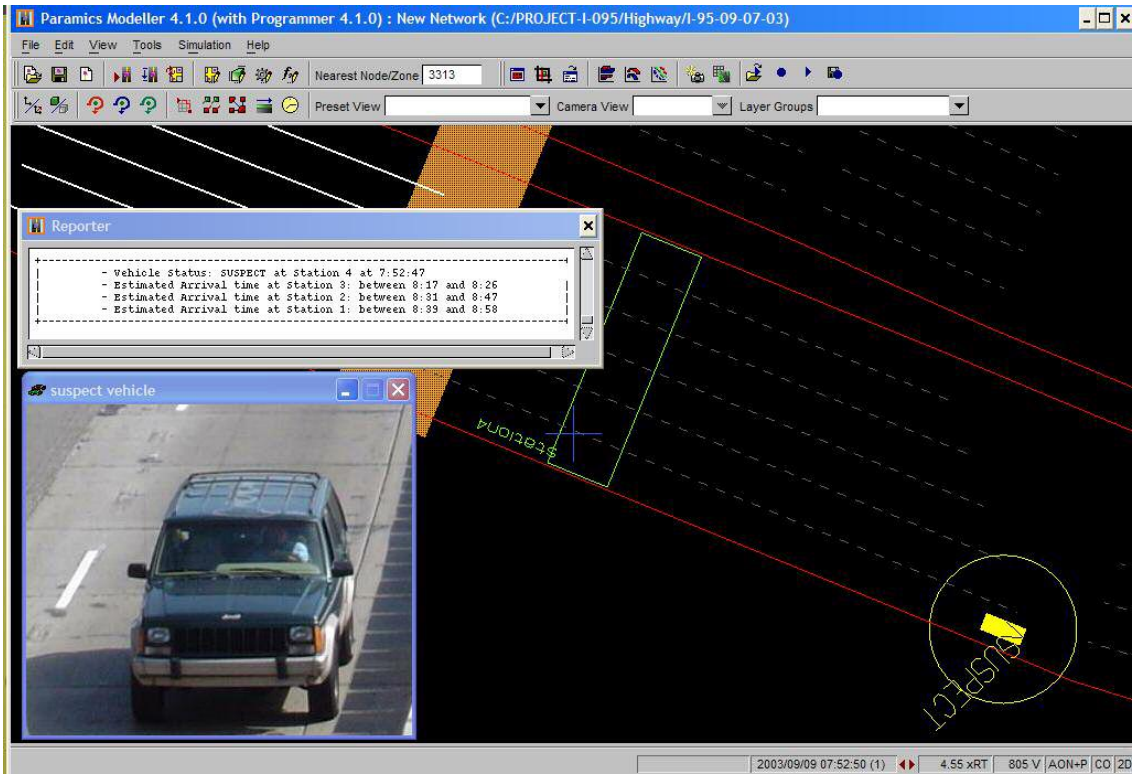

Figure 4.15: Target vehicle at Station 4

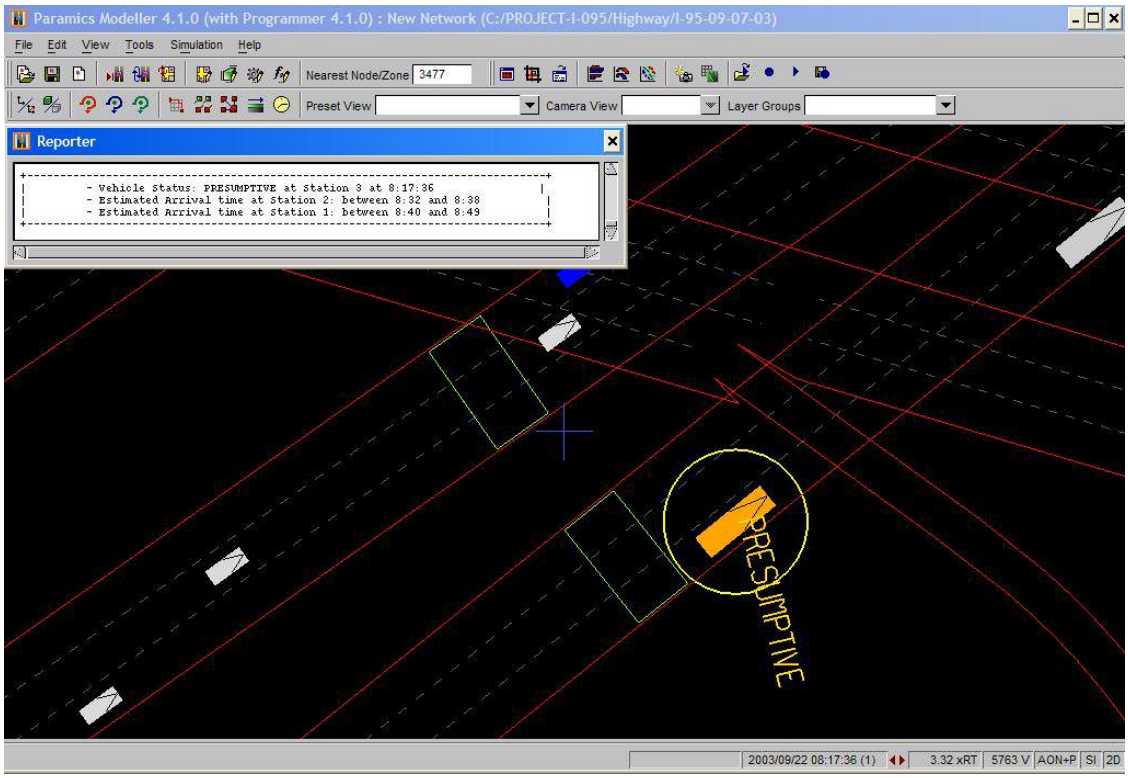

Figure 4.16: Target Vehicle at Station 3

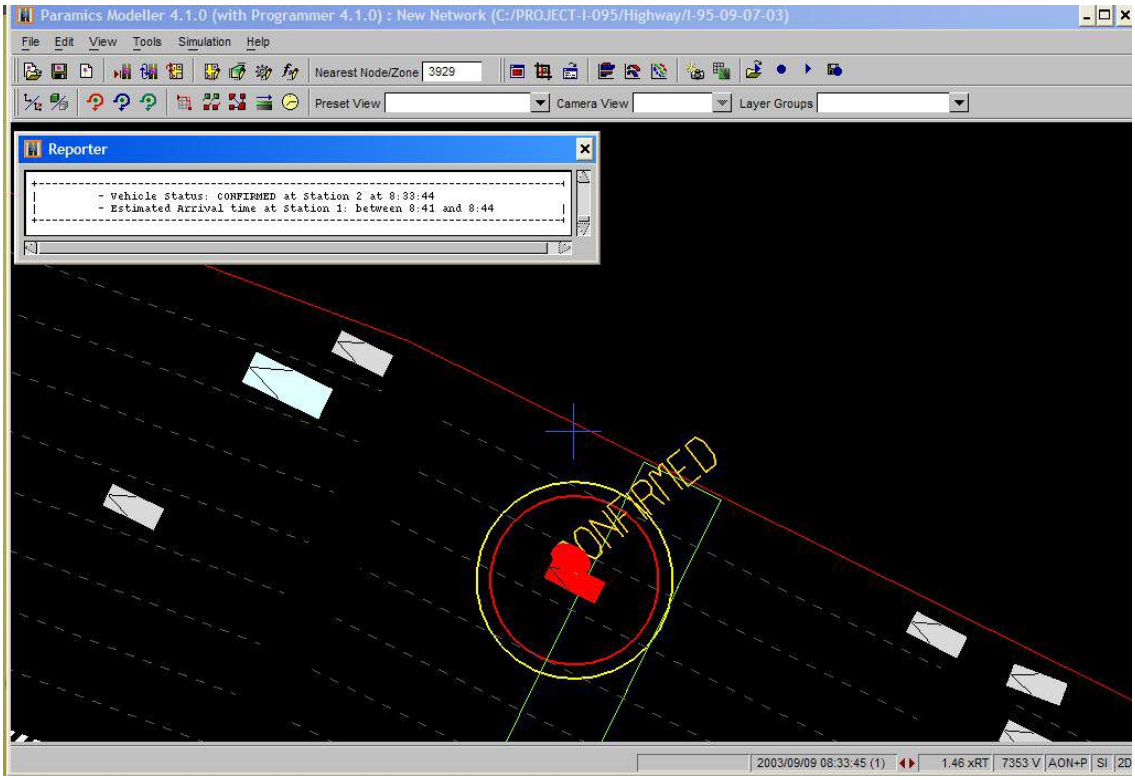

Figure 4.17: Target Vehicle at Station 2

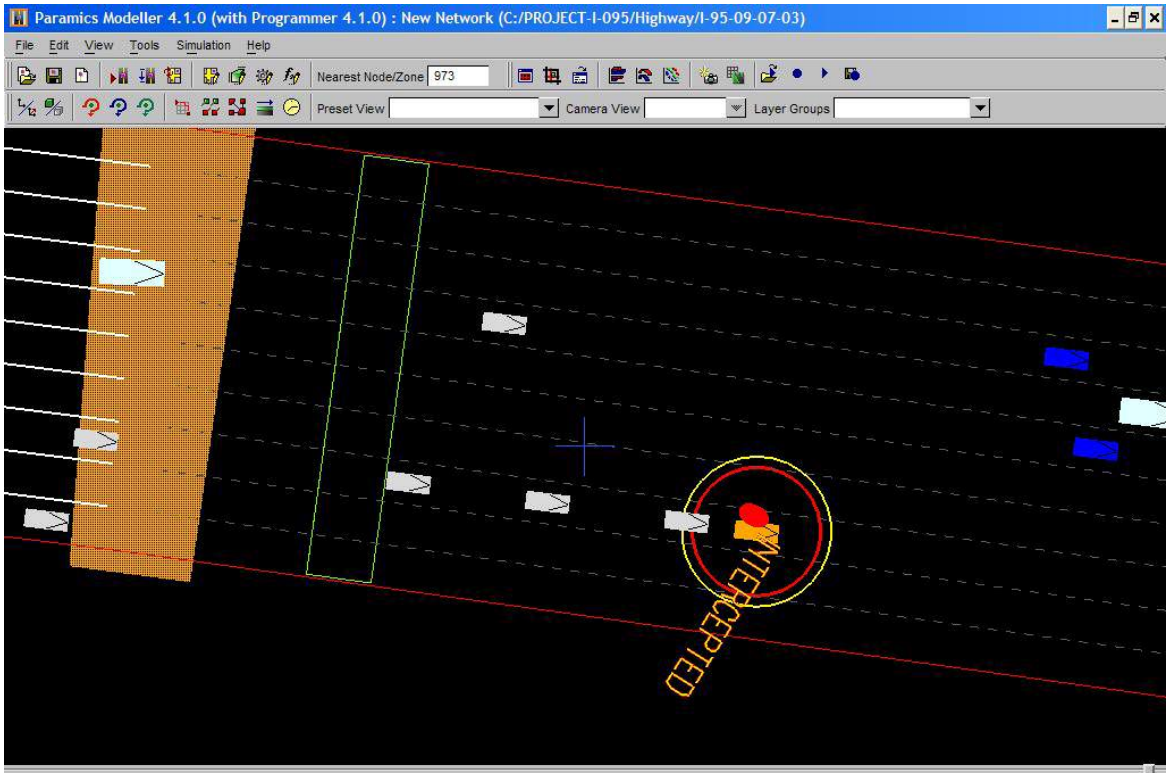

Figure 4.18: Target vehicle at Station 1

4.10 Travel time predictions

The last component of the enhanced-simulation tool to be validated was the travel time prediction module. This was done for the two demand patterns previously introduced, off-peak and on-peak traffic conditions.

The goal was to compare the predicted arrival times (based on the methodology described in sectin 4.3) and the actual arrival times as they occurred in the simulation.

The estimated arrival time window is updated each time the target vehicle passes a detection station. Therefore, arrival time at Station 1 (Holland Tunnel) will first be predicted when the target vehicle is detected at Station 4, then will be updated twice at Station 3 and Station 2; this series of predictions can finally be compared with the actual arrival time at Station 1, as it occurred in the simulation.

Figures 4.19 and 4.20 present, respectively for the low demand and high demand traffic patterns, the succession of travel time predictions made at each Station. At the top of the figure, the first window shows the reporter generated at Station 4. Going down the figure, the other windows show the reports produced at Station 3, then 2, then 1.

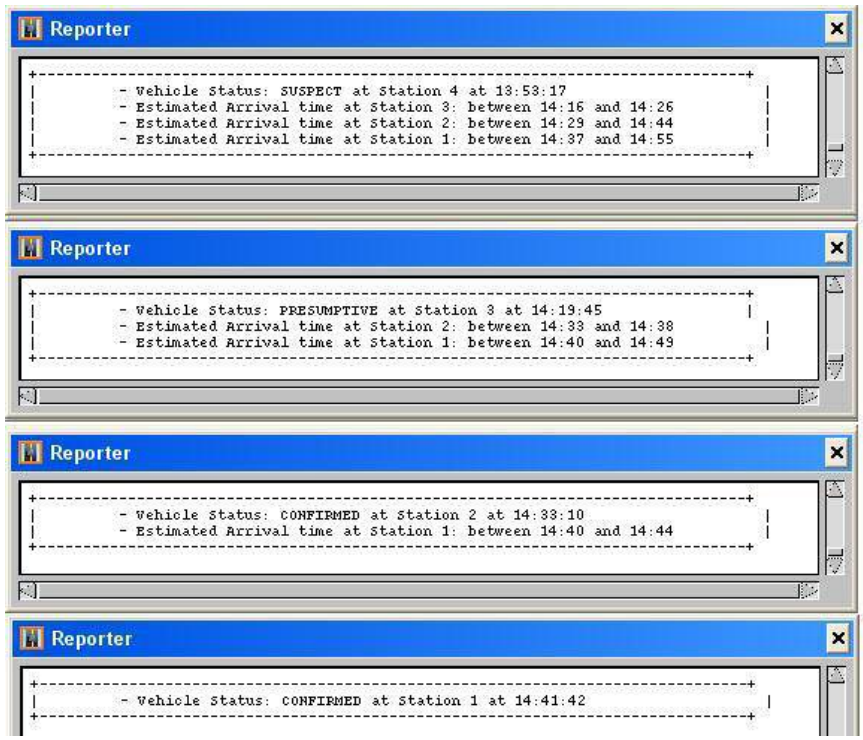

Figures 4.19: Arrival travel time estimates – Low demand case

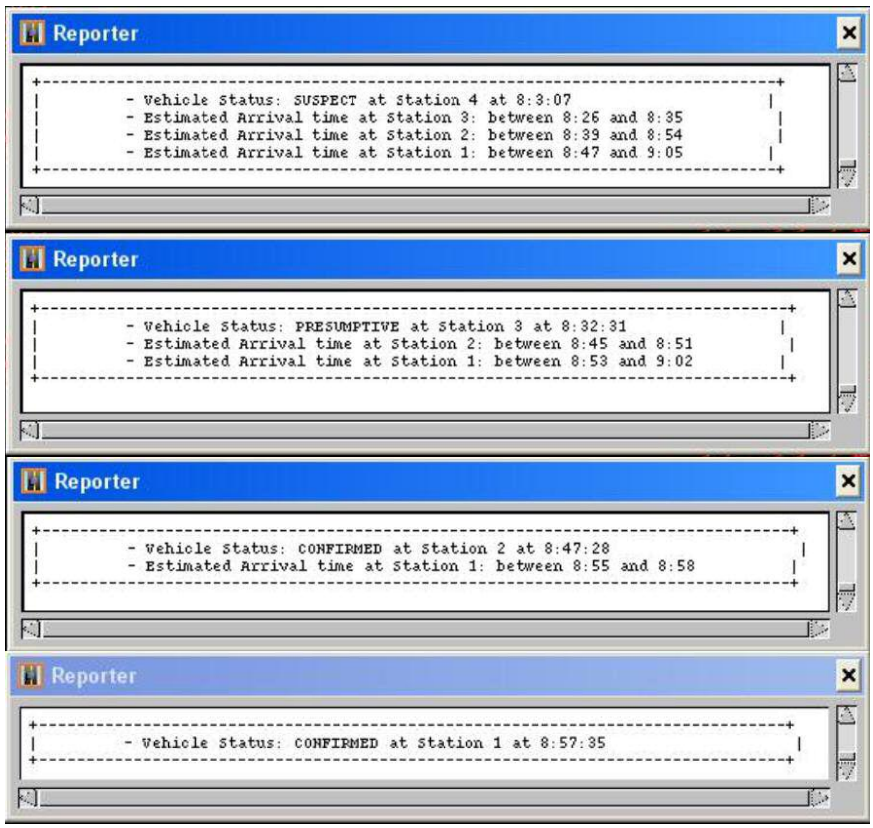

Figures 4.20: Arrival time estimates – High demand case

Figure 4.21 is another way of presenting the comparison between the predictions (from one station to the next) and the actual arrival times. It clearly shows, that under both traffic demand patterns, the travel time estimation tool was successful in accurately predicting arrival time windows for special target vehicles.

Figures 4.21: Comparison of predicted vs. actual arrival times

Summary of accomplishments

Our research team was able to perform the tasks identified at the beginning of the project, and listed in the introduction chapter of this report.

The project end-products consists of four main deliverables:

- A simulation model of the NJ Turnpike, capable of representing in details the geometry of the roadway facility, and a range of traffic patterns traveling through the modeled network. In addition, the base infrastructure required by the new detection system was replicated in the simulation model by installing four special sensor stations at specific locations.

- A detailed set of ConOps for the I95 corridor scenario that were closely coordinated with the local stakeholders that focused on the practical aspects of detecting, tracking and interdicting a terrorist threat.

- A analysis of the system operation that included the impact of false alarms, alert rates, time delays and police procedures.

- a report documenting the work performed, accompanied with a CD Rom containing a series of movie animations (AVI files) illustrating the different components of the simulation.

Recommendations for Future Work

The following issues are among those that should be pursued (by further analysis and interviews, through discovery workshops, and/or simulation) for the exemplary RADS deployment (or its equivalent elsewhere) as described herein:

I. ConOps Development

• Does the overall ConOps make sense to the people being asked to execute it?
• Do the assigned roles make sense? Could they be assigned differently to better effect?
• Are the assigned agencies capable of performing their required functions? How much special training or personnel additions might be required?
• Is the information communicated appropriate to the assigned task/decision? What other information would be desirable? Can critical information be communicated to patrol cars within the assumed time frames?
• Are the time intervals in this process realistic?

II. Organization of display and detection/alarm information to decision makers

- Are existing ITS tools adequate to address the real-time interdiction issues associated with a WMD threat?
- Could a code like Paramics be used to supplement target tracking in an operational system?

III. Characterization of nuisance sources

False-positive interdictions are costly in resources and in disruption to highway traffic.

False positives resulting from environmental noise can be reduced by obtaining multiple readings. However, false positives due to "nuisance sources" of radiation—benign, normal radiation sources being routinely transported on the highway—are not so readily controlled.

A necessary part of the solution to this problem is to obtain, by direct measurement, a good inventory of the actual nuisance sources that would be experienced on the subject highway—the number and type, together with other associated characteristics of interest.

IV. False positive rate reduction for nuisance sources

Several approaches to this problem should be pursued.

First, we can test for isotopic content of the material. When this can be done at a fixed RADS station, this will be helpful. When this is not possible, the interdiction force should make isotopic measurements. Although this will in itself obviously not reduce the false interdiction rate, it will be a crucial step in the process of confirming the type of cargo on a vehicle.

A second approach is to profile vehicles and occupants and exploit specific Intel. When this can be done with some confidence, it can be a useful contributor to lower false interdiction rates.

The third approach is potentially the most powerful one (but also most costly and complex institutionally). With a "radiation material authorization system," we would require all known radiation emitters transported on highways to be registered, and to obtain a special license for their movement. By requiring a special electronic license, e.g. via an E-ZPass tag with a special code, many nuisance sources could be identified and their authorizations validated by remote means. Special measures would have to be taken to make it difficult for this system to be abused, and it would be necessary to carefully test and vet any such method to assure its security.

V. Explore how alert level and Intel can be best integrated with other RADS data

VI. Explore how alert level and Intel influence the acceptable false positive rate in any specific system

Based on our experience in this project and the findings presented in the previous sections, we would suggest the following very rough order of operations for any next test deployment:

Identify a candidate region for the next test case deployment.
Identify the routes to be monitored in that deployment.

Identify participant agencies. Those would include agencies:
- directly involved in equipment operations (as was the NJ Turnpike Authority in the FY03 case);
- directly involved in the response (as was the NJ State Police in the FY03 case);
- determining the legal framework for interdictions, and
- whose jurisdiction would be impacted such that they need to be part of the deployment process.

Conduct interviews and (where called for) table top exercises with those agencies to:
- build their familiarity with the system;
- identify their concepts of operations (ConOps), and how those would be adapted to the RAD system, so the system can be designed to those ConOps;
- determine what legal framework can be developed, and any constraints placed upon the system by that framework;
- gain their buy-in for system deployment.

Key ConOps to explore: Two key ConOps sequences in the interdiction force:
- HQ Sequence: monitor - assess threat - decide- request interdiction (i.e., dispatch).
- Interdiction Sequence (field, cruising officer): receive information, seek, intercept,

pull over, engagement (i.e., approach, interrogate, visual scan, scan with hand-held monitor, decide action).

Key ConOps questions to explore:

- How would a Threat Assessment Officer monitor and act on RADS outputs?
- What interdiction reaction times could we expect?
- How many nuisance alarms per day would the interdiction force tolerate?

Our experience suggests that we need to develop detailed procedures to try out with the interdiction force in order to get well-thought-through responses to the above questions.

Characterize plausible detectors and detector suites.

Deploy background-data-collection detectors and detector suites along one or more of the routes as necessary to collect background data. Note that in this background-data-collection phase, some fairly extensive detector suites may be called for to collect several background radiation types/energies, though those suites may not be found to be cost effective for full deployment.

Analogously, characterize promising non-radiation detection algorithms, such as the one presented in the fourth example of Section 6.2, and so identify detection suites other than radiation detectors. The example suggests a detection suite including weight, axle spacing (treadle signals + speed), vehicle type based on automatic image processing, and loop detector signature. Then deploy that suite for background data collection.

Key question to be addressed early: Can the system discriminate vehicles well enough without a medallion system, or will it take a medallion system in order to make the system work acceptably well?

Based on the ConOps information and the background data collected, design a staged deployment of the system. The stages should be designed to learn-as-go, demonstrate concepts, develop ConOps as gain experience, and establish and collect effectiveness measures. One feature to explore: The development of nuisance-alarm and permitted-vehicle databases, to explore the feasibility of vehicle-matching to those databases to reduce nuisance alarms.

Additional Issues: The ConOps of the nuclear weapon specialist team, and interdiction force ConOps while awaiting nuclear weapon specialist team arrival, was not investigated in the FY03 effort. They need to be investigated on the matters that most directly interrelate with the operation of the rest of the system:

- What criteria should be used to decide when to call in the nuclear weapon specialist team?
- What information should be provided to the nuclear weapon specialist team?
- What response times can be expected?
- What should the interdiction force do until the nuclear weapon specialist team arrives? In particular, what is to be done with the vehicle/occupants and the media?