NCRST
NATIONAL CONSORTIUM ON REMOTE SENSING IN TRANSPORTATION
FREIGHT—PORTS

U. S. Department of Transportation
Research & Innovative Technology Administration

MeTrIS: Metropolitan Transportation Information System: Applying Space Based Technologies for Freight Congestion Mitigation

FINAL REPORT

Contract number DTOS59-07-H-0003
2007 July 1 to 2010 June 30

Michael F Goodchild, PI
Val Noronha, Project Director
with contributing authors
from partner institutions

National Consortium on Remote Sensing in Transportation
Department of Geography
University of California
Santa Barbara CA 93106-4060 USA

www.ncgia.ucsb.edu/ncrst

Table of Contents

Glossary

AA	Agreement Administrator (USDOT)
AOTR	Agreement Officer Technical Representative (USDOT)
CALMITSAC	California Marine and Intermodal Transportation Systems Advisory Council
Caltrans	California Department of Transportation
DGRC	Digital Geographic Research Corporation
DSRC	Dedicated Short Range Communication, a short-range roadside communication network contemplated for intelligent transportation systems
FHWA	Federal Highway Administration (USDOT)
GIS	Geographic Information System
GPRS	General Packet Radio Service, a data service on cellular telephone lines
GPS	Global Positioning System
ITS	Intelligent Transportation Systems
LADOT	City of Los Angeles Department of Transportation
LA/LB	Ports of Los Angeles and Long Beach
Land bridge	A shuttle service of terminal-contracted trucks that moves containers from a terminal to a near-port location for pickup by customer-designated trucks.
MARAD	Maritime Administration (USDOT)
MPO	Metropolitan Planning Organization
NCFRP	National Cooperative Freight Research Program
NCGIA	National Center for Geographic Information and Analysis
NCRST	National Consortia on Remote Sensing in Transportation
POLA	Port of Los Angeles
POLB	Port of Long Beach
RFID	Radio-Frequency Identification
RITA	Research and Innovative Technology Administration (USDOT)
SCAG	Southern California Association of Governments
SMS	Short Message Service, a cellular messaging system commonly known as “texting”
TRB	Transportation Research Board, National Academy of Sciences
Turn	Any of the following, depending on context: (a) a trip by a truck to the port area and back to its company yard, during which it may visit one or more terminals for various purposes, (b) a trip to the port area to pick up a container, then to a customer location for delivery, and back to the company yard, (c) a single visit to a terminal. Turn time and turns per day are measures of productivity.

UCSB	University of California, Santa Barbara
UW	University of Washington
VII	Vehicle-Infrastructure Integration, a USDOT program
WIM	Weigh in Motion

Acknowledgments

This report was authored by Michael F Goodchild, Principal Investigator, University of California, Santa Barbara (UCSB), Val Noronha, Digital Geographic Research Corporation. Richard L Church, UCSB, Anne Goodchild, University of Washington and William F Lyte, California Marine and Intermodal Transportation Systems Advisory Council (CALMITSAC). Patty Senecal and John Glanville served as topical industry consultants. Graduate students Linna Li, Ting Lei and Wenjuan Zhao were research assistants.

We are deeply indebted to many organizations and individuals for their role in the success of the project, in particular:

- The U.S. Department of Transportation, Research and Innovative Technology Administration, for funding and support of this project.
- Cost sharing partners and supporters: Tele Atlas North America, Digital Globe, APL, Port of Long Beach, California Department of Transportation, Los Angeles Department of Transportation.
- The 250 anonymous truck drivers and dozen motor carriers who volunteered their vehicles in support of this study; and several other trucking executives, notably Bob Curry (Cal Cartage), for advice and assistance.
- The port freight community: Geraldine Knatz and Mike DiBernardo (Port of Los Angeles), Don Snyder and Bob Kanter (Port of Long Beach), Jack Carbone (APL), Harbor Trucking Association, Harbor Association of Industry and Commerce, Jesse Glazer (Federal Highway Administration), Alan Hicks (Maritime Administration), Larry Orcutt and Doug Failing (Caltrans), Stephen Cheung (Office of the Mayor of Los Angeles), Hasan Ikhrata, Mike Armstrong, Annie Nam and others (Southern California Association of Governments), Michele Grubbs (Pacific Merchant Shipping Association).
- Our Steering Committee, who generously offered their time and effort without compensation: Kerry Cartwright, Don Cooke, Deborah Estrin, Jim Fitzgerald, David Maguire, Frank Quon, Randy Rogers and Linda Styrk.
- Consultants working on parallel projects: Gill Hicks, the Tioga Group, Cambridge Systematics, Ramon Grijalva, Jerry Wood, Kimley-Horn and Associates.
- Research partners who participated in earlier iterations of the proposed work: the late Reg Golledge, Kostas Goulias and Martin Raubal (UCSB), Dean Rosenberg (AirSis), Reg Souleyrette, Shashi Nambisan and the late Tom Maze (Iowa State University), Ed Van Eeckhout (Los Alamos National Lab), Demoz Egziabher (University of Minnesota), Genevieve Giuliano and Petros Ioannou (University of Southern California), Mark Bradley, John Rosati.
- Support staff at UCSB, most notably Beilei Zhang and Nancy Ponce (Department of Geography), the Office of Research, and Transportation and Parking Services.

Executive Summary

Port operations are at the heart of some of the most dynamic metropolitan centers in the world: London, New York, Los Angeles, Singapore and Hong Kong, to name a few. Ports are critical cogs in national and local economies, but their operations are associated with slow, heavy trucks and trains, congestion and pollution. The negative impacts on their surroundings hinder growth, jeopardizing sustainability of the economic benefits.

The twin ports of Los Angeles and Long Beach are perfect examples of these forces at work. They receive more than 40% of U.S. containerized imports, and support 3 million jobs nationwide. They also account for 50% of emissions in the Los Angeles basin. A 3× growth forecast for imports over the next decades raises doubts about the ports' capacity to accommodate further escalation in traffic. Expansion of port facilities requires tens of billions in infrastructure investment, and is opposed by neighborhood organizations that cite elevated cancer rates, noise, vibration, light pollution and traffic congestion.

This project set out to address this problem set in the national supply chain. A vision of a Metropolitan Transportation Information System (METRIS) was proposed by members of this research consortium in 2004, in which real-time data on the transportation system would create live information products, and in conjunction with optimization models and decision support systems, would streamline transportation operations, also addressing environment and security. Funding from the U.S. Department of Transportation's (USDOT) Research and Innovative Technology Administration (RITA) supported an implementation of METRIS in the San Pedro ports of metropolitan Los Angeles.

The Consortium was led by the University of California, Santa Barbara, with Digital Geographic Research Corporation, the University of Washington, the California Marine and Intermodal Transportation Systems Advisory Council (CALMITSAC), and consultants Patty Senecal and John Glanville. A Steering Committee, consisting of experts in port operations, highway operations, geographic information systems (GIS), and large scale tracking, assisted with strategic guidance. Private and public agencies signed up as cost-sharing partners.

Research Components

The research was built around 5 components:

1. Acquisition of GPS data on movement of goods off the ports; techniques to model port entities, operations and their data streams, and to relate GPS data to underlying geographic layers.
2. Basic time-space analyses on the GPS data, to develop information products that could address immediate concerns of port freight stakeholders, such as trucking firms, marine terminals, and local government agencies.
3. Optimization modeling to address inefficiencies in the management of empty containers, in particular their redundant haulage over tens of kilometers between the ports and hinterland warehouses.

4. Optimization modeling to utilize real-time truck location data to synchronize truck arrivals with port operations, particularly the time-consuming extraction of containers from grounded stacks.
5. Outreach to the port freight community, and investigation of a path towards commercial deployment of the analytical techniques and models.

Data Acquisition

About 250 drayage truckers operating principally in the San Pedro ports were recruited and instrumented with GPS and telematics devices, that communicated their location in real time. The principal hurdle in this process was the combined effect of a severe economic downturn, that dropped cargo volumes 20% in 2008-2009, and a protracted legal conflict between the ports and the trucking industry, triggered by the Clean Trucks Program. These and other forces resulted in attrition of roughly half the drayage trucking fleet operating in the San Pedro ports. By mid-2009 the industry found some stability, it was ripe for new technology to be installed in a fleet of new trucks, and recruitment was timely and highly successful.

This research component included two architectural components: (a) specification of a data model for port and highway operations, building on the Esri-UCSB UNETRANS data model of 2000-2001, and (b) development of conflation algorithms to relate GPS data to underlying GIS layers, a problem generally known as map matching. This latter research project was later developed into a separate research project funded by the National Geospatial-Intelligence Agency.

Basic Analyses

Several local agencies were eager for information products based on MeTrIS. Transportation researchers in planning and academic organizations have historically relied on periodic “intercept surveys” (i.e. conducted by intercepting drivers at rest stops or using police cruisers) to understand the flow of goods. Round-the-clock GPS data proved vastly superior as a data source. We produced the first maps of origins, destinations, routes, congestion spots, and specialized innovative products such as flow “drainage,” showing how port traffic along a freeway is distributed among major freeway exits. We generated maps showing variation in freeway travel time in the course of the day, and drive-time comparisons between routes.

For participating trucking firms, we generated real-time reports on vehicle location. Particularly useful to the trucking industry was a set of results documenting queue waits outside marine and rail terminals. This proved to be controversial, as some representatives of marine terminal operators were opposed to the release of such figures. Nevertheless, the industry as a whole demonstrated readiness for this flow of objective information to resolve long-standing disputes on queue time.

Empty Container Management

The principal storage site for empty containers is at marine terminal properties at the ports. Consequently, after containers are emptied in warehouse destinations in the Inland Empire, 40-80 km east of the ports, they are hauled back to the marine terminals. When exporters are in need of empty containers, they have them transported from a marine terminal, load and return them for export. This results in considerable redundant transportation of empties. At the height of San Pedro's trade activity in 2007, empty hauls represented 10,000 trips to the port daily, or 400,000 km daily in truck travel.

We proposed a set of inland Empties Storage Yards (ESY), that would store an inventory of empty containers for export use. The number and capacity of ESYs would be determined by optimization and decision support models, based on patterns of supply and demand of empties; supporting information systems would track and balance inventories. Our models show that just one or two ESYs would deliver considerable benefits: truck travel could be reduced by 75,000 km daily, resulting in 4,500 fewer port entries, savings of 20,000 liters of fuel and 50 tons of CO_2 per day.

A solution previously promoted by the ports, the Virtual Container Yard (VCY), was unsuccessful. We believe this was due to insufficient consideration of institutional and human factors, and that a physical yard, supported by appropriate policies, would receive greater industry acceptance.

Synchronization

The most time-consuming operation in a truck transaction at a marine terminal is the extraction of a container from a grounded stack. If the container is near the bottom of the stack, the boxes above it must be relocated to gain access to it. Trucks queues develop, causing waits that can extend for hours.

Clearly, containers should be sorted so that those most likely to be requested in the short term are shuffled towards the top of each stack. Knowledge of a truck's location within the metropolitan area and its progress towards the port is one valuable clue to whether or not the driver is likely to keep his appointment (if an appointment system exists) and how soon a container will be requested. Crane operators can use this information in deciding where to reposition boxes while sorting.

The principal difficulty with this proposition is that it appears to require a high penetration of GPS locators in the trucking industry to provide dense information on the sequence of truck arrivals, to be useful to a crane operator. This made it difficult to deploy a working prototype in the ports in the time-frame of the study. However, benefit models were developed based upon varying assumptions of penetration and the quality of information that could be derived from incomplete polling of the port truck fleet. The models predicted that in the best case, given real-time location and arrival sequence information on all trucks combined with optimal sorting of the stack, operations could proceed 15% faster. Moreover, the research concludes that some benefits can be realized even with low penetration of tracking and information flow.

Outreach and Commercialization

Outreach was an essential dimension of the project from the outset, to recruit participants for tracking, and to generate feedback from the freight industry for the modeling proposals that were to follow. Project staff frequently presented project overviews and updates to industry organizations, particularly the Harbor Trucking Association and the Harbor Association of Industry and Commerce. Federal, state and local government officials were briefed on progress. Two web sites were maintained, one with an academic flavor hosted by UCSB, that focused on the research components, the other with an industry orientation, hosted by DGRC. The consortium also undertook to design and maintain the web site for the broader remote sensing program on behalf of USDOT, to assist with its outreach goals.

The vehicle tracking component of the project has obvious commercial potential, that pre-dates the project and is not unique. The modeling components are original and unique, but as they represent a departure from current practice and some cost-benefit complexities, they require a period of consultation and gestation to gain industry acceptance. Commercialization strategies were explored, that would initially offer popular tracking services, and later combine modeling and planning. The analysis concluded that tracking on its own would be a difficult service to offer in that the field of competitors is large, but that in the long term the advantages of specific port-oriented efficiency offerings could differentiate the service.

Conclusions

This project advanced technologies and models of immense scope, which could have been the subject of several years of study. An early challenge was to define its boundaries. With a 2-3 year timeframe, the goal could not realistically have been to bring significant change in practice to the San Pedro ports. Only a small sample of vehicles could be tracked. The modeling methods relied on widespread subscription, which could not be achieved in the short term.

Hence the approach was to establish, by modeling, simulation and real data where possible, that significant benefits could be achieved by adopting the suggested strategies of extensive fleet behavior analysis, empty container management and port synchronization. While there were difficult institutional barriers to implementation, we took the position that the prospect of universal benefits—not just in traditional measures of transportation efficiency such as Vehicle Kilometers Traveled (VKT) and Average Annual Daily Traffic (AADT), but also by mobilizing billions of dollars' worth of stagnant goods inventories—would at least stimulate and sustain dialog.

The project exceeded expectations and was remarkably successful in evolving the mindset of the port freight community. Most major trucking firms participated, and continue to subscribe to MeTrIS. This was the first and still the only technology that could monitor queues outside marine terminals, objectively, accurately, continuously and cost-effectively. At least one terminal is reported to have changed its practices in response to MeTrIS information, expanding its land bridge service. The project was given generous coverage in the *Journal of Commerce* and the *Cunningham Report*, the two dominant industry reports. Federal, state, local and industry officials expressed unanimous support and approval.

However, this work is only a beginning. Major challenges remain to be addressed. Implementation of empty container management requires that the economic and other interests of all parties be addressed. Motor carriers derive revenue by moving empty containers; they are unlikely to support policies that stifle that revenue. A system of carbon credits, that offsets revenue losses incurred in the interests of environmental responsibility, may be the answer. Terminal operators need to be convinced of benefits to themselves, rather than just to trucking firms, by improving efficiency (the research establishes that this is the case); and ultimately the proposed solutions will succeed or fail depending on the acceptance by and training of crane operators and truck drivers.

A significant and deliberate omission was air quality considerations, such as truck emissions while idling, their impact on health, and implications for siting public facilities such as schools, playgrounds and health care facilities. It was felt that air quality was a substantive area of research in itself, and would constitute a distraction from other directions of this study.

At the time of writing this report, in late 2010, most of the trucking firms that had participated in the study had signed on for the commercial rollout of the MeTrIS tracking program. Separately, a consortium of ports, marine terminals, trucking companies and beneficial cargo owners were in discussions about purchasing regular reports on congestion levels in and around marine terminals, developed from MeTrIS observations. In short, the technologies proposed and developed in this study are already deployed, and for now are proving to be commercially sustainable.

Recommendations

We recommend the following to maximize the future benefits of the study:

1. The program of gathering location data from port trucks must continue. A commercial service has been launched, and motor carriers have responded enthusiastically.
2. Marine terminals, motor carriers, ports and beneficial cargo owners should agree on appropriate metrics of "turn time" and strategies to reduce it. At the time of writing, a Turn Time Stakeholders Group (TTSG) had been constituted and was in discussion with consortium members.
3. The scope of MeTrIS data should expand, to include payload and intended destination in real time.
4. A conference should be held among selected goods movement players, to address the establishment of Empty Storage Yards (ESYs) and incentives to use them.
5. Marine terminals and motor carriers should be encouraged to implement the synchronization proposals advanced by this research effort.
6. Port planners should be prepared to consider radical changes in the process of container pickup, including (a) a "taxi service" analogy in which, at peak congestion times, any truck is assigned the next available container, (b) land bridges, (c) container racks.
7. In cooperation with MARAD and FHWA, RITA should assist in the promotion of solutions to facilitate freight movement in the vicinity of this critical national facility.

The optimization models explored in this project necessarily simplify reality. However, they should not be dismissed on those grounds. They suggest that massive benefits can be realized. With $80 million in goods passing through the ports each working hour, the cost of inefficiency is exceedingly high, not just in economic terms, but also in terms of future readiness, environmental and security vulnerability.

1 — Introduction

Freight Congestion in Port Metropolises

Freight Growth and Congestion

Over the past 20 years, due to global manufacturing shifts and free trade agreements, the U.S. economy has grown increasingly intertwined with those of others, particularly in Asia, with China being a dominant player. From 1989 to 2009, the value of U.S. merchandise imports increased 3×, and the ratio of imports to GDP rose from 9% to 11%. While German and British imports grew about 3× each, imports from China rose 25×, and China's share of U.S. imports expanded from 3% to 19%, taking over from Canada the role of prime exporter to the United States. Exports grew 2–3× to most countries, and 12× to China, which now accounts for 7% of U.S. exports[1]. These statistics reflect both a massive swell in trade volumes, and a quantum shift in trade from east coast ports and even north-south land routes, to the west coast.

According to global trade forecasts, assuming stability of political and free-market forces, these trends are set to continue, raising questions about the adequacy of freight infrastructure: port and intermodal facilities, highways, rail, vehicle fleets, labor pools and support technologies. There will inevitably be short term congestion—there has been for several years—and consequences for metropolitan level productivity, the environment, and security, not just within freight facilities but throughout seaport/inland port metropolitan areas that form the bottlenecks in global trade flows.

This is the heart of the problem set that this project set out to address.

San Pedro Ports

Southern California is a prime case study of port growth and congestion. The ports of Los Angeles and Long Beach (LA/LB), together known as the San Pedro ports, serve as the landing zone for more than 40% of U.S. containerized imports, and support 3.3 million jobs nationwide, directly or indirectly[2]. The value of goods moved through the port complex is roughly $80 million each business hour. Until recently, a quarter of diesel particulate matter emissions in the Los Angeles basin were estimated to be related to port activity[3].

Other major container ports in the U.S. are New York-New Jersey, Savannah and Oakland. In 2007, LA/LB handled 2.7 times the volume of second-ranked New York-New Jersey. Expansion of the Panama Canal, currently under way, may lead to diversion of some trade to

[1] U.S. Department of Commerce: International Trade Administration; Trading Economics.com; The World Bank Group.

[2] Ports of Los Angeles and Long Beach and Alameda Corridor Transportation Authority 2007 "Economic Impact Study Finds Trade Moving Through Ports of Los Angeles, Long Beach and the Alameda Corridor Significantly Impact California's Economy," Press release, 2007-03-22.

[3] CALMITSAC 2006 "Growth of California Ports: Opportunities and Challenges: An Interim Report to the California State Legislature." California Marine and Intermodal Transportation Systems Advisory Council.

east cost ports[4], and Canadian and Mexican ports are slated for expansion, but the dominance of LA/LB and trade volumes handled by it are not expected to diminish.

Not surprisingly, trade statistics for the San Pedro ports mirror national trends in freight volumes. In fact, the San Pedro ports are the growth poles of freight expansion on the west coast, that wield enormous influence on national freight numbers. This raises troubling questions on growth capacity. The ports are competitive businesses, intent on expanding. But metropolitan Los Angeles already faces the most congested traffic in the nation[5]. The freeways leading out of the ports were designed four decades ago, prior to the Asian trade boom. Emissions from slow-moving and stationary vehicles correlate with elevated cancer rates in communities surrounding the ports. Clearly there is a need for more infrastructure, but given the scarcity of construction funds, short-term measures are required to utilize existing infrastructure more efficiently.

Several measures have already been implemented in the past decade: (a) on-dock rail lines at most marine terminals divert about 50% of container traffic off roads; freight is shuttled up the Alameda Corridor, a $2.4 billion sub-grade rail facility, to marshalling yards in central Los Angeles, (b) the Off-Peak program incentivizes port trucks to ply during evening and night hours, by levying fees on daytime truck transactions at the ports, (c) $32b in additional goods movement infrastructure is planned in the LA basin over the next decade, including grade separations, expansion of freeways, and consolidation of rail lines.

Freight infrastructure decisions in southern California are driven by the Southern California Association of Governments' (SCAG) periodic Regional Transportation Improvement Programs (RTIPs) and Regional Transportation Plans (RTPs), specifically the SCAG Heavy Duty Truck Model. SCAG partners with the Los Angeles County Metropolitan Transportation Authority (MTA or Metro) and the California Department of Transportation (Caltrans) to develop construction projects. At the time this study was conceived, patterns of truck movement in the LA basin were only anecdotally understood. The data source for the SCAG truck model was periodic intercept surveys and written questionnaires[6] (this remains accepted practice worldwide, and this comment is not intended to be critical). A need was felt for more current, continuous and comprehensive data on origins, destinations, routes and congestion hotspots, that might be provided by sensing technologies such as weigh-in-motion (WIM) devices, radio frequency identification (RFID) and global positioning systems (GPS). Numerous other aspects of freight movement, from enforcement (such as use of non-designated routes, illegal parking and speeding) to the study of futuristic policies such as congestion pricing, stood to benefit from real-time data on truck movement. Efforts by local government agencies and consultants to obtain GPS data on freight trucks had not been successful, for a variety of reasons.

For its part, the port drayage trucking industry was in need of better documentation of vehicle movement and delays, and improved communication with marine and rail terminals.

[4] Drewry Shipping Consultants Ltd 2008 "U.S. Transpacific Intermodal Today and Tomorrow."
[5] Texas Transportation Institute 2007 "2007 Annual Urban Mobility Report."
[6] Example: Meyer, Mohaddes Associates Inc 2003 "North Los Angeles County Truck Study Report (NCTS)." Report prepared for North Los Angeles County Subregion and the Southern California Association of Governments.

MeTrIS Vision

In 2004, members of this research consortium developed a vision of a futuristic freight transportation system supported by rich, real-time geographic information. The Metropolitan Transportation Information System (MeTrIS) would be fed by tracking devices in vehicles, as well as other roadside sensors such as loop detectors, traffic cameras and license plate readers. Sensors on fixed infrastructure (e.g. bridges) would simultaneously report their status. MeTrIS, in conjunction with appropriate visualization, analysis, modeling and communications, could inform real-time operations and security, as well as longitudinal planning of infrastructure and transportation policy. Implementation of MeTrIS would require technical development, coordination with other technical efforts, notably Intelligent Transportation Systems (ITS) and related technologies, Vehicle-Infrastructure Integration (VII) or IntelliDrive, as well as overcoming a set of forbidding institutional hurdles, data confidentiality being of paramount concern.

The MeTrIS vision was presented, developed and refined over two years, in consultation with state, local and private agencies. It became apparent that in view of their congestion issues and enormous economic impact, the San Pedro ports were an urgent candidate for a first implementation.

Project Background

Commercial Remote Sensing and Spatial Information Technologies Program

The U.S. Department of Transportation (USDOT) was tasked under the Safe, Accountable, Flexible, Efficient Transportation Equity Act—A Legacy for Users (SAFETEA-LU) of 2005 to administer a research program in Commercial Remote Sensing and Spatial Information Technologies (CRS-SIT). In 2006, the Research and Innovative Technology Administration (RITA) issued a Broad Agency Announcement for proposals from universities on research in space based technologies, that would address Freight Congestion Mitigation, among other application areas.

Proposal

Our proposal to USDOT was to implement MeTrIS in the ports, to gather data on truck movements, to develop a data model, to analyze the data, generating a series of useful information products to serve local authorities, and to develop optimization models to ease congestion and to facilitate freight flows. Two process models in particular were proposed: the first sought to minimize the transportation of empty containers—an economically unproductive activity, that needlessly generates traffic flows and environmental impacts. The second model focused on interactions between trucks and marine terminals, in particular the process of picking up containers from grounded stacks. The proposal included a component to explore commercialization opportunities. In addition there were tasks of administration and outreach.

There were multiple challenges in the project: to recruit a significant base of trucking firms or individual owner-operators to participate, to develop technical specifications, analyses and models, and to overcome institutional resistance to changing established practices.

Consortium Members

The University of California, Santa Barbara (UCSB), in cooperation with Digital Geographic Research Corporation (DGRC), the University of Washington (UW), the California Marine and Intermodal Transportation Systems Advisory Council (CALMITSAC) and consultants Patty Senecal and John Glanville, proposed this project. UCSB and UW developed the data models and optimization models, DGRC was responsible for industry liaison, vehicle instrumentation, data management and analysis, CALMITSAC and Patty Senecal provided industry advice and contacts, and John Glanville offered advice on commercialization.

As proposed, the consortium appointed a Steering Committee, with representation from the Burlington Northern Santa Fe (BNSF) railroad, Caltrans District 7 (Los Angeles), Environmental Systems Research Institute (Esri), Port of Los Angeles, Port of Seattle, TeleAtlas, USDOT Maritime Administration, and University of California, Los Angeles.

Report Structure

This report begins with the technical architecture of geographic data models in which MeTrIS data are embedded, in Chapter 2. Chapter 3 discusses industry issues and recruitment of motor carriers, and basic analyses performed on the data. The next two chapters delve into process models: empty container management is covered in Chapter 4, and Chapter 5 describes truck-terminal synchronization. Chapter 6 reports on outreach and commercialization activities, and Chapter 7 offers conclusions and recommendations.

2—Data System Architecture

This chapter explains the data architecture and fundamental analytical processes related to processing geographic data, and GPS data in particular. The first section covers a data model appropriate to port objects and activities. Next we discuss practical considerations of vehicle tracking instrumentation. The third section describes the process of conflation, which relates GPS data to underlying geographic databases, which is generally the same problem as that of relating two disparate geographic databases to each other.

Port Data Model

As the problems faced by society have grown ever and ever more complex, the research projects that address these problems have become increasingly multidisciplinary, often involving collaborations not only between disciplines, but also between researchers separated by large geographic distances. The approach to science that has emerged in the past two decades entails the sharing of data and tools to an unprecedented extent, and requires that researchers establish and follow a variety of standards, covering the syntax and meaning of data, and the functionality of tools. In general these requirements are referred to as interoperability, and in the geospatial world the primary developer of standards and specifications has been the Open Geospatial Consortium.

A project such as MeTrIS involves large amounts of disparate data, and numerous methods for analyzing, modeling, and transforming these data. The collaborating team is spread over several institutions, and includes people with training in several disciplines, from geography to operations research. To support the work, therefore, we chose to invest heavily in data modeling, in order to establish a common baseline of terms, classes, and data that would support the project and be compatible with emerging standards in transportation data modeling.

One such standard is UNETRANS, an object-oriented design developed in collaboration with Esri, the leading developer of GIS software. In the late 1990s Esri adopted an object-oriented approach to the design of its databases, and provided supporting tools in its main product, ArcGIS Version 8. Subsequently, it sponsored the development of a series of standard data models for different GIS application domains, in collaboration with the community. UNETRANS was conceived as such a data model for the specific needs of transportation applications of GIS. It was developed by Esri and UCSB, with input from several workshops and meetings of the transportation GIS community.

The MeTrIS project adopted UNETRANS, and extended it in several ways to make it suitable for the management of GIS data about ports, and support for port operations. The following sections first describe GIS data modeling in general, with specific reference to Unified Modeling Language (UML); then discuss UNETRANS; and then describe the extensions necessary to accommodate the needs of MeTrIS.

Object-oriented data modeling

A data model can be defined as a template for the data relevant to a specific computer application. Consider, for example, the task of capturing the data needed for a GPS navigation system in a car. To provide advice on how to navigate from an origin to a destination, the system would have to know first about all of the streets and other rights-of-way. The street network might be represented as a collection of connected lines, each line representing the stretch of a street between two adjacent intersections. In addition to the geometry of each such line, the system would need to know about the line's attributes: its name, whether or not it was one-way, the number of lanes, etc. It would also be important to know about the intersections that connect these lines, along with each intersection's attributes: whether there are traffic lights or a four-way stop at the intersection, for example. The database might also include a collection of points of interest, such as hotels, restaurants, schools, or churches, in order to recognize these when a user specifies a destination, and to treat them as landmarks in driving directions.

This discussion has identified three *classes*, or types of objects essential to the task, along with the attributes needed to distinguish one member of a class from another. The first principle of object-oriented data modeling is that all objects are members of classes. In UML, a visual language for the description of an object-oriented design, each class is represented as a box divided into three boxes, one above the other. The top row defines the name of the class. The second is used to define all of the attributes of the class. The bottom row is used to accommodate any methods defined for the class, but is often ignored. It is also possible to think of a class as a table, with each row defining one member of the class and each column defining one attribute.

In a GIS application each class of objects will belong to one of the fundamental data types of GIS: points, lines, and areas. In the example of the GPS navigation system, the lines of the street network are a specific application of the general line data type of a GIS. Intersections are an example of the point data type, as are points of interest. We term this a *specialization* and say that the specific application *inherits* the properties of the general class, and symbolize it in UML with an open triangle pointing from the specialized class to the general one. Thus the lines inherit the properties of GIS lines. In general some attributes will be defined for the general class (in the case of lines, this would include length, since all lines in a GIS have length) and some will be defined only for the specialized class.

Four additional types of relationships are defined in object-oriented modeling for GIS. The lines of a transportation network may have *subtypes*, which might include railroads, bicycle paths, and canals in addition to streets. An object in one class may be related to an object in another class through an *association*. For example, the connectivity between lines at an intersection might be represented as associations between the line class and the intersection class. Classes can also be defined as *compositions* or *aggregations* of other classes; for example, a collection of street lines might be defined as aggregating to form a route.

UNETRANS

The UNETRANS model is designed to accommodate a wide range of transportation applications of GIS. Its goal is to provide a common framework, by establishing classes and relationships that can be incorporated many different projects, and serve as a *lingua franca* between them. Structuring a GIS database with a data model allows its contents to be shared, edited, and described in a simple, uniform way. Attributes of defined classes can be added to accommodate the needs of specific applications, and new classes and subtypes can be defined to enrich the existing ones by adding greater specialization.

A suitable test of a data model is to ask whether it includes places to store every item of data that is important to a given application. Any application, or use case, must provide certain defined outputs. These outputs depend on inputs, and also on functions that can be used to transform the inputs. Thus a key to data modeling for a given use case is that every needed item of input finds some place in the data model where it can be stored. The use cases that provided the input to the design of UNETRANS were drawn from a range of applications, from public transit planning to wayfinding and highway maintenance. While new use cases may force extensions of the model, the range of use cases collected during the design of UNETRANS provide some assurance that such extensions are not likely to require a fundamental reorganization of the model.

MeTrIS

The majority of the MeTrIS project needs fit easily within the UNETRANS framework. In one key respect, however, the data model needed to support port operations requires a substantial rethinking. Data modeling for transportation normally assumes a linear, network structure to the transportation system, composed of lines and their specializations. But in some use cases it is possible for vehicles and other moving objects to leave the linear structure, moving freely in two- or three-dimensional space. This occurs in shipping when ships deviate from normal shipping lanes; in aviation when aircraft deviate or take shortcuts; and in trucking when vehicles move freely around ports or parking lots. One solution to the problem is to linearize the uncontrolled spaces. For example, we might replace the open space of a port facility with a finite set of connected lines. Map matching takes on new dimensions in these situations, since there is no longer a linear network to which to match a GPS track.

The solution adopted for the project allows vehicles to define their own tracks, without map matching, when the observed GPS track deviates by more than a critical threshold distance from the linear network. This “floating” section of track is anchored at both ends when the truck returns to the linear network. Between anchor points, the GPS track is taken to be the best estimate of the actual path, and used to compute parameters of distance, speed, direction, and acceleration, without reference to any part of the network.

Conclusion

The MeTrIS project has provided an interesting problem in geospatial data modeling, with an important new characteristic in the possibility of unmatched tracks. The data model established for the project is a simple extension of the UNETRANS model, and accommodates this important extension. Continuing work is generalizing this extension to other use cases that involve similar floating elements.

Vehicle Tracking: GPS and Communications

Tracking a vehicle with GPS and communicating the data to a server are now well established technologies. The principal variables are the cost of equipment and data transmission, and the quality of data. Currently, the GIS industry estimates that at least half of all commercial Class A trucks are tracked by GPS. However, the penetration rate for the drayage industry is far lower, because (a) drayage is a relatively low-budget operation, (b) trips are usually within metropolitan areas, they are shorter, driver contact is frequent, and there is a lower perceived need for vehicles to report their location automatically. For these reasons, commercial vehicle tracking systems tend to be oriented towards long-haul trucking. Polling intervals are typically in the range of 15 minutes, well suited to a trip between Los Angeles and Memphis, but far too coarse for a trip from the port to a local warehouse.

GPS devices typically output latitude, longitude, altitude, time, speed, heading, and various measures of signal quality and satellite-specific parameters. Systems vary with regard to how many of these fields are transmitted and recorded. Storage of all fields offers advantages, but at a cost.

Tracking devices can store data on board for periodic physical retrieval, or they can transmit data by satellite, cellular or short-range (WiFi or Bluetooth) wireless communications. Other technologies such as WiMax and Dedicated Short Range Communication (DSRC) require infrastructure that is not yet widely established. Some of the factors to be considered in the selection of a tracking device are: cost of equipment (one-time), cost of data transmission (recurring), location polling interval, communication interval, communication technology, unit size, durability and environmental tolerance (temperature, vibration), safety, power and antenna cabling requirements, antenna sensitivity, and maintenance features.

A survey of the GPS tracking market was undertaken prior to project commencement, because the instrumentation strategy would be a strong determinant of cost, which needed to be specified in the project budget. Two commercial tracking firms, Qualcomm and AirSis, both based in San Diego, were approached as potential partners, and specifications of a number of vehicle tracking devices were evaluated.

Qualcomm declined to participate in the project. AirSis was cooperative, and offered favorable rates and cost match. DGRC had independently developed prototype MeTrIS tracking technology of its own. This solution was found to be more cost-effective and adaptable than any alternative device, and the 2006-2007 proposals to USDOT called for AirSis to provide technical support to install and troubleshoot these units. In the course of the project, installation and maintenance proved not to be major challenges, hence the AirSis role was obviated and assumed by DGRC.

DGRC's tracker was produced in two models. One employed local storage. On-board flash memory provided adequate storage to capture several months of vehicle activity. A second model employed wireless communications. Based on prevailing data costs, a location sampling interval of 12 seconds was implemented initially. This compared extremely favorably against 15-minute polling by most other systems, and could be achieved at far lower cost.

Local storage and retrieval

Within 45 days of the award, a first truck was instrumented with a local-storage tracking unit. GPS points were written to the storage device at 5 second intervals. Local storage has advantages: there are no wireless data transmission costs, hence density of sampling is limited only by availability of sufficient storage capacity.

On the other hand there are drawbacks to local storage. There is no immediate feedback on the health of an instrument—among other hazards, it could be deliberately disconnected by the truck driver or mechanics. A physical visit is required to interact with on-board instrumentation, to download accumulated data. The problem is compounded when multiple trucks at multiple companies are involved.

Wireless transmission

Wireless transmission does have ongoing costs, several potential modes of failure and unreliability, compared with local storage. However it has the advantage of real-time data monitoring, enabling a variety of applications, particularly in operations and security. It could be argued that many of the objectives of this project were long-term in nature and could have been addressed using local storage and periodic retrieval. One of the useful early findings was that wireless transmission and real-time access to data enable a variety of applications of interest to several potential end-users, and the cost and associated problems are well justified in many cases.

Data flow

Data from MeTrIS trackers are transmitted to a cellular tower, and relayed by the cellular carrier to a designated server. Data are sorted and archived by automated processes. Real-time and batch processes produce information outputs detailed in the next chapter.

Conflation

Geospatial data conflation provides data support for transportation analysis and modeling in the MeTrIS project.

Good geospatial data are critical for transportation systems

Legitimate scientific research and wise decision-making require high-quality data as input, and usually need data from a variety of producers, since it is not realistic to collect all data directly. For transportation analysis and modeling, accurate road network data are a vital component to ensure model effectiveness and subsequent conclusions and recommendations. It is not uncommon that a transportation project, e.g., traffic assignment, involves integration of street network data from two sources: one dataset has good spatial accuracy, and the other has better non-spatial attributes. The increasing and rapid development of remote sensing and other technologies as well as the growth of the Internet provide abundant opportunities to collect and access vast volumes of geospatial data. In addition to well-known datasets provided by government such as US Census TIGER/Line files and free data services like Google Earth, large amounts of geospatial information are being produced by commercial companies, such as Tele Atlas, Navteq, and Digital Globe. The OpenStreetMap project is also building a database of detailed geographic information through the efforts of a worldwide network of volunteers. Large volumes of geospatial data have the potential to benefit transportation research, policy development, and decision making. However, it is not always straightforward to take advantage of this abundance due to inconsistency, incompatibility, and heterogeneity among various datasets.

Conflation is a solution to create better data from existing sources

Rather than a visual overlay of data from diverse sources, automated conflation of heterogeneous datasets provides a better solution since it opens possibilities for updating, averaging to obtain better estimates, and analysis and modeling. Conflation is a process of combining information from two or more related datasets, and providing the user with a new dataset derived from input datasets according to the specified requirements. We use the term conflation to distinguish this task, which is defined for vector data, from other data integration tasks, such as the fusion of raster information from overlapping images. The difficulty of conflation depends on many factors such as complexity of representation and the volume and accuracy of the datasets involved. Specifically, incompleteness and inaccuracy of the original datasets, different reference systems, distinct generalizations and representations of reality, semantic issues of terminology and classification, various scales, and different purposes, as well as various time frames all create challenges in the use of geospatial data from heterogeneous digital sources.

The main reason why geospatial data conflation is necessary is imperfection and limitation in representation of a continuous enormous geographic world in a discrete digital machine that is only capable of storing limited information. Due to inevitable simplification and abstraction of geographic data, uncertainty is inherent in all geographic datasets. Therefore, when relevant information for a particular task is scattered in diverse data sources, conflation is required to combine them into a consistent structure in spite of existing discrepancies that are due to approximation of the reality in each source dataset. In this project, we investigated the nature of geospatial data conflation, designed a conceptual framework for conflation, and developed new methods for improving conflation results with particular emphasis on feature matching. We fully studied the necessity and significance of geospatial data conflation and developed a general framework for conflation that defines geospatial features as 4-tuples and the process of conflation as a linear combination of geometric and semantic similarity measures with corresponding error estimation (Adams et al., 2010).

The need for conflation occurred in two somewhat independent ways in this project. First, we had the opportunity to combine data on street networks from various sources, to obtain the most useful combination. Second, we needed to match the tracks of vehicles to a street database, a task sometimes known as map matching. Both tasks are examples of the same general problem: in the second case, a GPS track of a vehicle must be matched to a database. However the first problem is normally tackled in batch fashion, conflating two or more existing databases, while the second task is normally considered to require a real-time solution.

There are generally two major steps in feature matching: first, we choose a similarity measurement to be used as a criterion for matching; second, we identify all matching pairs of features using this selected similarity criterion. If two features in different databases are represented similarly in terms of positions, shapes, and relationships with surrounding features, it is probable that they represent the same entity in the real world and that the small difference between them is caused by different data schemata or uncertainty introduced in the data creation process. Samal et al. (2004) summarized a set of possible similarity measures that have been developed in a variety of disciplines and might be useful in conflation, including categorical similarity, string similarity, and shape similarity. More broadly, similarity measures commonly used in feature matching can be classified into three types according to whether they are based on similarities of geometry, attribute, or topology, or combinations of these.

Optimized feature matching

Although the criteria for feature matching vary in different applications (e.g., Saalfeld, 1988; Samal et al., 2004; Beeri et al., 2004; Safra et al., 2006), a common strategy in most work is the sequential workflow of matching that we call a greedy method: pairs of matched features are identified one after another. An obvious issue with this kind of method is that once a feature is matched to a wrong feature in the other dataset, no remedy can be made to correct this error. In this project, we developed a new strategy for automatic feature matching called optimized feature matching that relies on a global measurement of similarity by regarding feature matching as an assignment problem that takes into account all corresponding pairs of features simultaneously (Li and Goodchild, 2010a). This new method was compared with widely used greedy methods on both synthetic point datasets and real street-network datasets. As a result, the optimization method works better than the greedy methods in terms of percentage of correct matches. Figure 1 demonstrates the relationship between percentage of correct matches and the number of points for 19 pairs of synthetic datasets. When the points are sparsely distributed in the area, both greedy and optimization methods achieve a good percentage. However, as the density of features increases, the percentage of correct matches by the two greedy methods decreases drastically. The percentage of correct matches by the optimization method is relatively stable, close to 100%.

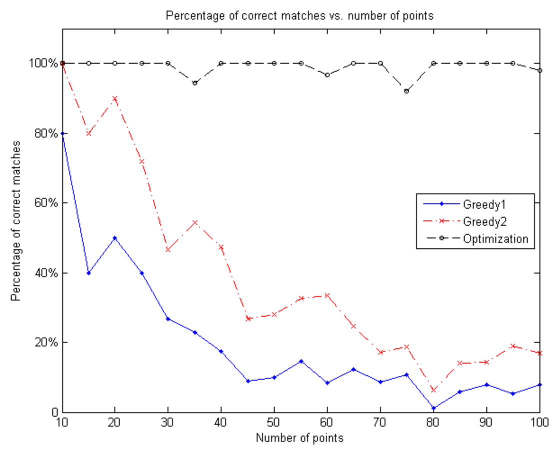

Figure 1. Percentage of correct matches vs. the number of points

These methods were also applied to two real polyline datasets obtained from different agencies which prepared them according to different standards. The criterion for feature matching in this example is the Hausdorff distance between polylines since it characterizes the proximity of two linear features particularly well (Abbas, 1994). Two versions of the same street network of a neighborhood in Goleta are displayed in Figure 2(a), and two versions of streets in downtown Santa Barbara are displayed in Figure 2(b). As demonstrated in Table 1, the total distance between matched pairs using the optimization method is smaller than that using the greedy methods. Furthermore, the mismatch rate is significantly reduced using the optimization method in both cases.

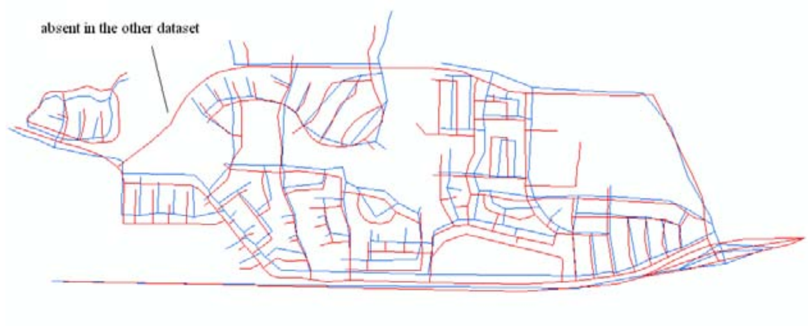

Figure 2. Street networks in a neighborhood of Goleta, CA (more than 200 features)

Figure 3. Street networks in downtown Santa Barbara, CA (more than 1000 features)

To extend this method to address a broader range of datasets, we proposed a new dual strategy for automatic feature matching: the matching process is formulated as an optimization model that takes into account all potentially matched pairs simultaneously by maximizing the total similarity of all matched features; and two approaches are defined based on the nature of the positional distortions: datasets with independent distortion and datasets with autocorrelated distortion. Autocorrelation of distortion is to be expected, since if points were disturbed independently the result would be an unacceptably chaotic representation. Spatial autocorrelation of distortions is an instance of the widely recognized principle known as Tobler's First Law of Geography: "nearby things are more similar than distant things."

For datasets with autocorrelated distortion, we take into account rectification of those positional offsets by incorporating an affine transformation into the optimized feature matching model to improve the matching performance. In addition, this model takes advantage of the asymmetry of a dissimilarity metric—directed Hausdorff distance—to address 1:m correspondences in linear feature matching. In our study, a similarity index is created by a combination of geometric and semantic information: directed Hausdorff distance, angle, and dissimilarity between feature names. The goal of our model is to find the global optimal solution from all possible choices, by maximizing an objective function:

$$\text{Maximize } \sum_{i=1}^{p}\sum_{j=1}^{q} s_{i->j} z_{i->j}$$

where i,j are indices for the features in the first and second dataset respectively, p and q are the number of features in each dataset, and $s_{i \to j}$ is the directed similarity from feature i in the first dataset to feature j in the second dataset. The objective function maximizes the total similarity between all matched feature pairs. The variable $z_{i \to j}$ represents a match between feature i and feature j, taking value 1 if a match is made and 0 otherwise, i.e.

$$z_{i->j} = \begin{cases} 1, & \text{if a match is made from feature } i \text{ to feature } j \\ 0, & \text{otherwise} \end{cases}$$

This model is subject to a cardinality constraint and a length constraint. Consequently, the solution of this model successfully identifies 1:1, 1:m, m:1, 1:0, and 0:1 correspondences between features without aggregation or splitting of features in the preprocessing before matching (Li and Goodchild, 2010b). We tested two approaches for automatic feature matching on six test areas (Figure 3) that contain real street-network data for Santa Barbara, CA, extracted from two different sources. These test areas represent two major types of streets in a street database: streets in urban areas and in rural areas. Table 1 demonstrates the results for the four test areas using the approach for datasets with independent distortion. The percentage of correct identifications varies from one test area to another and mostly depends on the spatial pattern of features, feature density, and discrepancies between the two input datasets. The average percentage of correct identifications is 97.18%. The experiments show that the spatial pattern of the data is more important than the number of features in affecting the performance of feature matching. Both approaches were applied in Test Area 5 and Test Area 6 that have a typical autocorrelated distortion between two datasets. Test Area 5 achieves 100% correct identifications using the second approach, a significant improvement compared to the results without consideration of autocorrelated distortion in the first approach (91.36% correct identifications). In Test Area 6, only two pairs of corresponding features are not correctly identified using Approach 2 due to the difference in feature names, which may be solved if we use a more accurate metric for name dissimilarity.

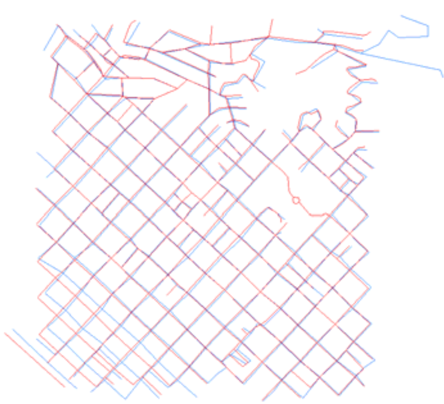

Test area 1

Test area 2

Test area 3

Test area 4

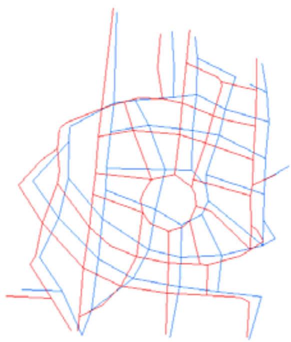

Test area 5

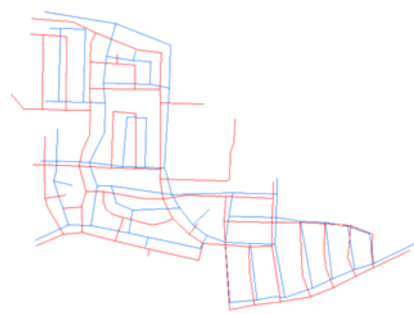

Test area 6

Figure 4. Test areas of street networks in Santa Barbara, CA

Table 1. Results of optimized feature matching for datasets with independent distortion

	Test area 1	Test area 2	Test area 3	Test area 4	Total
Number of features in Dataset 1	434	308	377	344	1463
Number of features in Dataset 2	423	264	374	322	1383
Number of corresponding pairs and singles	450	330	419	362	1561
Number of correct identifications	441	322	410	344	1517
Percent correct identifications	98.00%	97.58%	97.85%	95.03%	97.18%

Table 2. Matching results for test areas 5 and 6

	Test area 5	Test area 6
Number of features in Dataset 1	81	84
Number of features in Dataset 2	80	77
Number of corresponding pairs and singles	81	84
Number of correct identifications using Approach 1	74	79
Number of correct identifications using Approach 2	81	82
Percentage of correct identifications using Approach 1	91.36%	94.05%
Percentage of correct identifications using Approach 2	100.00%	97.62%

Conflation provides a better foundation for transportation information systems

Feature matching is one of the crucial components in conflation. Without correct identification of matched features in different datasets, the subsequent steps such as feature transformation will not be executed properly. The product of conflation not only meets the requirement of creating higher-quality data for transportation information systems from multiple sources, but also provides a great potential to utilize rich yet incompatible geospatial data in order to facilitate spatial analysis and reasoning in transportation modeling. As more road network data are being created by volunteers using GPS tracks, digitized remote sensing images on the Web (e.g., OpenStreetMap), etc., this new source of road data may significantly benefit transportation research and policy making if it can be effectively integrated into existing datasets by conflation. Our research on conflation provides a better foundation for model construction, verification, and validation in the current project; meanwhile, these methods are readily implemented to obtain better input into spatial models and analysis in other projects that rely on geospatial data, especially linear network data.

References

Abbas, I. 1994. Base de donnes vectorielles et erreur cartographique: problmes poss par le contrle ponctuel; une mthode alternative fonde sur la distance de Hausdorf. Computer Science. Paris, Universit de Paris VII.

Adams, B., Li, L., Raubal, M., and Goodchild, M.F. 2010. A general framework for conflation (extended abstract), GIScience 2010, Zurich, 14-17 September, 2010.

Beeri, C., Kanza, Y., Safra, E. and Sagiv, Y. 2004. Object fusion in geographic information systems. Proceedings of the Thirtieth International Conference on Very Large Data Bases - Volume 30, pp. 816-827 (Toronto, Canada).

Li, L. and Goodchild, M.F. 2010a. Automatically and accurately matching objects in geospatial datasets. In: Proceedings of Theory, Data Handling and Modelling in GeoSpatial Information Science. Hong Kong, 26-28 May, 2010.

Li, L. and Goodchild, M.F. 2010b. An optimization model for automatic feature matching in conflation. *International Journal of Geographical Information Science*, in review.

Saalfeld, A. 1988. Conflation automated map compilation. *International Journal of Geographical Information Systems*, 2(3), pp. 217-228.

Safra, E., Kanza, Y., Sagiv, Y. and Doytsher, Y. 2006. Efficient integration of road maps. Proceedings of the 14th Annual ACM International Symposium on Advances in Geographic Information Systems, pp. 59-66 (Arlington, Virginia, USA).

Samal, A., Seth, S., and Cueto, K. 2004. A feature-based approach to conflation of geospatial sources. *International Journal of Geographical Information Science*, 18(5): 459–489.

Several parties could be considered potential "customers" to the project. On one hand, trucking firms that have their vehicles tracked receive benefits in terms of up-to-date information on truck location and daily activity logs. They could potentially pay for this service. However, commercial tracking has existed for several years, and the novelty of the project does not lie in real-time location reporting. It lies in aggregate analysis of travel patterns, congestion delays, and activity sets. The customers for this information are the ports, local governments, industry associations, and the trucking industry.

Recruitment

Confidentiality

By far, the primary barrier to recruitment is in assuring participants of confidentiality of data. There are several dimensions of confidentiality:

- Drivers are concerned about traffic violations such as speeding, illegal parking, and use of non-designated routes.
- Drivers are further concerned that all activities are recorded: use of questionable repair facilities, personal stops, whether for meals or at the dentist's, are available for scrutiny.
- The data reveal the trucking firms' customers, and freight volumes associated with each.
- Delays at privately owned port and rail facilities are documented.
- The data can indicate patterns of freight movement that could potentially be used against local or national interests. Competing ports could point to inefficiencies in San Pedro, or attacks on facilities could be planned with the help of such data.

In addition, the trucking industry is concerned that information generated by analyses, even in the aggregate, can be used against the interests of the industry, e.g. in studies of congestion pricing.

Commercial tracking firms are anecdotally known to sell truck tracks for research purposes, the only stipulation being that "identities" of vehicles, i.e. identifiers such as license plates and company names, are masked. In the case of GPS data, this is entirely inadequate. Analysis of origins and destinations can easily reveal identities of operators and their business relationships.

Human Subjects provisions required each truck owner to provide signed consent to have the vehicle tracked. To comply with Human Subjects requirements, and to assure the trucking industry of confidentiality standards it required to sign up, several security precautions had to be implemented, and access to data heavily regulated. Outputs available to the public were masked to conceal elements that could potentially compromise business interests. Even within the research team, player identities were protected against discovery. It could be argued that these measures detracted from the utility of the data, and perhaps hindered outreach among some potential users. On the other hand, there was broad industry support for our standards of data protection. Given the difficulties other agencies have had in the past, recruiting carriers in GPS-based freight studies, we were privileged to have gained the trust, support and goodwill of the industry.

Unanticipated Barriers to Recruitment

The project suffered from two unexpected problems in developing the fleet sample. The economic downturn of 2008-2009 was reflected in a 20% drop in imports in the San Pedro ports, and at the same time the Alameda Corridor increased its share of container traffic, to as much as 50% at some marine terminals. These factors led to a significant fall in truck traffic.

The second problem was the fallout from implementation of the Clean Trucks Program (CTP) by the ports of Los Angeles and Long Beach in October 2008. The CTP not only mandated the retirement of older vehicles, but also imposed concession arrangements on truckers operating in the ports, and required drivers to be employees rather than contractors. Inevitably there were disagreements and protracted legal battles between ports and carriers, that created considerable uncertainty in the industry. (This is a thumbnail summary of a complex situation, and necessarily simplifies and omits details).

The CTP, economic downturn and rail expansion resulted in attrition of 50% of the trucking industry. The number of drayage trucks plying in the ports dropped from 16,800 in 2007 to approximately 8,000 in 2010.

There are other, older trucks that still work on drayage duty; an emerging practice is for carriers to dispatch newer trucks to the port, and to turn containers over to older trucks outside the port boundaries. These practices, which are reflected in the GPS data, distort the fundamental patterns of drayage trucking.

Sampling

Within 45 days of project commencement, a first truck was instrumented with MeTrIS tracking equipment. Within 6 months, a dozen trucks were transmitting data regularly. Recruitment suffered significantly due to the CTP and recession in 2008, but recovered in 2009. Two years into the project the sample size had grown to 120, and the project ended with 250 trucks from a dozen trucking firms, representing 3% of the port drayage fleet.

The extent to which this is a representative sample depends on the information items of interest. There is no known systematic bias in the sampling, nor is there an explicit attempt to develop a random sample, because the universe of drayage trucking is highly amorphous and poorly understood, and because of severe recruitment constraints. MeTrIS was adopted largely by referral. Many motor carriers were initially naturally mistrustful, and it was only after a track record had been established with other carriers that they signed on. Moreover, with individual owner-operators there was a higher risk of loss of contact with the subject, and loss of equipment. Only over the past two years, pursuant to the CTP, has there been a registry of drayage trucks, and lists have been developed of firms applying for loans to finance new fleets. It is known anecdotally that about 20% of the registered motor carriers carry 80% of containers off the port, hence enforcement of randomness in vehicle selection for its own sake translates to tracking a smaller proportion of traffic.

In early 2010, ten motor carriers formed the Clean Trucks Coalition (CTC), and received permission from the Surface Transportation Board to share equipment and business facilities. The carriers were some of the largest in the industry, with a combined fleet of 800-1,000. More than half of the CTC member companies participate in MeTrIS, and correspondingly constitute about half the MeTrIS companies. Other participants are smaller firms with 10-50 trucks, and two very small carriers with 5-10 trucks. In this regard, our sample is a representative cross-section of the industry.

Maps and Analyses

Routes

Simple plotting of GPS traces (position histories) produces effective maps of routes followed by the drayage fleet. Symbolizing the tracks with color or line style to represent travel speed, the maps clearly distinguish between freeways and surface streets (Figure 5), and indicate congestion spots on the freeway system.

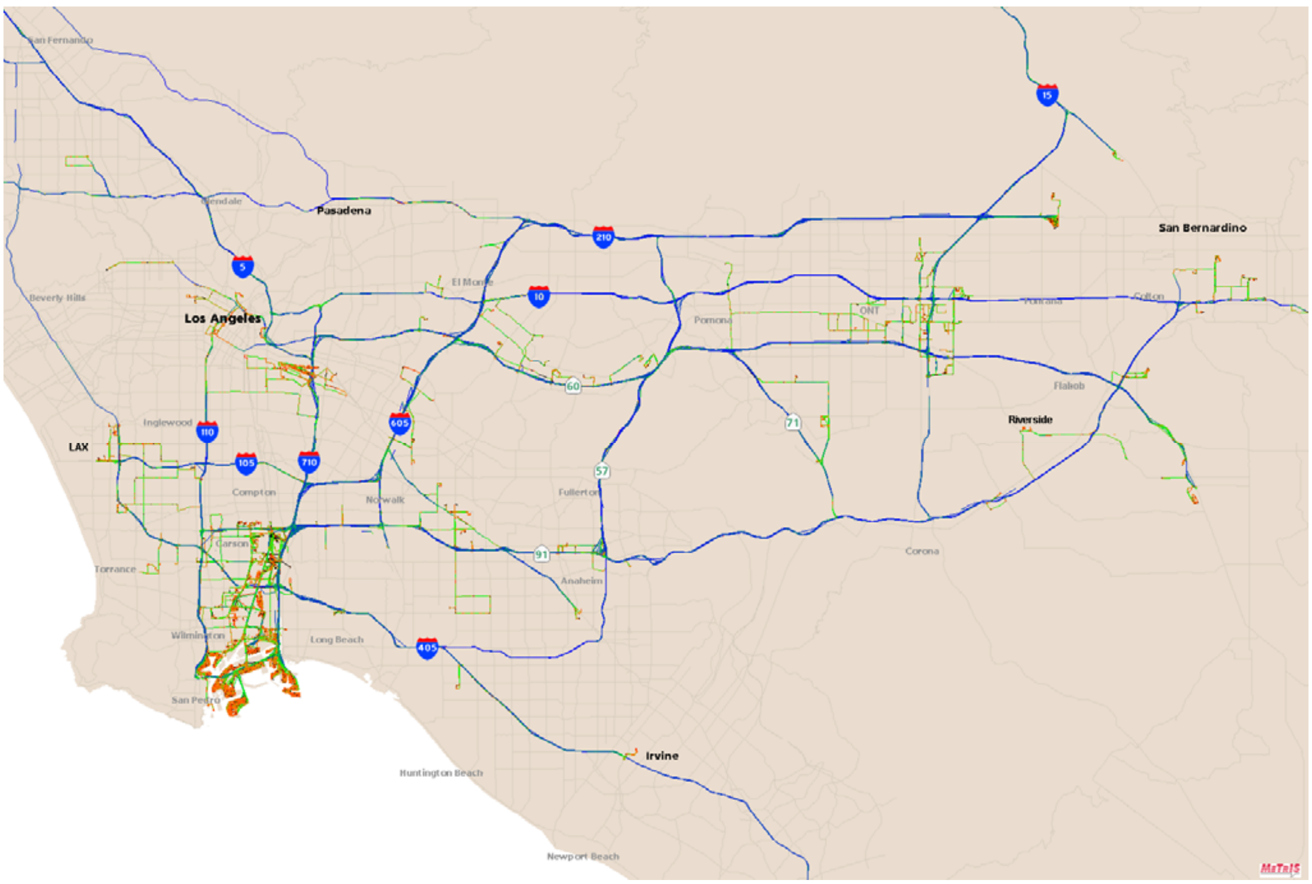

Figure 5. LA basin routes followed by drayage trucks, May 2008. Colors represent speed.

Within small areas such as marine terminals, motor carrier yards and intermodal yards, the geography and time-stamps on traces, combined with aerial imagery if necessary, reveals the business processes of queuing, admission, physical transactions (loading and unloading of containers) and exit (Figure 6). A composite map from all trucks provides an illuminating picture of port layout and circulation (Figure 7).

Instantaneous or real-time mapping of vehicle locations show the density of freight activity in different parts of the LA basin. Clearly there is a lot of traffic is in the ports and the neighborhoods of Wilmington and Carson immediately north (Figure 8). Activity on routes to the Inland Empire appears relatively light, but that is partly because of differences in geographic scale: over greater distances, vehicle icons are spaced further apart, but a count of vehicles on those routes shows more activity than the graphic suggests.

Figure 6. Path of a truck through a marine terminal

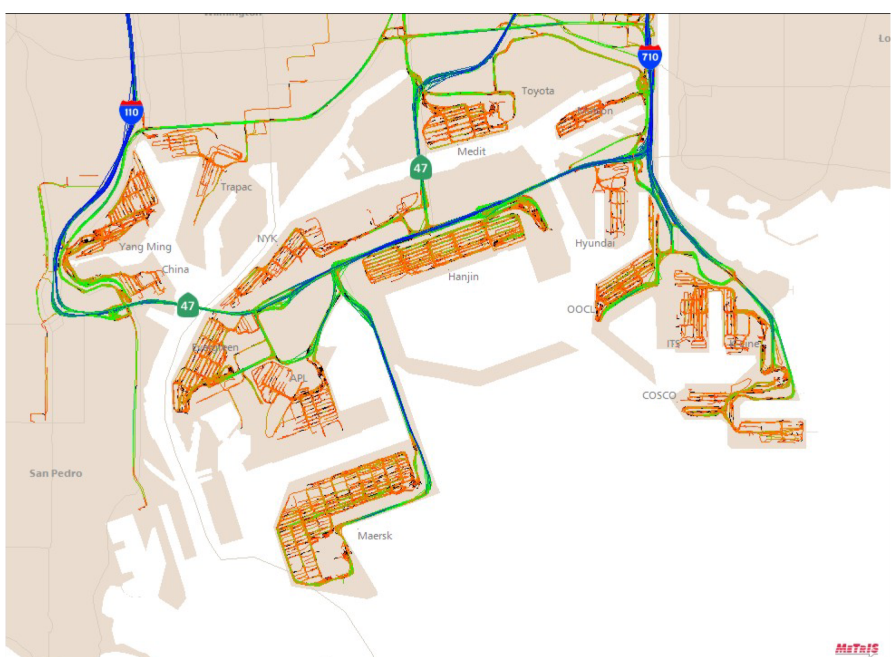

Figure 7. Port layout and circulation as revealed by MeTrIS traces

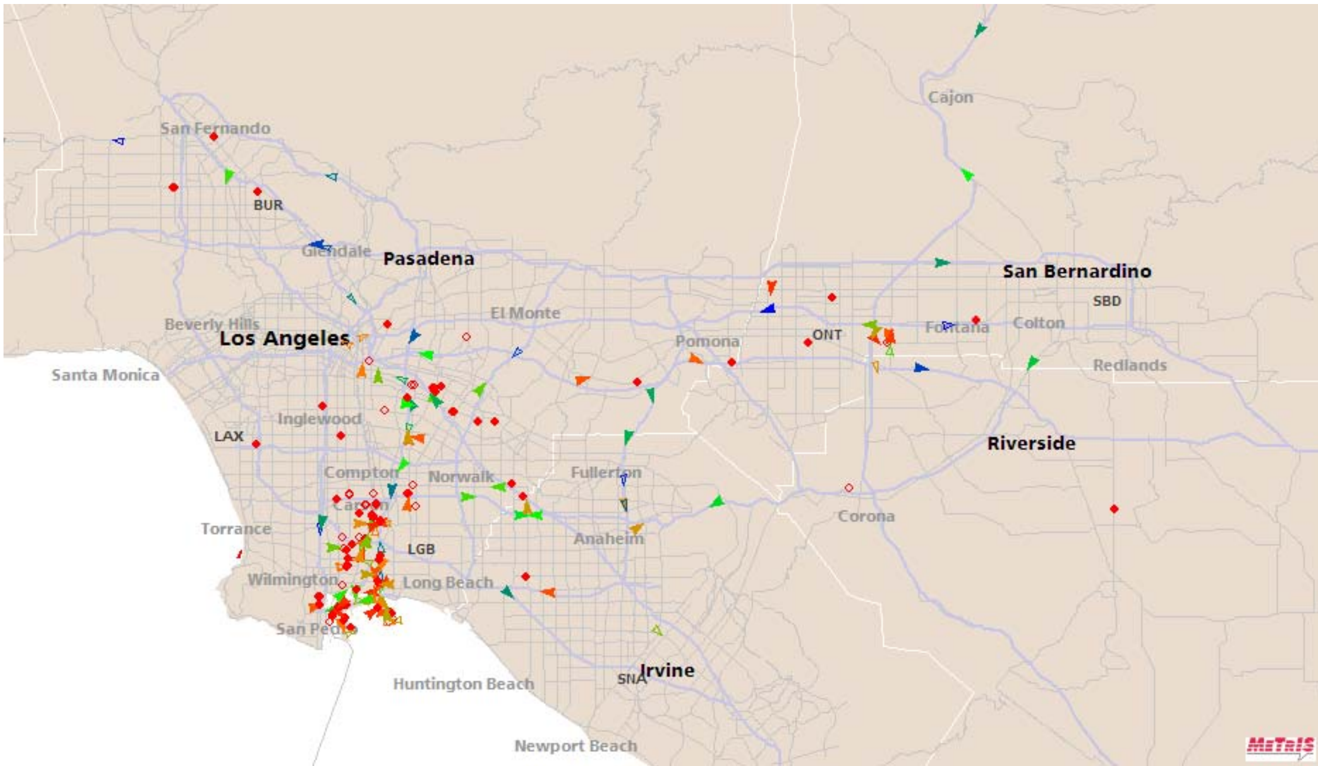

Figure 8. Real-time map of truck movements shows reach and density of port freight activity

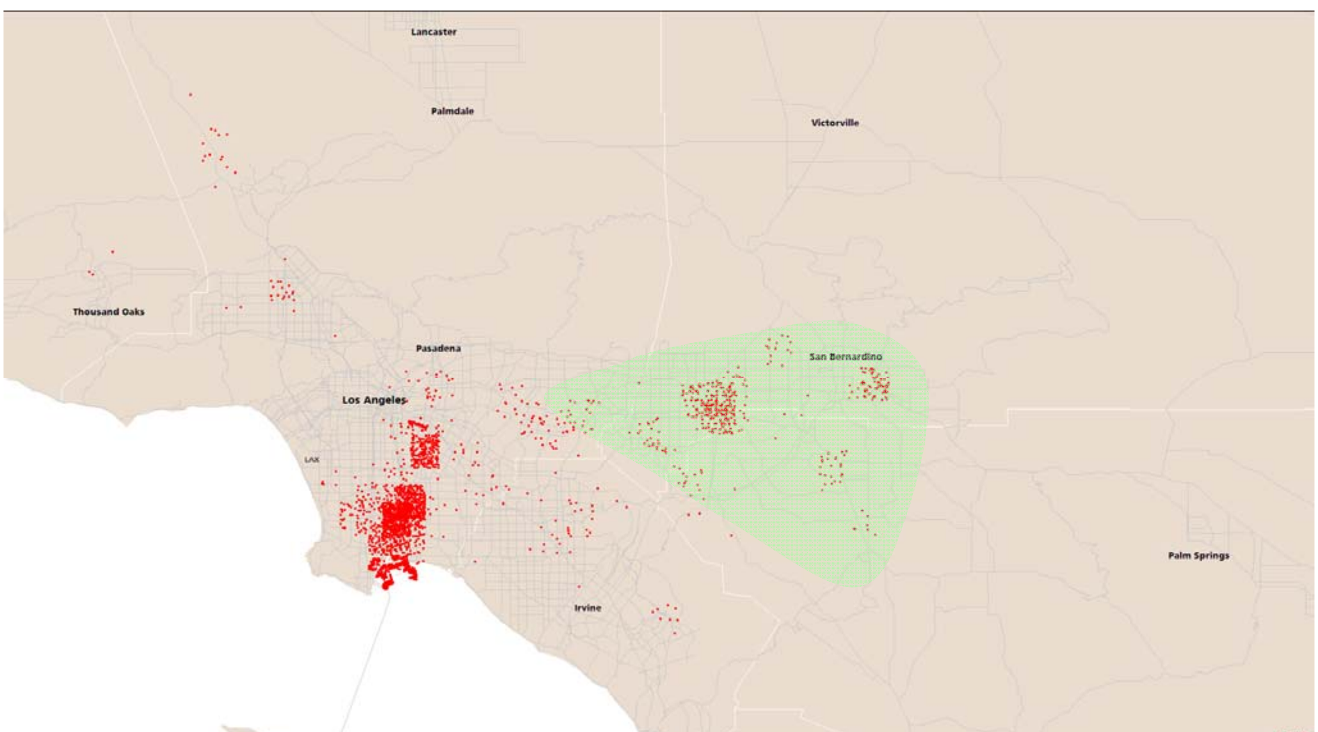

Figure 9. Origins and destinations, with some locations perturbed to protect confidential data. Based on MeTrIS data, the Inland Empire import/export warehousing district appears in green.

Origins and Destinations

Analysis of vehicle speed produces tables and maps of common origins and destinations (Figure 9). Because there is currently no documentation of trip purpose in the track data, destinations may be port terminals, rail terminals or warehouses, or service points such as fuel stations, restaurants or maintenance yards. Since the latter class of stops are likely to be en route to goods destinations, it can generally be inferred that points furthest from the ports are true destinations, and this can be confirmed by examination of maps and aerial imagery.

Such destination maps effectively reveal the functional geographic structure of the basin with respect to freight. Motor carrier yards tend to congregate in the Carson area, the Hobart rail yards account for trips to central Los Angeles, while import and export warehouses tend to lie in the eastern regions of the basin, commonly known as the Inland Empire.

Travel time

Analysis of travel time along major arteries typically shows a peak at 3 pm, and a secondary peak at about 9 pm as night port operations peak. MeTrIS GPS analysis can support complex travel time queries, such as comparison of two or more routes from an origin to an intermediate or final destination. Figure 10 shows two routes, indicated by red and blue stippled lines, from the intersection of SR-91 and I-710 to the intersection of SR-57 and SR-60. More trucks use the blue route traveling eastbound, yet the red route tends to be slightly faster. Drivers explain that the red route has steeper gradients, better suited to empty containers, which are hauled at higher speeds.

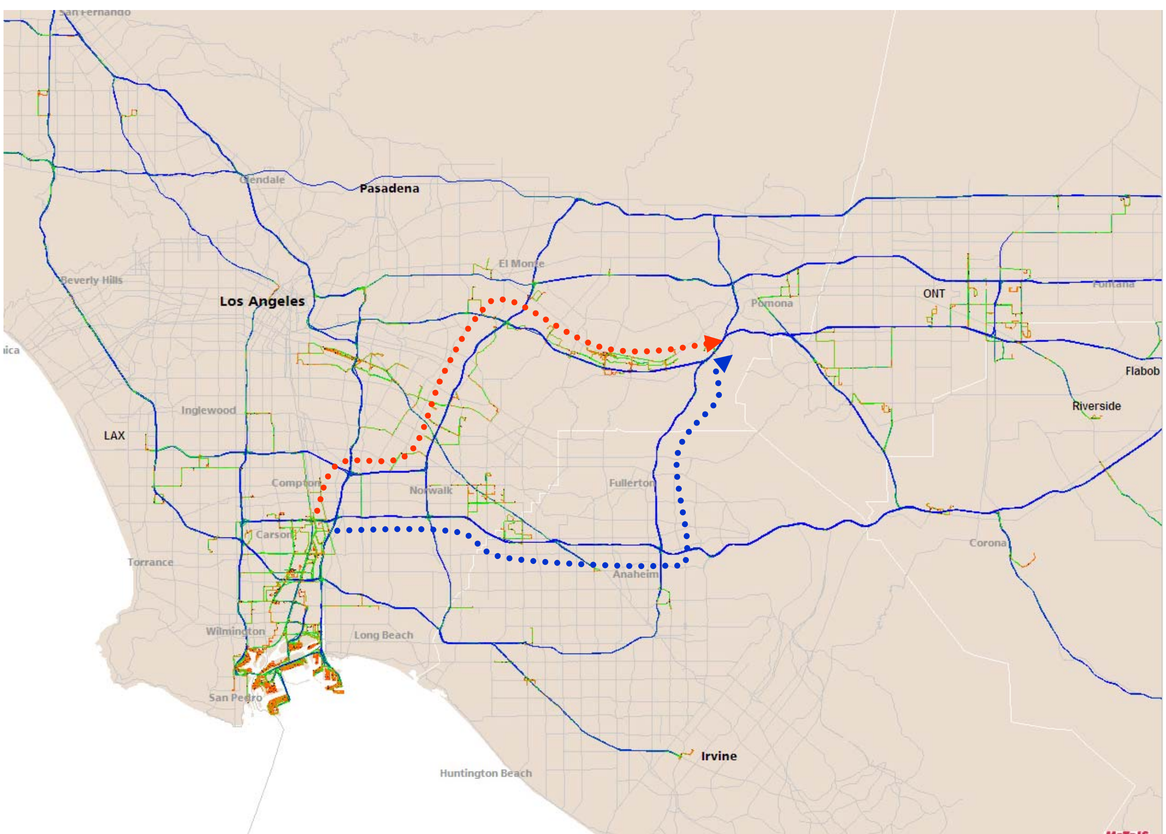

Figure 10. Two routes to the same destination

Terminal Queues and Turn Time

By far the greatest concern of the freight industry in the San Pedro ports, whether truckers, marine terminals, labor unions, or beneficial cargo owners, is the length of time taken to pick up a load at a marine container terminal. Motor carriers generally refer to a "turn" as a trip from the carrier's yard to the port. On that trip, a truck would typically drop off an empty container at one terminal and pick up a full import container at the same or another terminal. Sometimes the truck may have to drop off or pick up a chassis.

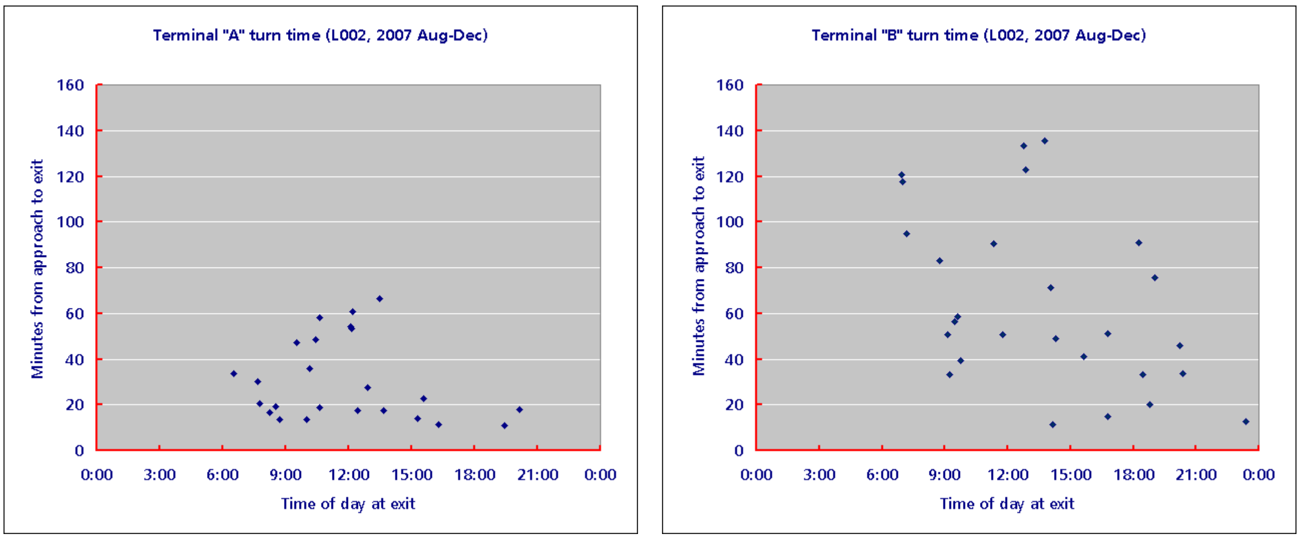

Figure 11. Turn time at two terminals, measured in 2007

From an analytical standpoint, a turn can be defined and turn time can be measured, but the value of the measurement is dubious, because turns measure neither revenue nor efficiency. A more useful set of metrics is the time spent at a marine terminal, including the wait queuing outside the gate, the admission process, transaction time delivering or picking up the container, queuing for exit, and egression. With knowledge of terminal layout, and adequate sampling of vehicle trajectory, these are easily measured by GPS.

Over more than a decade, motor carriers and marine terminals have disagreed over measurement of turn time. Motor carriers rely on accounts of truck drivers, who are apt to recall the worst cases. Terminals time-stamp a transaction at the entrance and exit gates, and hence do not include queue time outside the entrance. MeTrIS has provided the first objective measurement of turn times. Figure 11 shows early observations of turn time at two marine terminals, identified as Terminals "A" and "B," based on data from one truck. At Terminal A, the entire visit, from the first point of queuing to exit, is never longer than 70 minutes; however, at Terminal B, transactions are considerably longer, with a maximum of nearly 140 minutes. This is one important aspect of the problem: durations vary widely among terminals. They differ in corporate ownership and structure, land ownership and physical practices. For example, the APL terminal has a dedicated exit off SR-47, and any traffic along that route is clearly and exclusively APL-bound. Queuing for the Evergreen terminal, on the other hand, is on a city street, accessible to other vehicles. When queues back up, trucks are diverted to a holding area at some distance from the terminal.

Objective measurement of turn time, particularly queues, has been a unique contribution of MeTrIS. However, release of findings has been highly controversial, and we have faced implied threats of legal action from representatives of marine terminals, questioning our right to embed devices in trucks on terminal property without prior authorization.

At the time of writing this report, the legal dimension of this controversy has subsided, and the terminals, motor carriers and ports are working cooperatively towards objective measures of terminal waits. However, the issue illuminates an area of concern in the rapidly advancing field of location monitoring and Location Based Services. Location privacy is typically considered from the standpoint of the subject bring tracked, and Institutional Review Board (IRB) mechanisms are in place to enforce this. For example, prior to instrumenting vehicles, we were required to obtain signed consent from drivers and/or motor carriers. However, there is no requirement to obtain consent from privately-owned or privately-leased places of business visited by the trucks, although the tracking devices gather data that can be processed into comparative metrics of performance.

Reports to Motor Carriers

As an incentive to motor carriers to participate in this research, DGRC provided real-time reports on trucks, and daily summaries of their encounters with landmarks of interest: freeway intersections and port features. Due to the dense sampling of MeTrIS data, these reports provided a high quality trace of truck movements, and for motor carriers never previously exposed to tracking, it was a free preview of the benefits of telematics-based fleet management. In the perpetual back-and-forth among carriers and terminals regarding queue delays, carriers were able to challenge terminals' accounts of queue duration. In one instance reported to us, a terminal was so convinced by the MeTrIS data record that it accepted the carrier's complaint and greatly expanded its land bridge, reducing truck visit time considerably.

Conclusions

This chapter has described relatively straightforward analytical processes using standard mapping techniques and GIS functionality. Yet there has been great receptivity in the freight industry for the findings. This reflects the three words of advice received from the Steering Committee, which pointed out that the greatest value of this project lay in "data, data, data."

MeTrIS has produced the first objective data on freight activities, and in doing so it lays a foundation for a broader, long-term focus on performance metrics, improvement and even algorithmic optimization. Coinciding as it does with the Clean Trucks Program, that has changed the tenor of drayage trucking in San Pedro, MeTrIS has been welcomed by the industry, from individual truck drivers to motor carriers and marine terminals.

There are important privacy aspects of MeTrIS: there are individual and corporate behaviors and even national strategic interests reflected in the data, and protection of those interests demands that access to the data be carefully managed.

Given that MeTrIS has made invaluable contributions to the understanding of freight movement over a period of 3 years, we recommend, unsurprisingly, that it continue to do so, providing an uninterrupted longitudinal stream of data to inform operations, planning and policy.

Introduction

The port complex of Los Angeles and Long Beach (referred to here as the Port Complex (PC)) provides an intermodal interface between marine and land based container shipping. The majority of incoming foreign-sourced containers that are handled by this complex exit the terminal on a trailer/chassis being towed by a drayage truck. The remainder of the incoming containers leaves the PC on rail, destined for delivery outside the Southern California region. The port complex handles on an average day approximately 23,000 inbound TEUs (twenty foot equivalent units). This translates to 11,500 containers (mostly 40 ft units) being taken per day from the ports, 9,400 by truck and 2,100 containers per day handled by on dock rail.

Although there is the expectation that on-dock rail will play a larger role in container movement in the future, and help to reduce the proportion of containers handled by truck, the total volume of containers to be handled by the port is forecasted to grow at a large enough rate that the amount handled by trucks will continue to increase. There are many geographical reasons for why this will occur. The biggest reason for this is that many of the local destinations are not connected to rail. There are, however, a host of other issues that ensure that trucks will always play a major role in port operations. First, many of the consignees are too dispersed to be efficiently served by rail. Second, recent analyses demonstrate that loading containers on rail at the dock and then transporting them to an inland LA regional intermodal facility for final delivery by truck does not pencil out economically (Tioga, 2008). Finally, the LA region serves as a major inventory processing point, where marine containers are unloaded and the contents are sorted and repacked onto trucks for shipment to store distribution centers. The bottom line is that drayage trucks will remain the major delivery vehicle for containers destined to southern California, once the ship is unloaded for the foreseeable future.

The bulk of marine containers taken from the Port Complex by truck are delivered to a local consignee. Figure 12 depicts several possible alternatives in the manner in which a container headed to a local consignee may travel by truck. In Figure 12, the movement of a full container is represented by a solid line and the movement of an empty container is represented by a dashed line. Figure 12a, depicts the case where a full container is offloaded from a ship and then transported by truck to a consignee. After the container is emptied, it is taken back to the port for short term storage. Note that Figure 12a depicts the case where the empty container is subsequently picked up at the port and taken to a local shipper, where it is filled for export. Then, the now-filled container is taken by truck to the port for export. Figure 12b depicts a different trajectory for the local container after it is emptied at the consignee. Instead of being hauled back to the port, as in Figure 12a, the empty container is transported by truck directly to a shipper, and then after it is filled for export it is then taken by truck to the port. Figure 12b represents a better alternative than the transport pattern of Figure 12a, as the total amount of distance involved in transporting the container may be considerably less, not to mention the reduction in traffic coming into and out of the port complex. This second alternative is called a “street turn.” It takes special coordination of information to accomplish a street turn. To support this alternative, it is necessary to know when a container has been emptied at a consignee and is available for transporting to a shipper, whether it is the right size and type, and whether it will involve the same terminal operator (inbound & outbound). This type of information is supported by web-based software called a "virtual container yard" (VCY) (see e.g. Theofanis and Boile, 2007). Although there are a number of merits for using a VCY, there are also barriers to its use. (1) If a container is not returned to the marine terminal within a certain time window, per diem charges accrue. (2) Truckers are concerned with losing some part of the transport fee, as they have delivered the container, but not necessarily picked it up for the street turn. (3) The VCY lacks structure and predictability. A driver has to visit different addresses each day to pick up empties; a container may not be available as advertised, or there could be roadability or cleanliness issues. Whatever the case, street turns, although efficient are used only about 10% of the time in the LA basin.

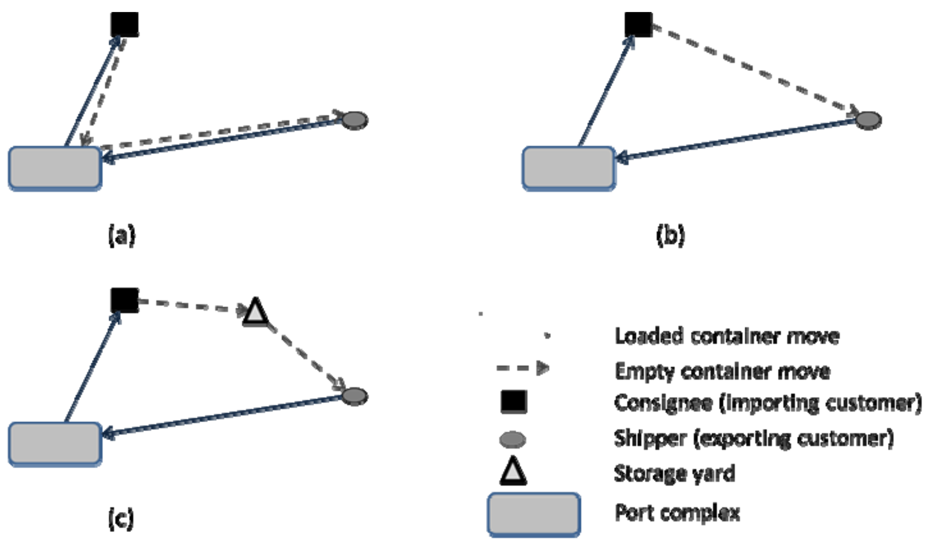

Figure 12. Truck delivery and pickup patterns for marine shipping containers

The third option in handling an import container is depicted in Figure 12c. Here, the import container is trucked from the port to a local consignee. There it is emptied and then transported to a short term empties storage yard (ESY) that is not at the port complex. Should the ESY be located somewhere between consignees and shippers, then the distance of travel for the empty will be less than what would be needed for the case depicted in Figure 12a. Although there would be additional handling of the container (dropping off and picking up at the ESY) as compared to a street turn alternative, the handling is no different than the alternative depicted in Figure 12a. The benefits of maintaining an ESY away from the port are based on reducing the traffic into and out of the port, as well as potentially reducing the distances that are involved in trucking the container. Although the options depicted in Figure 12 involve the case where a container is used for both import and export, it should be mentioned here that some containers are used for import only, returned from the consignee empty and then placed on a ship for global repositioning. So, in that case, the handling options are fixed and movement cannot be reduced.

There is another element that can be observed in Figure 12, when focusing on full container moves. Essentially, between the various options of handling containers in Figure 12a, b, and c the moves of full containers remains the same. That is, the move made in trucking the full container from the port to the consignee and the full container move from the shipper to the port is fixed. Because of this property, past work has concentrated on optimizing the repositioning of empty containers, rather than all of the trip segments that a container makes. There are two reasons why this can be done, without loss of overall accuracy. The first reason is that the truck that delivers the full container to the consignee usually drops it off and then departs, as it takes time to empty the container. The second reason is that the trucker that delivers an empty container to a shipper quite often does not wait for that container to be filled. Although one may exchange an empty for a full at a shipper, this does not change the basis under which the empty container has been handled. Thus, from a strategic level we can separate the analysis of empty container management from the overall flows of both empty and full containers. This perspective changes when moving from a strategic level of analysis to that of a daily operational role, where it would be important to coordinate all transports, empty or full, so that we can utilize trucks as efficiently as possible. This basic perspective is not new as the research literature on container management can be easily divided into the strategic and operational roles. Examples of strategic models can be found in Crainic et al. (1989); Crainic et al. (1993a); Crainic et al. (1993b); Gendron and Crainic (1995); Xu (1999); Boile et al. (2008); Mittal (2008) and a review in Lei and Church (2010). Examples of operational models can be found in Namboothiri and Erera (2004); Jula et al. (2005); and Imai et al. (2007). The work here concentrated on improving the efficiency in handling of empties from a strategic perspective.

Now that we have a basic understanding of how marine containers are handled outside the ports in Southern California (and similarly elsewhere), how big is this problem and is there something that can be done to improve the efficiency of the trucking operations? Figure 13 depicts the Los Angeles region in terms of the distribution of shippers and consignees, aggregated in terms of the approximately 4,000 Transportation Analysis Zones (TAZs). If a TAZ contains a shipper or consignee it is mapped as a circle or square respectively. From a recent survey (Tioga, 2008) we have derived an estimate of daily flow totals of full containers entering and leaving specific TAZs. Since the bulk of empties are transported back to the port before being transported to a consignee, we can use this information to estimate the total amount of travel involved in the repositioning of empty containers. Overall, on the average day approximately 8,400 empty containers are drayed to the port from consignees and 3,500 empties are drayed from the port to shippers, where the remainder (4,900 empties) is loaded onto ships for global repositioning.

To make a realistic estimate of empty transport, we also need to have some idea as to which routes would be traveled by truck in delivering or picking up containers. This can be done by analyzing the truck tracking data generated by this project and discussed in section 2 (destinations and routes). Overall, we found that trucks were routed between origins and destinations, principally along highway segments, and using routes which minimized estimated travel time based upon data provided by TeleAtlas. For the region, we have 206 origin (or consignees) TAZs for sources of empty containers and 208 destinations (or shippers) TAZs for demands for empty containers. We generated using TeleAtlas data the travel times between all consignee TAZs and the port complex, and the travel times between the port and all shipper TAZs. Weighting these travel times/distances by the number of trips made between the port and these TAZs allows us to estimate the total amount of travel distance (or time) involved in repositioning empties in the LA basin currently. This estimate was 473,417 kilometers of travel per day in moving empty containers. This does not include any additional inefficiencies involved in coordinating truck trip segments (which could also be a staggering amount). This movement involves the consumption of 32,900 gallons of fuel (9 mpg), the production of more than 1,873,000 lbs of carbon dioxide, and handling more than 3,500 empty containers every day at the port that are that are both checked in and subsequently checked out on a subsequent day. This is a staggering amount of resources (trucks, fuel, air pollution emissions, etc.) devoted to the repositioning of empties. One of the major objectives of this project was to find a cost effective way in which this could be substantially reduced.

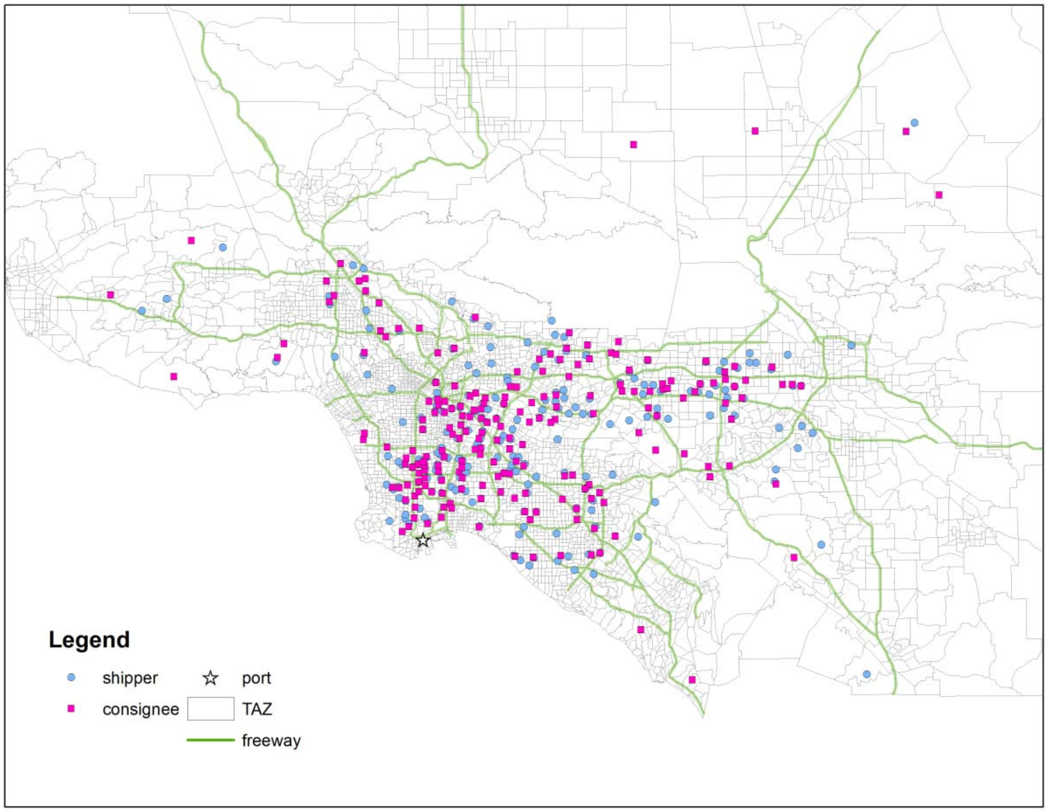

Figure 13. Distribution of shippers and consignees by transportation analysis zones in the LA region

Model Specification

As described in previous sections, the main issue of improving the container movement at a strategic level is to improve the operations of repositioning empty containers. Full container moves can be considered fixed except for operational level issues. At this level of analysis, empty container movements can be viewed as a problem of locating a set of storage facilities to facilitate flows of empty containers through the local transportation network. Consignee sites are locations where full containers are unloaded and empty containers are generated. Therefore, they can be viewed as a geographically dispersed set of sources of empty containers. Similarly, shipper sites are locations where empty containers are needed and they are sinks of empty containers. Storage facilities (including the port complex) can be viewed as intermediate locations that empty flows from sources to sinks must transition through, unless street turns are used. In the big picture, the goal of the empty storage problem is then to find good storage locations so that the total system cost of flow for the empty repositioning activities can be minimized. To implement this vision, we developed three strategic level models (we describe two here) using Integer Linear Programming to analyze the efficacy of maintaining such storage yards away from the port. The first model is based upon determining a system-optimal solution. The second model is defined to identify solutions that tend to be oriented towards user-optimality or based upon where drayage drivers might pick up or drop off empty containers. Thus, model solutions can be compared in terms of efficiency and possible operating policies. We begin with the formulation of the system optimal model.

The first model is based upon the assumption that there is one firm which makes all container transports. We know that this assumption does not hold for LA or for other ports for that matter, but it does represent the case where we wish to identify the absolute minimum movement possible when establishing storage yards. This assumption has significant ramifications in the operation of any storage yards away from the port. If a storage yard is used as a drop off location more frequently than it is as an empty pickup location, then an surplus of empties will tend to grow at that location. Then, ultimately the growing surplus of empties will need to be repositioned to either the port for global repositioning or another storage yard that requires more empties to satisfy its demand. In general, repositioning between storage yards is more expensive than taking the container directly to the storage yard where it will be needed. This is due to the fact that there are no economies of scale in the local repositioning of empties when using trucks. There are, however, options where the repositioning of empties may involve some economies of scale, e.g. using a special dedicated system, like rail or maglev. The system objective of the first model is to locate one or more away-from-port empty container storage yards so as to minimize the total distance (or travel time) involved in moving empty containers between consignees, yards, and shippers, while keeping empty containers entering/exiting the port complex to a minimum. To formulate the first model, we need the following notation:

i = an index used to refer to consignees, where $i = 1, 2, 3, \ldots, n$

j = an index used to refer to shippers, where $j = 1, 2, 3, \ldots, m$

k = an index used to refer to storage yard sites, where $k = 1, 2, 3, \ldots, K$.

d_{ik} = the distance or cost to transport a container from consignee i to location k.

d_{kj} = the distance or cost to transport a container from location k to shipper j

d_{tk} = the distance or cost to transport a container from location t to location k

S_i = the average daily volume of containers emptied at and transported from consignee i

D_j = the average daily demand for empties at shipper location j

D_{port} = the demand for empties at the port for purposes of global repositioning

a_i = is the fraction of empties at consignee i that is transported directly to a shipper (i.e. street turns to shippers).

β_j = the fraction of demand at shipper j that is met by direct transport from consignees (i.e. street turns from consignees).

p = the number of storage yards to be located

We also need to define the following decision variables:

$$y_k = \begin{cases} 1, & \text{if a storage yard is located at site } k \\ 0, & \text{otherwise} \end{cases}$$

x_{ik} = the fraction of the daily volume of empty containers generated at consignee i that are transported to storage yard k

e_{tk} = the average daily volume of empty containers repositioned from storage yard t to storage yard k.

For convention we will index the port as storage site $k=1$. The first empty storage location model (OSEM1) can then be formulated as:

$$\min Z = \sum_{t=1}^{K}\sum_{k}^{K} d_{tk}e_{tk} + \sum_{i=1}^{n}\sum_{k=1}^{K} d_{ik}(1-\alpha_i)S_i x_{ik} + \sum_{k=1}^{K}\sum_{j=1}^{m} d_{kj}(1-\beta_j)D_j x_{kj} \qquad (1)$$

subject to

$$\sum_{k=1}^{K} x_{ik} = 1, \quad \text{for each } i = 1, 2, 3, \ldots, n \qquad (2)$$

$$\sum_{k=1}^{K} x_{kj} = 1, \quad \text{for each } j = 1, 2, 3, \ldots, m \qquad (3)$$

$$x_{ik} : y_k, \quad \text{for each } i = 1, 2, 3, \ldots, n \text{ and } k = 1, 2, 3, \ldots, K \qquad (4)$$

$$x_{kj} : y_k, \quad \text{for each } j = 1, 2, 3, \ldots, m \text{ and } k = 1, 2, 3, \ldots, K \qquad (5)$$

$$\sum_{t \neq k} e_{tk} + \sum_{i=1}^{n}(1-\alpha_i)S_i x_{ik} - \sum_{t \neq k} e_{kt} - \sum_{j=1}^{m}(1-\beta_j)D_j x_{kj} = \begin{cases} D_{port}, & \text{if } k=1 \\ 0, & \text{otherwise} \end{cases}, \quad \forall k \qquad (6)$$

$$\sum_{t \in K} e_{tk} + e_{kt} \leq M \cdot y_k, \quad k = 1, 2, 3, \ldots, K \qquad (7)$$

$$\sum_{k=1}^{K} y_k = p \qquad (8)$$

$$y_1 = 1 \qquad (9)$$

$$x_{ik} \geq 0, \; x_{kj} \geq 0, \; e_{tk} \geq 0, \; y_k = 0,1 \quad \text{for appropriate values of } k, t, i, \text{and } j \qquad (10)$$

The objective function (1) minimizes the total mileage (or time) involved in moving empties between consignees and storage yards, between storage yards and shippers, and between storage yards for inventory rebalancing. This model is based upon a given level of "street turns." The level of street turns in the model is represented by the parameters of α_i and β_j. If $\alpha_i = 0$ for each i and $\beta_j = 0$ for each j, then no street turns are assumed, and the model meets all demand by transporting empties from short term storage. If the values of α_i and β_j are greater than zero, then demands and supplies of empties are reduced by the assumed level of street turns. Note that the values of α_i and β_j may vary between locations but must be such that the total reduced demand for containers equals the total reduced supply of containers. It is important to note that the objective does not include the distances involved in making street turns, as it is assumed that these are exogenous to the model. Constraint (2) maintains that all empty containers at a consignee (minus the ones that are devoted to street turns are transported to one or more storage yards (including the port). Constraint (3) ensures that the demand for empties at a shipper is met from empties transported from storage yards (after accounting for the demand that is supplied by street turns). Constraint (4) requires that empty containers generated at a consignee can only be shipped to storage facilities that have been selected. Similarly, constraint (5) requires that empty containers used by a shipper can only be supplied from located storage yards. Constraint (6) is a "flow conservation" constraint, which requires that in the long term, the empty container flows into and out of each potential storage site should be balanced except for the port which absorbs the imbalance of empty containers in the entire region (empties for global repositioning). Constraint (7) requires that any rebalancing flows must be between open storage yards. Constraints (4), (5) and (7) ensure that no containers are handled at sites that have not been selected for storage yards. When a site has not been selected for a storage yard, then the conservation of flow constraint (6) will automatically be met as no containers can be taken to that site or picked up at that site, ensuring the inventory will always be zero for that site. For constraint (7) the value of M must be a sufficiently large number to represent any rebalanced flows that are associated with a given storage site when it has been selected for a yard. Finally, constraint (8) establishes that exactly p storage yards will be located and constraint (9) maintains that the port is always chosen as one of the storage yards for empties. Constraint (10) defines the range and the integrality conditions for each decision variable.

It should be pointed out that even though the above model allows rebalancing between storage yards, an optimal solution will never allow an inventory imbalance between what is delivered to a storage yard and what is to be supplied from that storage yard. Instead, it will maintain inventory balances by optimizing empty assignments from consignees to yards and empty containers to shippers from yards. This is a property that has been discussed in the literature and is due to the fact that there are no economies of scale in rebalancing between yards when using a truck. Consequently, it is possible to eliminate all of the and variables from the problem as well as constraint (7) from OSEM1 without any loss of generality. The main reason why we have included rebalancing in the model is that the second model discussed here will require this option.

The first model (OSEM1) is a system-optimal model. It is quite similar to the model of previous work such as Crainic et al. (1989) in that it minimizes overall system transport cost and allows balancing flows between depots. It differs from that of Crainic et al. (1989) in that street turns are represented explicitly. Costs of storage could be easily included in the objective function of OSEM1, but our main objective was to identify the level of empty transport mileage based upon the number of located storage yards. One of the reasons for doing this is that transport mileage is a major determinant of other concerns, like air quality and carbon emissions, and that costs for these factors were not available.

The second model, OSEM2, is based upon the assumption that there are multiple trucking companies and independent truck drivers who carry out container transports. This model represents a distinct departure from past research on container operations, in that past work assumes a central, one-company perspective. We believe that past work is not very applicable for those cases that involve many companies as well as a host of independent contractor operators. We assume here that the drayage driver/dispatcher is quite greedy in the use of any storage yards. Assume for the moment that a drayage driver has picked up an empty at a consignee and needs to drop it off somewhere. Here we assume that the driver will take it to the nearest established yard and dump it so to speak as quickly as possible. After all, this strategy will involve minimizing the distance that an empty is hauled by the driver, a sort of "user optimal" strategy. Additionally, when a driver is headed to supplying a shipper with an empty container we will assume that the driver will pick it up at the storage yard that is closest to the shipper's location. This will in effect help to keep trip mileage to a minimum on the part of the drayage truck, but then may contribute to an imbalance in container inventories that will need to be rebalanced. We can add this type of behavior to the model by forcing container drop off assignments to the absolute closest yard from the consignee or by forcing the pickup of an empty from the storage yard that is the closest to the intended shipper location. To add this drop off and pick up behavior to the model we need to introduce the following additional notation:

$$C_{ik} = \{ \text{ sites } q \mid d_{iq} < d_{ik} \text{ or } d_{iq} = d_{ik} \text{ and } q < k \}$$

$$C_{kj} = \{ \text{sites } q \mid d_{qj} < d_{kj} \text{ or } d_{qj} = d_{kj} \text{ and } q < k \}$$

The set C_{ik} represents the set of storage sites that are closer than site k is to consignee i, and the set C_{kj} represents the set of storage sites that are closer than site k to shipper j. We also need to change the definitions of the x_{ik} and x_{kj} variables to represent the case that the assignment is discrete and to a specific yard based upon whatever is closest. Consider, then:

$$x_{ik} = \begin{cases} 1, & \text{if storage yard } k \text{ is the closest located yard to consignee } i \\ 0, & \text{otherwise} \end{cases}$$

$$x_{kj} = \begin{cases} 1, & \text{if storage yard } k \text{ is the closest located yard to shipper } j \\ 0, & \text{otherwise} \end{cases}$$

The change in the above transport variables is based upon making "greedy-like" assignments in terms of dropping off empties at the closet yard or picking up empties at the closest yard. Given the previous notation in OSEM1, the added set notation, and the new decision variables we can formulate OSEM2 as follows:

$$\min Z = \sum_{t=1}^{K} \sum_{k}^{K} d_{tk} e_{tk} + \sum_{i=1}^{n} \sum_{k=1}^{K} d_{ik} (1-\alpha_i) S_i x_{ik} + \sum_{k=1}^{K} \sum_{j=1}^{m} d_{kj} (1-\beta_j) D_j x_{kj} \tag{11}$$

subject to (2) through (9),

$$\sum_{q \in C_{ik}} x_{iq} + x_{ik} \geq y_k, \quad \text{for each } i = 1, 2, 3, \ldots, n \text{ and } k = 1, 2, 3, \ldots, K \tag{12}$$

$$\sum_{q \in C_{kj}} x_{qj} + x_{kj} \geq y_k, \quad \text{for each } j = 1, 2, 3, \ldots, m \text{ and } k = 1, 2, 3, \ldots, K \tag{13}$$

$$x_{ik} = \{0,1\},\ x_{kj} \in \{0,1\},\ e_{tk} \geq 0,\ \ y_k = 0,1 \quad \text{for appropriate values of } k, t, i, \text{and } j \tag{14}$$

The formulation above is quite similar in concept to OSEM1, except that transport assignments between consignees and yards and transport assignments between yards and shippers are discrete (zero-one) and must be made to the closest located storage yard. This is enforced using constraints (12) and (13), respectively. It should be noted that in (14), even though the assignment variables x_{ik} and x_{kj} are defined as binary variables, it is not necessary to maintain these variables as zero-one integer when the model is solved. This is because the closest assignment constraints will force the integrality of assignment variables when the location variables, y_k are zero-one integer (see Lei and Church, 2010). This is a valuable property when solving this problem using some form of linear programming with branch and bound, as the number of needed integer variables can be kept relatively small.

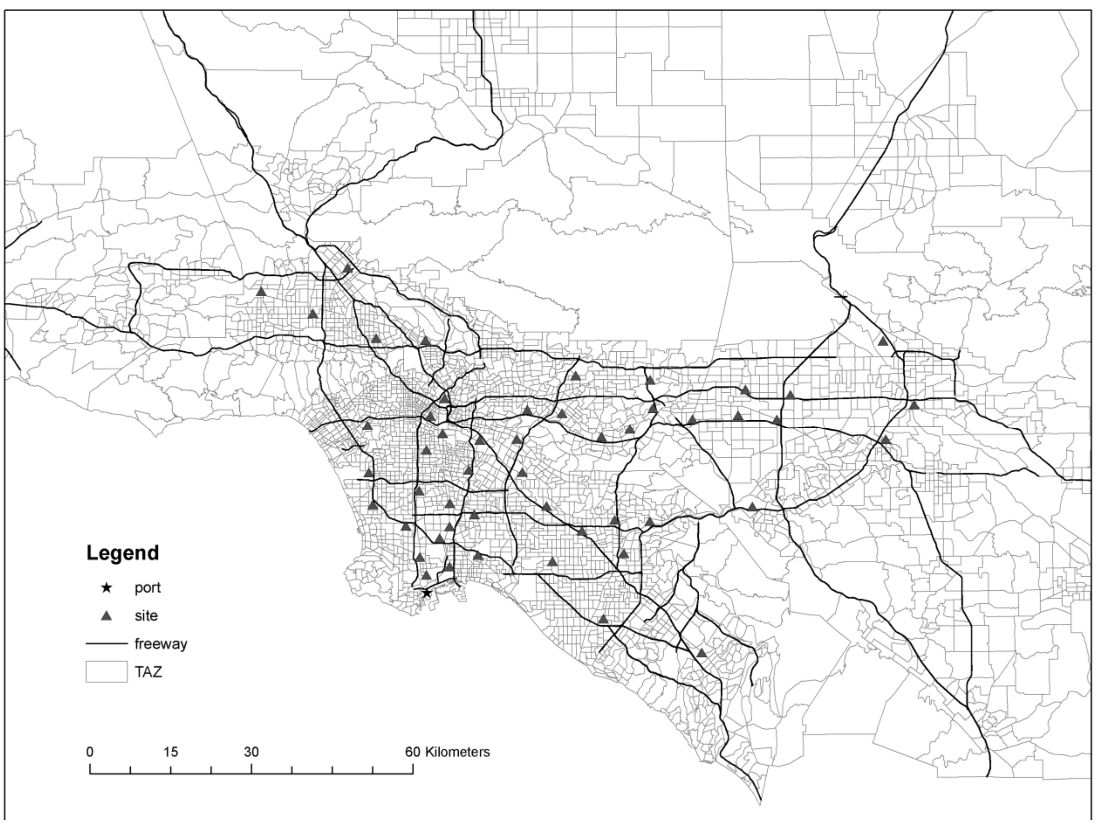

Figure 14. Candidate sites for storage yards

While OSEM1 represents a system optimal perspective, the OSEM2 model represents a user optimal perspective. This second model represents the first time in which an attempt has been made to capture user characteristics in an empty storage model, as virtually all past work assumes a one-company operation. OSEM2 captures the operating nature of the drayage business in the U.S. with respect to port operations, while OSEM1 fails to represent the combined nature of a large number of individual trucking company making independent decisions.

Model Validation/Results

In this section, we present results from our strategic planning models applied to the port complex of Los Angeles and Long Beach and the Los Angeles basin. There are three basic components that are necessary in making a realistic analysis: 1) a road network data base, 2) a truck routing program that realistically determines routes that represent what a drayage driver would take, and 3) volume data on consignee destinations and shipper locations. For road network data we used the TeleAtlas network for the Los Angeles region. As described earlier, a sample of drayage trucks routes based upon GPS tracking data were used to test and verify the accuracy of a truck routing/path model. We compared the actual truck routes to that of a shortest time path model applied to TeleAtlas data and found similar routes and times, if not exact routes and similar times. Differences seemed to occur only when two competitive and equal route segments existed. The volume data on consignees and shippers and their locations were based upon a survey from Tioga Group (2008). Overall the LA region is comprised of 4109 TAZs. Of these, 206 TAZs contained active consignees and 208 TAZs contained active shippers. Each TAZ centroid was digitized and used as the location for consignees or shippers within the TAZ. Candidate locations were chosen from locations that were near highway and highway intersections and where there exists a high concentration of distribution centers based upon a map presented in the Tioga Group (2008) report. In total 50 candidate site locations were identified. Figure 14 shows the locations of these candidate sites represented as triangles. The figure also shows the highway system in thick lines and the TAZs as polygons.

We implemented the two OSEM models using OPL 5.5 and CPLEX 11.0, a sophisticated linear/integer programming system now marketed by IBM. In our analyses, we ran the OSEM1 and OSEM2 models on the LA basin data with street turn levels of 10% and located up to 10 empty storage yards with each model. It should be noted each model maintains an empties storage yard at the port. Thus, any model application associated with locating 1 storage yard (i.e. $p=1$), will be forced to place the storage yard at the port. This solution would then reflect current operations of the Los Angeles basin. For values of p greater than 1, solutions contain an empties storage yard at the port and $p-1$ storage yards away from the port. In solving the model, we minimized travel time involved in moving containers as this reflects a driver perspective. However, for the output we report both travel times and travel distances.

Table 3 shows the results for running OSEM1 on the LA data with $\alpha_i = 0.1$. Each line of the table is associated with solution derived for a specified number of empty storage yards, p. The β_j values were adjusted so that the street turn volume as a fraction (α_i) of consignee generated empties equals the street turn volume as a fraction (β_j) of the shippers demand. For each level of deployment, p, the optimal objective value which represents the smallest total transport travel time of empties possible and is given units of (daily) container minutes. The tables also include total empty container transport distances associated with each solution (in kilometers) as well as empty transport distances broken into different travel segments. The total Empty Kilometers Traveled (EKT) in moving empty containers is given in column 3, followed by the total distances involved in draying empties from consignees, draying empties to shippers, and inter-depot empty rebalancing flows. Finally, the table also gives the (daily) number of empty containers entering and leaving the port, listed as "Port Inflow" and Port Outflow." Figure 15 depicts the trend, in empty transport distances, by the total of trip segments and individual trip segments, as the number of storage sites varies from 1 to 10.

Table 3 OSEM1 model results when all $a_i = 0.10$

p	Travel Time (min)	EKT (km)	Consignee travel (km)	Shipper travel (km)	Inter-depot	Port Inflow	Port Outflow
1	346392	473416	343979	129437	0	9484.2	3189.2
2	243982	324773	246249	78523	0	6427.3	132.29
3	213459	274807	209626	65181	0	6423.5	128.53
4	202255	257281	199798	57483	0	6295	0
5	197756	249083	196702	52380	0	6295	0
6	193824	245844	204602	41242	0	6295	0
7	190461	240120	201755	38364	0	6295	0
8	189092	238970	199664	39305	0	6295	0
9	187898	237655	201093	36562	0	6295	0
10	187052	236289	200136	36152	0	6295	0

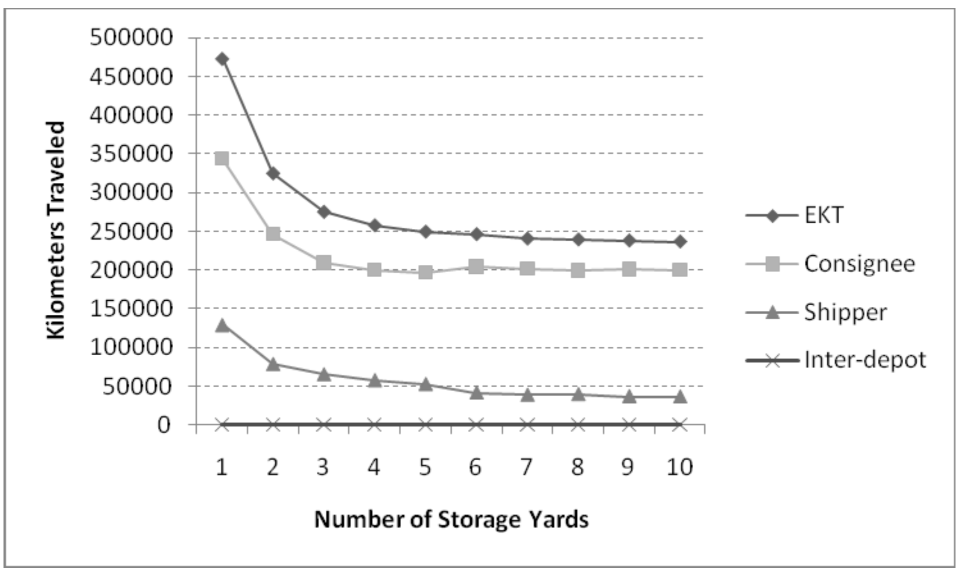

Figure 15. Empty container transport breakdown for OSEM1

Table 3 indicates that the level of empty container travel in the Los Angeles basin on a daily basis equals approximately 474,000 km a day (this does not include the distances traveled in street turns). With one away from port storage yard, EMT can be reduced to nearly 325,000 km per day. This represents a reduction of approximately 150,000 km per day, reducing port entries and exits, as well as traffic in the region immediately adjacent to the port. Consequently, the savings in hauling empty containers can be substantial by maintaining one yard away from the port.

In Table 3, Figure 15, the case when p = 1 corresponds to the base case scenario in which no away from the port depot is located where all empties are stored at the port complex. We can observe that total empty travel (in terms of time and in distance) decreases with increasing values of p values. The marginal value of savings in time and distance tends to decrease as the number of storage yards increases. Past the level of 3 or 4 storage yards (not including the port itself) additional savings are somewhat negligible. Similar trends can be observed for the consignee and shipper components of the travel costs. Note that since this model represents a system optimal perspective, inter-depot flows are zero. Regarding empty traffic at the port, the total volume of empties entering the port is quickly reduced to the extent that empties no longer are hauled away for the port. Empty outflow reaches 0 when 3 yards are used (1 at the port and 2 away from the port). The rest of the Port inflow reflects the local volume that is headed for global repositioning.

Figure 16 shows the optimal 3 storage yard solution presented in Table 3. In the figure, the boxes represent sites selected for depots. Other points represent locations of consignees (small boxes) and shippers (circles). The lines connecting consignees and shippers to selected depots represent the "assignment" of flows. The actual routes are calculated using the shortest (time) route on the road network. From this figure, it seems that the depots are located in "median" positions. We can also observe that certain demands, e.g., some on the north-west corner (in Ventura County, CA) are not assigned to their closest open depot but are assigned to the port instead. This is an example of a system-optimal environment, where a user may be assigned to farther facilities for the sake of system optimality.

Table 4 presents the results for the OSEM2 model. Note that the results for p = 1 are not included, as it represents the base case that is given in Table 3 for p = 1. This model represents a "greedy view" or a user perspective where all empties are dropped off at the closest storage yard from the consignee or picked up at the closest yard to the shipper. For this model there are positive inter-depot flows, and this was expected as the drop off and pickup rules will likely generate inventory imbalances that will need to be corrected by inter-depot transport. Similar to Table 3, the total EKT amounts in Table 4 decrease with an increasing number of established storage yards. Consignee and shipper travel similarly decreases with increasing values of p. The inter-depot flow of empties shows a slight increasing trend with each additional increment of p. Note that port outflow is not reduced to zero with high values of p, as the port appears to serve as a storage facility for shippers near the port. The travel breakdown in terms of the travel segments for empty transport for the OSEM2 solutions given in Table 4 are presented in Figure 17.

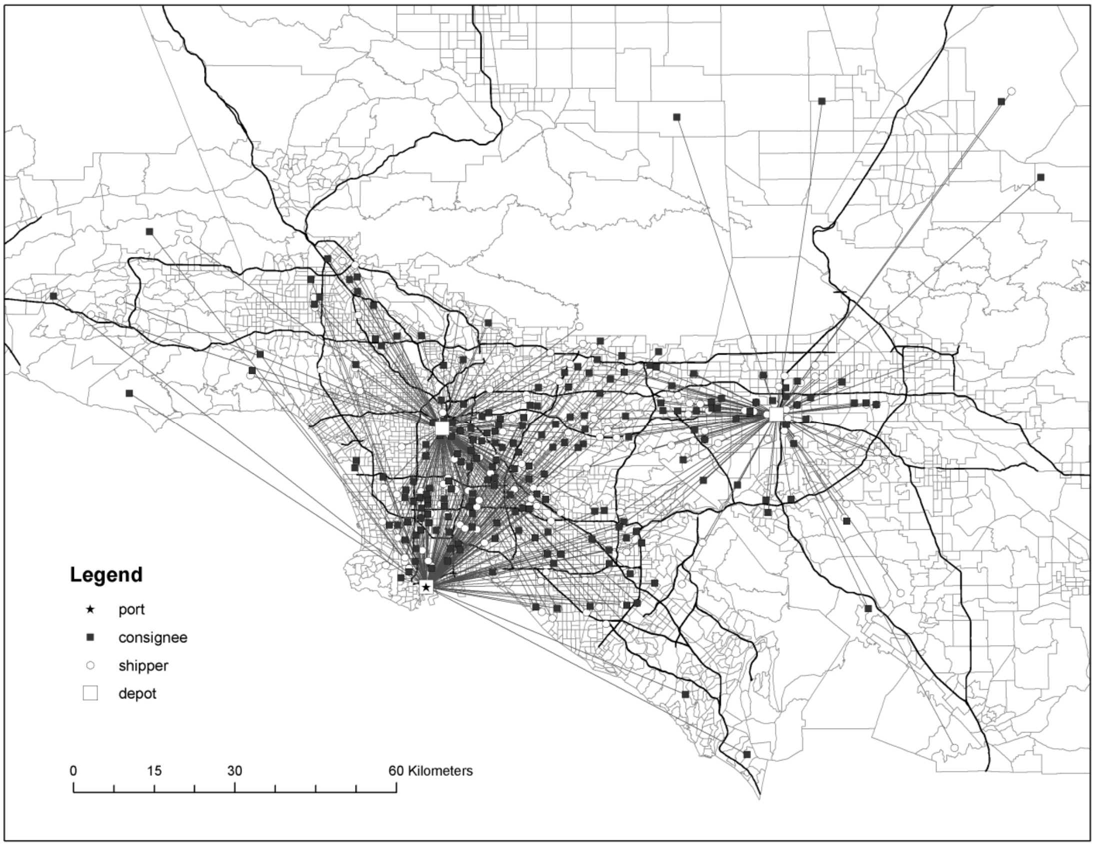

Figure 16. OSEM1 solution for 10% street turn and 3 storage yards

Table 4. OSEM2 model results when all $a_i = 0.10$

p	Travel Time (min)	EKT (km)	Consignee travel	Shipper travel	Inter-depot	Port Inflow	Port Out-flow
2	300715	399094	196009	71284	131799	6820.4	525.39
3	282124	368052	160576	54956	152518	6820.4	525.39
4	264668	334699	132020	47174	155504	6494.2	199.18
5	255771	325655	120448	38418	166788	6494.2	199.18
6	250031	316724	104575	32668	179480	6494.2	199.18
7	244817	301822	92107	26201	183513	6498.7	203.69
8	240056	298464	90905	26432	181126	6494.2	199.18
9	236208	292271	84698	23809	183763	6494.2	199.18
10	233488	289497	83391	23572	182533	6494.2	199.18

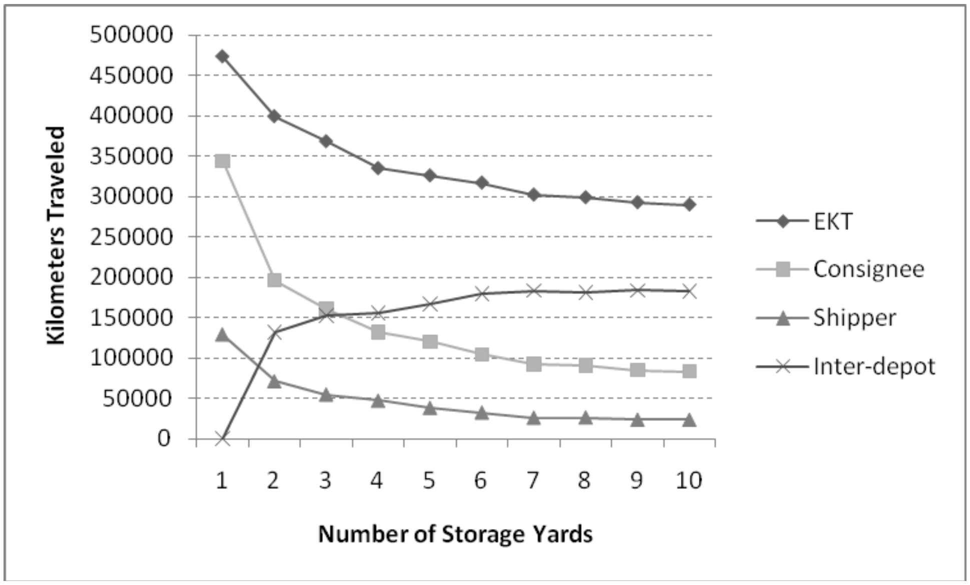

Figure 17. Empty container transport breakdown for OSEM2

Figure 18 shows the OSEM2 solution for locating 3 depots with a street turn rate of 10%. Compared with Figure 16, a notable difference is that customers are assigned to the closest open facilities as this model attempt to represent a user perspective. Consider again, the peripheral customers in the north-west (in Ventura County, CA). These positions are no longer assigned to a yard for inventory balance purposes as in Figure 16, but now assign to their closest empty storage yard. It should also be noted that the locations of the depots are somewhat different from the OSEM1 solution presented in Figure 16. Finally, Figure 19 presents results for both OSEM models in terms of total EKT for moving empties to meet demand across the basin.

There is one other point that is important to make from the analysis of drayage operations. This involves the traffic generated on interstates 110 and 710 south of the 405 in Los Angeles and Long Beach. Table 5 presents results tallied in terms of traffic flows on these two interstates.

The values in this table indicate that one or two storage facilities in addition to the storage already provided by individual terminal operators can help reduce traffic flows by nearly six thousand trips a day. It should also be mentioned that with construction projects planned for these two interstates, any reduction in traffic flow will help to reduce traffic congestion now as well as during construction. It is also important to recognize that the mileage on these two interstates can be reduced by 50,000 km a day when establishing only one away from port storage facility.

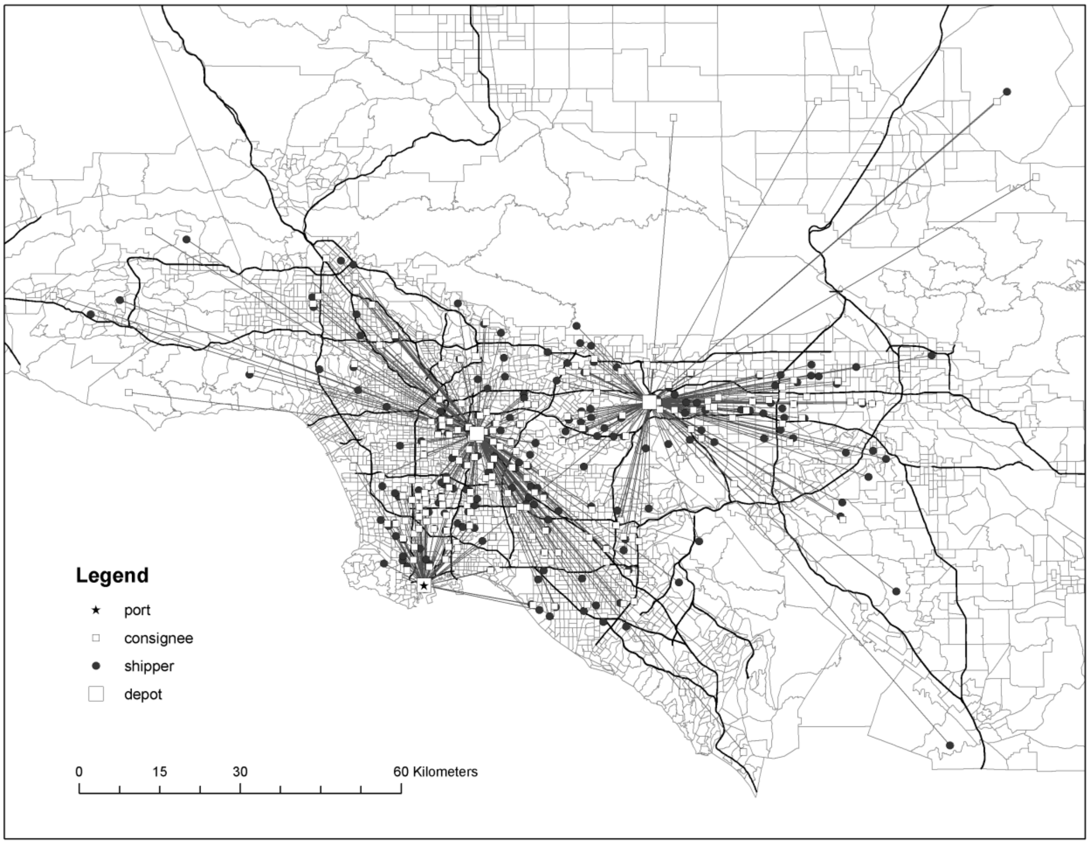

Figure 18. OSEM2 solution for 10% street turn and 3 storage yards

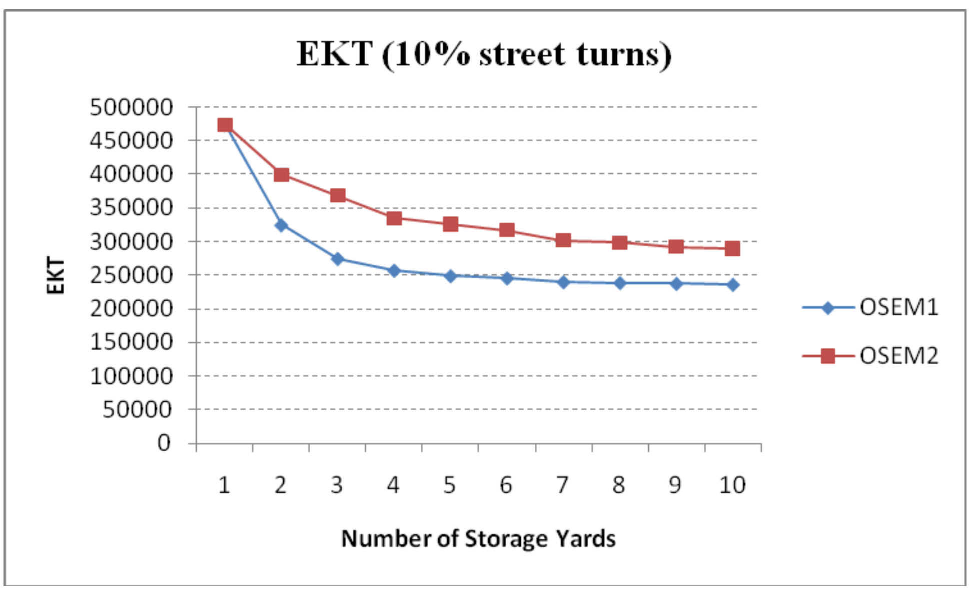

Figure 19. Total empty kilometers traveled (EKT) for different OSEM models

Table 5. Traffic flows along interstates 110 and 710 south of the 405 near the ports

P	Total empties transported on I-710 & I-110 per day	# Empties being hauled from consignees	#Empties being hauled to shipper
1 (port)	12270	9220	3049
2 (port +1)	6363	6164	199
3 (port + 2)	6359	6160	199
4 (port + 3)	6355	6031	324
5 (port + 4)	6355	6031	324
6 (port + 5)	6224	6031	193
7 (port + 6)	6227	6031	196

Overall, the results are quite interesting from a basin-wide perspective. It is clear that the marginal benefits of establishing more than 4 yards diminishes to the extent that it would not be practical to consider such an investment. It is also clear that system optimal values tend to indicate that an additional storage yard could result in the reduction of EKT by nearly 150,000 km per day. This is a substantial amount of truck travel reduction per day in the basin. Second, depending upon how drayage drivers and their companies utilize the storage yards, reductions in travel per day may be closer to 75,000 kilometers per day (OSEM2) or even as low as 40,000 kilometers per day (as determined in this project). The results clearly suggest that it is easy to overestimate the savings in transporting empty containers when solving the problem from a system optimal perspective, given that individual drayage companies may use them to their utmost advantage. Technically, results of OSEM2 indicate that drayage companies could save significant time and distance in moving empties with a few well-located storage yards. The real issue is that these savings are offset in part by the need to reposition empties between yards. Considering the development of a few empty storage yards is a complex task, however, this analysis sheds light on possible problems with the development of such a system without establishing specific operating agreements in advance. For example, if on the average each drayage company was required to take from a yard approximately the same number of containers as it drops off at the same yard, then it would not be required to rebalance inventories. Second, it may be possible to establish a “virtual container yard” for the expressed purpose of “trading” drop off rights, or “pick up” requirements among drayage companies, in order to efficiently meet the operational rules of the storage yards.

The bottom line is that this is the first analysis of container management that involves looking at both perspectives: user optimal and system optimal. It is the first to involve a complete basin wide analysis for Los Angeles, and it is the first to estimate the total amount of empty container repositioning that takes place in the LA basin on a daily basis. It also demonstrates that planning for an empties storage policy away from the port must involve all parties, terminal operators, the port owners, and the drayage companies.

Recommendations and Further Work

Based upon the modeling and analysis performed as a part of this research project, we now know that there are a number of possible options in reducing vehicle transport and increasing vehicle utilization. For example, coordinating incoming trucks with outgoing containers may yield better stack management and lower service times within the port itself. Reducing the volume of empties being carted to and from the port can also be substantially reduced. Together, these two possibilities can help to increase the number of effective trips per truck and reduce trip mileage. This seems like a situation in which it should be relatively easy to move forward with some type of implementation. That fact is that there is a large amount of distrust among the parties (drayage companies, independent drayage truck owner/operators, port and terminal operators, local citizenry, local governments, Caltrans, shippers, consignees, etc). None will act unless they can be assured that the changes do not place their business at risk, require them to invest without being compensated in some manner, while at the same time environmental groups and citizenry want substantial reductions in traffic and air pollution. The complexity of the situation, especially in terms of the number of interested parties, makes this a challenging problem at the very least. To move forward there are several main questions to answer:

- What should we do?
- How do we pay for it?
- Who wins and who loses, with an attempt to mitigate the losses?

To answer these three questions, we need to:

- Provide convincing reasons why this option needs to move towards implementation
- Involve all of the stakeholders
- Provide final recommendations and a pathway towards implementation

To address these issues we need to engage all stakeholders. Although we cannot do this as university researchers, we can provide the very basis under which a meeting of the minds can take place. That is, with the support of the Ports, Caltrans, and local agencies including SCAG and SCAQMD, we can provide the technical support to getting to "Yes." Taking a "value-based" engineering approach led by one of the major agencies in LA involving representatives of the various stakeholder groups, we can provide a focused approach that is supported by a spatial decision support system. The basic idea is that all groups need to understand how impacts and costs change with a new storage system and have a basis for understanding how much they gain or lose with the various options. What is particularly important to understand is the fact that current operating policies by terminal operators make it nearly impossible to move forward, yet local populations are subjected to higher levels of pollution and traffic than necessary. To support a meeting of the minds it is important to develop a decision support system, which demonstrates the impacts and benefits associated with specific operating policies and storage yard locations in the basin. The idea is that a consensus can only be reached with adequate supporting information, so that all parties can understand the nature of the problem and the impacts to each stakeholder group.

To move forward there are several additional work elements that needs to be addressed. To support a the development of a SDSS we need:

- An accurate geographical account of container use by type and ownership: refrigerator units, standard 20's, standard 40's, ownership, etc.
- Bobtail mileage estimates and the extent of trip coordination
- Average match of drop-offs with pick-ups
- Full for empty at consignees
- Empty for full exchange at shippers
- Empty for full or full for empty at a port terminal or port complex
- Chassis management must be included in the models—this is currently a fluid problem because steamship lines are in the process of adjusting their policies.

We also need to derive addition data and metrics (especially associated with bobtail moves):

From truck tracking data:

- Estimate when a truck appears to drop off a container and pick one up at the port terminal (empty for full exchange or full for full exchange or full for empty exchange) from tracking data

From terminal operators:

- Survey terminal operators for a gross estimate as to the frequency of trips that involve an approach by a bobtail or a departure by a bobtail

From drayage companies/drivers:

- Estimate the frequency of trips to and from port terminals that involve a bobtail approach or a bobtail departure

From a sample of large and small shippers and consignees:

- Estimate the frequency of trips that involve an approach/departure that involve a bobtail.
- Estimate the level of interim storage possible

The results to date from the analysis of storing empties away from the LA and LB ports are significant, and the implications for moving forward are clear and unambiguous. It is within reason that drayage truck mileage on an average day can be reduced by 150,000 km. This estimate is generated from a "system"-optimal model and overlooks the nature of the user-oriented drayage business. User-optimal models demonstrate, in terms of traffic and mileage, that savings may not be realized without appropriate policies and consensus on key operational components.

This project has shed considerable light in terms of the prospects for moving forward. With additional funding and cooperation among stakeholder parties, we can provide a basis in which stakeholders can negotiate in a value-based engineering setting and move forward to development and implementation.

References

Boile, M., Theofanis, S., Baveja, A., Mittal, N., 2008. Regional repositioning of empty containers: case for inland depots. Transportation Research Record: Journal of the Transportation Research Board 2066, 31-40.

Crainic, T.G., Dejax, P., Delorme, L., 1989. Models for multimode multicommodity location problems with interdepot balancing requirements. Annals of Operations Research 18, 279-302.

Crainic, T.G., Gendreau, M., Dejax, P., 1993a. Dynamic and stochastic models for the allocation of empty containers. Operations Research 41, 102-126.

Crainic, T.G., Gendreau, M., Soriano, P., Toulouse, M., 1993b. A tabu search procedure for multicommodity location/allocation with balancing requirements. Annals of Operations Research 41, 359-383.

Gendron, B., Crainic, T.G., 1995. A branch-and-bound algorithm for depot location and container fleet management. Location Science 3, 39-53.

Imai, A., Nishimura, E., Current, J., 2007. A Lagrangian relaxation-based heuristic for the vehicle routing with full container load. European Journal of Operational Research 176, 87-105.

Jula, H., Dessouky, M., Ioannou, P., Chassiakos, A., 2005. Container movement by trucks in metropolitan networks: modeling and optimization. Transportation Research E 41, 235-259.

T.L. Lei and R.L. Church, 2010. Locating short-term storage facilities for empty transport containers. Manuscript submitted to Transportation Research Part E.

Mittal, N., 2008. Regional empty marine container management. PhD dissertation, Rutgers, The State University of New Jersey, New Brunswick, USA.

Namboothiri, R., Erera, A.L., 2004. A set partitioning heuristic for local drayage routing under time-dependent port delay, Transactions of IEEE International Conference on Systems, Man and Cybernetics, 3921–3926.

Theofanis, S., Boile, M., 2007. Investigating the feasibility of establishing a virtual container yard to optimize empty container movement in the NY-NJ region, Final report submitted to the 2005 UTRC Research Initiative, region II, Rutgers center for advanced transportation and infrastructure, Rutgers, The State University of New Jersey, Piscataway, NJ.

Tioga Group, 2008. Inland port feasibility study, Final Report (No. 06-023) to the Southern California Association of Governments. Tioga Group, Los Angeles, USA.

Xu, Y., 1999. A discrete choice based facility location model for inland container depots. West Virginia University, Morgantown, USA.

Introduction

Container terminals are large, interdependent and complex systems, with many processes and activities. They are generally located in urban areas, and have little available land for physical expansion or inland transport network improvement. Because of growing international trade volumes and increasing pressure to accommodate more container traffic, container terminals have been facing increasing congestion and have become bottlenecks in the global supply chain.

To deal with those challenges, container terminals around the world have been pursuing different strategies to improve terminal productivity, such as automation of terminal operations using emerging technologies, extended gate hours, reduced container dwell times, and gate appointment systems. Some of those strategies have been effective, while others have not met expectations for improvement. For example, appointment systems were perceived by the trucking industry as ineffective in reducing truck turn time and a wasted effort by many terminal operators. Figure 20 shows the results of a field survey conducted at the Ports of LA & LB regarding the effectiveness of gate appointment systems, illustrating that most systems were perceived by trucking firms as less than marginally effective in reducing the truck turn time (Giuliano and O'Brien, 2007).

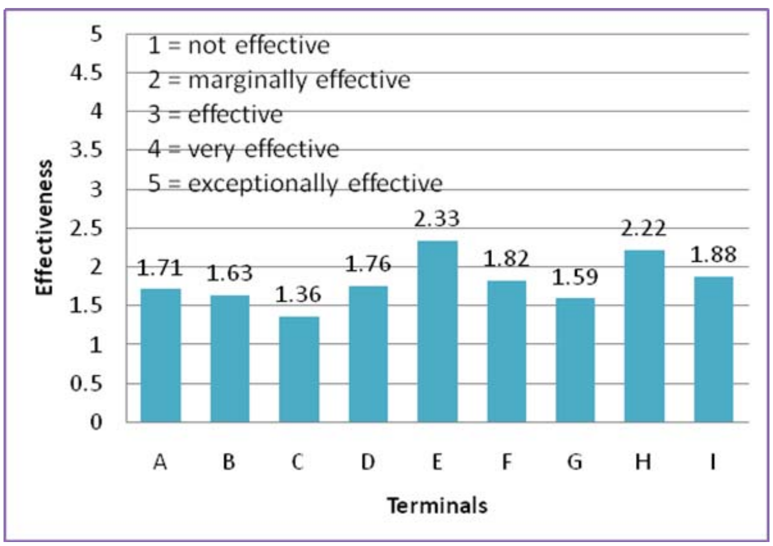

Figure 20. Effectiveness of gate appointment systems in reducing truck turn time at Ports of LA & LB

If utilized effectively by terminal operators, truck pre-arrival information obtained from the implementation of a gate appointment system should allow for greater terminal operating efficiencies, and therefore an improvement in truck wait time. In addition, the use of GPS on drayage trucks also provides opportunities for terminal operators to acquire more accurate information about the truck arrival time at the terminal gate. The research described in this report addressed two problems: 1) whether and how truck arrival information can be used to improve the drayage truck/ container terminal interface, and 2) whether historical GPS data can be used to measure truck drayage network reliability and forecast truck arrival times at the port.

This research only considers those terminals which stack containers on the ground at their container storage yards. In those terminals containers are often required to be relocated in order to access the desired container which may be buried underneath others. That activity is called container rehandling. This is unproductive work but unavoidable since truck arrivals are a stochastic process, and the truck arrival sequence seldom matches the container storage sequence. In current practice the containers are usually relocated to the nearest available stack, limiting the distance traveled by the crane to finish one rehandle operation. However, the storage location of rehandled containers affects the number of future rehandles.

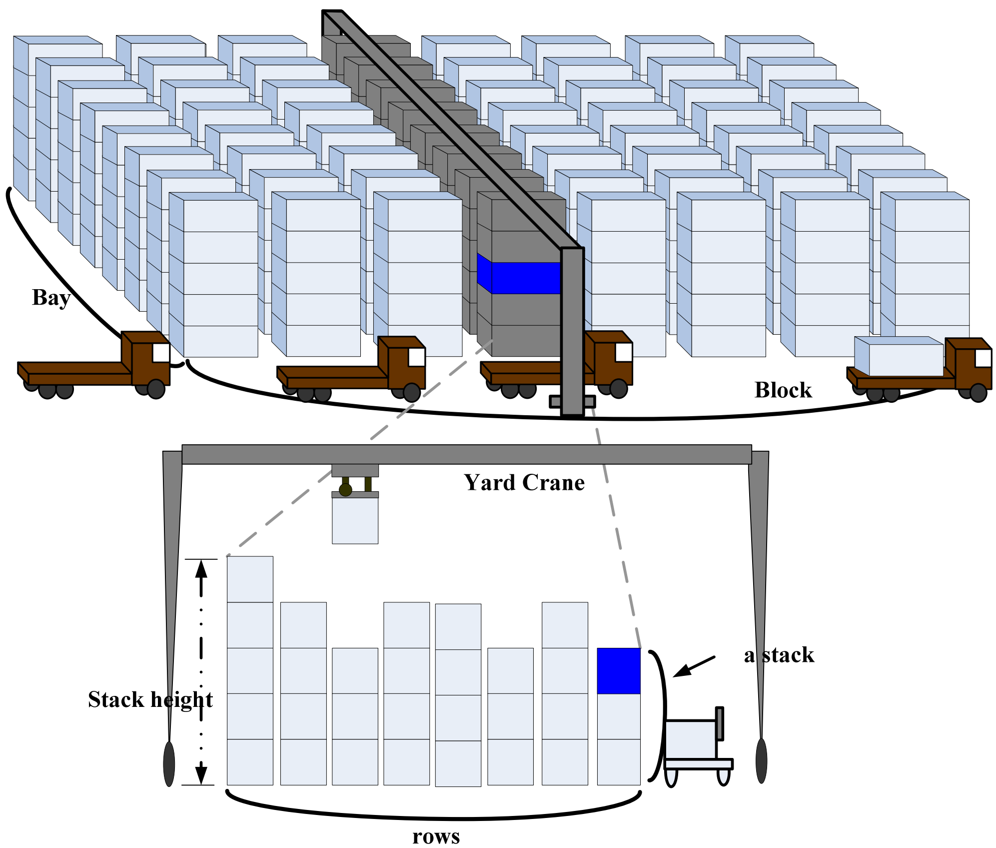

Figure 21. Container block, bay configuration and yard crane positioning

Consider a container bay with eight stacks and six containers in each stack (see Figure 21), and assume the containers to be retrieved are randomly distributed and rehandled containers are always relocated to the nearest available stack. Define yard crane efficiency as the ratio of productive crane moves to total crane moves as follows:

$$\text{crane efficiency} = \frac{\text{productive crane moves}}{\text{productive crane moves} + \text{unproductive crane moves}}$$

Productive crane moves are ones in which a desired container is moved. Unproductive crane moves are rehandles, or moves that relocate an undesired container in the process of obtaining the container of interest. To pick up all the containers from the bay, the expected number of unproductive crane moves aver a thousand iterations is 78, while the number of productive crane moves is 48, equal to the product of stack height and stack number. Crane efficiency is therefore 38 % for this case. This case, where there is no pre-planning of container storage, provides a lower bound on crane efficiency. This bound is not intended to represent expected terminal operations, but provide an upper bound on rehandling activity. In current terminal operations, rehandles still represent a significant level of effort in the terminal. By reducing container rehandles, the terminal could improve yard crane productivity, reduce truck transaction and delay times, and improve container throughput.

For each container stack, if the truck arrival sequence equals the sequence of containers in storage from the top of the stack to the bottom of the stack, rehandling activities can be completely eliminated. This provides a lower bound on rehandling activity. If the truck arrival sequence is known but does not match the storage pattern, during the process of retrieving required containers for waiting trucks, the storage location of rehandled containers can be carefully determined to avoid being rehandled again. Therefore, if truck arrival information along with container details is known in advance, more advanced container handling strategies can be used to reduce container handling work and truck delays. The objective of this research is to identify the truck information requirements for achieving a significant improvement in truck transaction time and yard crane productivity, evaluate the impact of different yard configurations on the effectiveness of this truck information, and assess the use of historical truck GPS data in providing truck arrival times.

Literature Review

There is a wide body of literature which considers improvements to marine terminal operations that is peripherally related to this research. Here we focus on the more closely related research. The most closely related paper was written by Jones and Walton (2002). They studied whether and how more accurate and timely information about the departure times of containers can be used to more efficiently and effectively manage import containers in stacked storage yards. They developed an event-based simulation model capturing the interactions among a port's various subsystems to evaluate the impact of using this departure information on the number of container rehandles, ship turnaround time, and average cost per container moved through the port. Their study assumes that the import container departure time has been acquired by the terminal operator prior to the ship unloading, and they used this information to determine the container stacking sequence on the yard during ship unloading process. While the overall intent is the same, to reduce rehandling activity, Jones and Walton study a different component of the terminal operations (unloading container from the vessel to stacks), and solve a different mathematical problem. In this research we assume the truck arrival time is obtained after import containers have been stored on the yard, to mimic the practice of having real-time, rather than strategic information.

Some research has focused on optimizing container storage and stacking logistics to reduce container rehandling work and improve yard operation efficiency. Kim and Hong (2006) proposed two methods for determining the locations of rehandled containers to minimize the number of rehandles during the pickup operation given the container retrieval sequence. First a branch-and-bound algorithm is suggested and then a decision rule is proposed. Although in numerical experiments the branch and bound (B&B) algorithm outperforms the heuristic algorithm, the computational time of the B&B algorithm exceeds the level appropriate for real time usage when problem size increases. Aydın (2006) studied the same problem as Kim and Hong (2006), but he considered minimizing not only the total number of rehandles, but also the total distance travelled by the crane. He first solved the problem using the B&B algorithm and the heuristic algorithm proposed by Kim, and also suggested two other alternatives, a greedy heuristic and the difference heuristic. His experimental results indicate that the solution gap between the heuristic and optimal algorithms is within 8%.

Some other researchers have studied how to reduce the truck transaction time at a container yard by better utilizing the current system or improving operational methods. Huynh (2008) studied regulating the number of trucks that can enter the terminal to make the gate appointment system effective. He proposed a methodology, which is a combination of mathematical formulation and computer simulation, to determine the maximum number of trucks allowed to enter the terminal while maintaining a target truck transaction time. Kim et al. (2003) studied sequencing trucks for container transfer operations to minimize truck delay at the container yard. A due time for transfer service is assumed for each truck, and delay of a truck beyond the due time incurs a penalty cost. A dynamic programming model was developed to minimize the total delay cost, and a learning-based method for deriving decision rules was suggested to solve the model. Kim and Kim (2002) studied optimizing the size of terminal storage space and number of yard cranes for handling import containers and developed an analytical cost model which addresses terminal space cost, investment and operating cost of yard cranes, and waiting cost of outside trucks. In that model truck cost was estimated based on truck transaction time, and transaction time was evaluated by formulating the container transfer operation for trucks as an M/G/1 queuing model.

The studies by Aydın (2006) and Kim and Hong (2006) are closely related to this research. Their research contributed to developing efficient algorithms to minimize the rehandling work given complete container retrieval sequence information. However, they didn't address the problem given incomplete information, nor do they evaluate the benefit to the terminal from adopting those strategies. This research addresses the problem how truck arrival time information with different levels of quality can affect container handling efficiency.

With regard to traffic prediction, its approaches can be categorized into three broad areas: (1) statistical models, (2) macroscopic models, and (3) route choice models based on dynamic traffic assignment (Akiva et al., 1992). In addition, researchers have applied artificial neural network (ANN) techniques for predicting roadway travel times (Park et al., 1999). Here a review is provided on literatures which attempt estimating or predicting vehicle travel times based on information collected from probe vehicle technologies.

Park et al. (1999) examined how to use real-time information collected from ITS technology for predicting link travel times for one through five time periods ahead (of 5-min duration). They employed a spectral basis artificial neural network (SNN) that utilizes a sinusoidal transformation technique to increase the linear separability of the input features. Actual link travel times from Houston that were collected as part of the AVI system of the Houston Transtar system were used as a test bed. It was found that the SNN outperformed a conventional artificial neural network and gave similar results to that of modular neural networks. The results of the best SNN were compared with conventional link travel time prediction techniques including a Kalman filtering model, exponential smoothing model, historical profile, and realtime profile, and it was found that the SNN gave the best overall results.

Rice and Zwet (2004) proposed a method to predict vehicle travel time on a freeway segment when its departure is at a certain time in the future. The prediction is based on the current traffic situation in combination with historical data, and the prediction method arises from the empirical fact that there exists a linear relationship between any future travel time and the current status travel time. Consequently, a linear regression model with time varying coefficients is developed for predicting travel times on freeways.

Chien and Kuchipudi (2003) applied Kalman filtering algorithm for predicting travel time based on real-time and historic vehicle information, which was collected by road side terminals. Factors that would affect the prediction results are explored, such as historical seeds. The results reveal that during peak hours, the historic path-based data used for travel-time prediction are better than link-based data due to smaller travel-time variance and larger sample size.

However, most of those studies focus on short-term travel time prediction and require extensive historical data as well as real time information to train the model. This research is not intended to provide real-time travel time prediction or routing guidance for truck drivers; instead this research is aimed to evaluate how predictable the truck arrival time at terminal gates is and how it varies across temporal and spatial extent of the drayage network.

Methodology

A Hybrid Approach of Simulation and Queuing Models

To address the question of whether truck arrival information can be used to improve the drayage truck/container terminal interface, this research considered the retrieval operation of import containers performed by the yard crane within a container block to serve drayage trucks (Figure 21). Some assumptions are made as follows:

- The yard crane serves the drayage trucks by the first-in-first-out rule (FIFO);
- Re-handled containers are relocated to a slot within the same bay;
- No additional containers are added to the block during the container pick-up process;
- Truck arrivals can be modeled by a Poisson process;
- The location of the container requested is randomly distributed;
- The location of each container in the block is known in advance and tracked throughout the pickup process;
- Truck arrival information includes the container to be retrieved.

Under the second assumption container bays are independent of each other; thus the analysis for container re-handling work is performed for one bay by one crane and the result is the same for any bay within the block. The results for the operation of one yard crane within a container block can be extended to the whole container yard with multiple yard cranes given identical assumptions for each crane. In that situation the container yard can be segregated into multiple sub-areas and each sub-area is assigned to one yard crane, with each crane modeled as an independent system.

Truck information is considered for container retrieval within the same bay. Based on the amount of known truck information and whether the information is static or updated in real time, six scenarios are defined to represent situations with various information qualities (Table 6).

Table 6. Scenario Definitions

Scenario	Definition
No truck information	No truck information is available
Static group information	The terminal knows which of several groups a truck will arrive in, but not of the exact order of truck arrivals within any group. For example, trucks can be assigned to two groups, A, and B. The terminal knows which trucks are in group A and which trucks are in group B, and that all trucks in group A will arrive before any truck in group B. But the exact arrival sequence of trucks within group A or B is not available. “Static” means information is provided before any truck arrives, and is not updated over time.
Static partial sequence	The terminal knows which of several groups a truck will arrive in, and the exact order of truck arrivals for the first group. Information is not updated over time.
Dynamic group information	The terminal knows which of several groups a truck will arrive in, and the group information is updated over time. Every time all the trucks in the first group are exhausted, the terminal receives information about the arrival group of the next N trucks, where N is the number of trucks in the original first group.
Dynamic partial sequence	The terminal knows which of several groups a truck will arrive in, and the arrival sequence of the first group. After a truck in the first group is served, information about the first truck within next group becomes available, and this truck joins the first group.
Complete sequence	The complete sequence of truck arrivals is known.

Without truck arrival information, rehandled containers can be relocated to the nearest available stack. This strategy minimizes the travel distance of the yard crane and is used widely in container terminals. This strategy will be called the nearest relocation rule and applied to the scenario with no truck information for container relocation. For all the other scenarios with some truck information, a solution approach called revised difference heuristic is proposed for using truck information to reduce container rehandling work. This algorithm requires each container’s retrieval order number as input. The retrieval order number can be obtained by relating the truck arrival sequence (or group) to the container of interest. Given truck arrival information, the revised difference heuristic can be applied to determine the best storage location (the one that incurs fewest future rehandles) of the rehandled container. Let X denote the order number of the container to be rehandled, the revised difference heuristic is described below.

Revised Difference Heuristic

Step 1: When relocating container X, search for a stack with container Y whose order number is smallest in its stack and still bigger than X. In this way no additional rehandles will be necessary for container X. If multiple stacks satisfy this condition then the stack containing smallest Y is chosen. If such stack does not exist, go to step 2.

Step 2: Search for a stack in which the container with the smallest order number is the same as X. If multiple stacks satisfy this condition then randomly select one. If such stack does not exist, go to step 3.

Step 3: Search for a stack with container Z which is accessible by the crane and has an order number smaller than X. If multiple stacks are found, then the one with largest Z is chosen to minimize the difference between X and Z. If such stack does not exist, go to step 4.

Step 4: Search for a stack to minimize the difference in order number between its top container and X.

Decisions are made sequentially regarding relocations using the revised difference heuristic, from the top container on the target stack (the stack in which the requested container is located) to the one just above the required container.

To evaluate the effectiveness of using truck arrival information, a hybrid approach of simulation and queuing was developed. As illustrated by Figure 22, a computer based simulation was first developed to model container retrieval for a bay of containers under the revised difference heuristic algorithm, which determines the optimal storage location of rehandled containers. That simulation model takes the container block design and truck arrival information as input, and ouputs the expectation and variance of the number of rehandles under different scenarios. Then an M/G/1 queuing model is formulated by considering the yard crane as the single server and the drayage trucks as the customers whose arrival follows a Poisson distribution. This queuing model takes the simulation model output and truck arrival rate as inputs and outputs the crane productivity and truck turn time. Detailed explanations of the simulation model and queuing model are provided in following sections.

Computer Simulation

Computer simulation is developed to model the container movements and evaluate the container rehandling work under different block designs and truck information qualities. The truck arrival information and container bay configuration are taken as program inputs, and the parameters used to define the two sets of inputs are listed in Table 7 and Table 8.

The parameters listed in Table 7 and Table 8 are user-defined inputs into the programs. Four parameters are considered for each scenario to define the information quality (Table 7): the length of known subsequence, which refers to the number of trucks within the known arrival sequence; the number of groups; each group size; and the information updating rule, representing whether the information is real-time updated or not. Another five parameters are required to define the bay configuration (Table 8): number of stacks, stack height, stack storage capacity, bay balancing condition, and bay loading percentage.

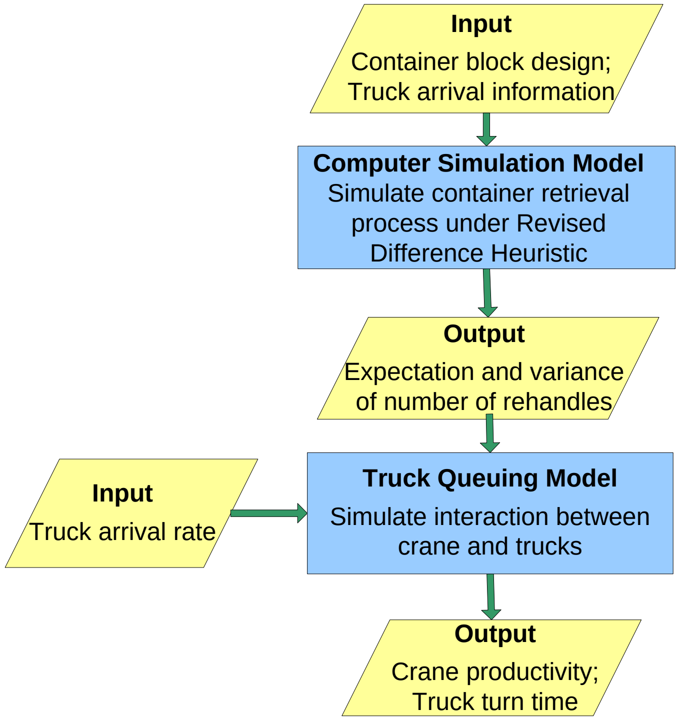

Figure 22. Hybrid approach of simulation and queuing model

Table 7. User-defined settings

Scenarios	**Parameters**			
	Length of known subsequence	**Number of groups**	**Group size**	**Information update rule**
Static group information	0	Within [2, total truck pool size]	Within [1, total truck pool size]	No update
Static partial sequence	Equal to the size of the first truck group	Within [2, total truck pool size]	Within [1, total truck pool size]	No update
Dynamic group information				Updated in terms of group size
Dynamic partial sequence				Updated in terms of truck unit
Complete sequence	{Number of stacks} × {Stack height}	—	—	—

Table 8. Parameter settings

Description	Parameter setting
Number of stacks	Within the range [2, 12]
Stack height	Within [2, min(number of stacks, 6)]
Stack storage capacity	Maximum stack height + 1
Bay balancing condition	Balanced bay, or unbalanced bay
Bay loading percentage	33.3%, 50%, 66.7%, 83.3%, 100%

The computer programs are written in Matlab, and the container bay is modeled using arrays to represent the storage locations. The stacking sequence of containers in the bay is randomly generated, with containers represented by retrieval orders and stored in an array. The truck arrival sequence or groups are generated according to the specified value of parameters (those in Table 7). Two different functions are written for determining the storage location of the rehandled container, respectively representing the nearest relocation strategy and RDH. The main program simulates the container pickup operation under each solution approach by calling the corresponding function when a container is required to be rehandled and updating its storage location in the array. One counter is used to track the total number of rehandles, and updated whenever a rehandle occurs. Many problem instances can be specified, and the program evaluates and outputs the expectation and variance of number of rehandles for each scenario with specified parameters based on simulation results, which serve as the input into truck queuing model.

Truck Queueing Model

A truck queuing model was developed to evaluate crane productivity and truck transaction time. For a yard crane working within a block of inbound containers, the container retrieval operation can be modeled as an M/G/1 queuing system, with the yard crane being the single server and the arriving trucks as customers (Figure 21).

The server's parameters are determined by decomposing the crane service time into inter-bay travel time, container rehandling time and handling time and estimating the variance and expectation of each time component. Based on the assumption that trucks are served FIFO and the requested container location is randomly distributed, the inter-bay travel time between two random retrievals can be easily calculated based on statistical derivation. With regard to crane rehandling time, we assume that the number of re-handles and the time to re-handle one container is independent, and consequently the expected re-handling time can be calculated as the product of expected number of re-handles and the expected time to re-handle one container. The variance of container rehandling time is derived by assuming the variance of the time to rehandle one coantainer as zero because its value is small enough that its impact on model output can be neglected. The crane handling time can be calculated in a similar way as the rehandling time.

After determining the server's parameters, the service rate of the queuing system can be calculated using following formula:

$$\rho = \lambda \cdot E(T_c)$$

The expected truck transaction time can be estimated as:

$$E(W) = E(T_c) + \frac{\lambda \cdot V(T_c) + \rho \cdot E(T_c)}{2(1-\rho)}$$

And the crane productivity is:

$$P = 1/E(T_c)$$

where T_c is crane service time;

ρ is the service rate of the queuing system;

λ is the truck arrival rate.

Truck Arrival Time Prediction

To understand whether historical GPS truck data can be used to estimate truck arrival times, we considered the regional truck drayage network serving the Ports of Los Angeles and Long Beach. Three months of truck GPS data (January to March, 2010) collected through the project was used for analysis. To protect the privacy of businesses, the truck trip was truncated at the freeway interchanges closest to the origin/destination location. The truck information obtained from the GPS dataset includes the truck trip ID, location (coordinates), time and date stamp.

To evaluate the port drayage network reliability, the coefficient of variation was used as the reliability measure:

$$COV = \frac{\text{Standard deviaiton of travel time}}{\text{Mean travel time}}$$

This reliability measure was used to examine how the truck travel time varies across the drayage network and during different times of day and days of the week.

To forecast the truck arrival time at container terminals, a simple method was proposed to predict the 95% confidence interval of truck travel time between given OD pairs. This method assumes that the truck travel time on each roadway link is independent, and thus the mean travel time (or variance of travel time) of a path can be estimated by summing up the mean travel times (or variance of travel times) on the connecting links of this path. Two months' GPS data was used for predicting travel times (data collected from January to February), and last month's GPS data was used for validating the accuracy of prediction results. Then this method was applied to the entire drayage network to evaluate how the truck arrival time window at terminals varies with the trip's origin and departure time.

Research Results

The first part of this section presents the estimated improvements in crane productivity and truck transaction time if a terminal utilizes truck arrival information to reduce rehandling work, and the second part presents a reliability analysis of the port drayage network and the results of the travel time prediction.

Impact of truck arrival information on system efficiency

The impact of various information qualities, truck arrival rates, and block configurations on drayage truck/ yard crane system performance was evaluated to identify the effectiveness of truck information under different system configurations.

The analysis demonstrated that truck information can generate significant benefits for both the marine terminal and drayage trucks, and updating information in real time can lower the requirement on information quality. Figure 23 and Figure 24 provide an example of crane productivity improvement and truck transaction time reduction under various scenarios for a block configured with 40 bays, 6 rows, and 5 containers in each stack, and truck arrival rate of 6 per hour. Notice the similarities between the two figures, indicating that change in truck information quality has similar impacts on both crane productivity and truck transaction time. Two other observations can be made from Figure 23 and Figure 24. First, given static information, the value of truck group information is maximized when the sizes of two groups are equal. The value of partial sequence information grows steadily with the length of sequence. Second, updating information in real time can lower the requirement on information quality. For the scenario with dynamic group information, peak benefit is realized at a much smaller first group; for the scenario with dynamic partial sequence information, significant benefit is achieved from knowing 1/6 of the total sequence and little additional value is generated from a longer sequence. Therefore, a complete sequence is not required for significantly improving system performance if real time information is available.

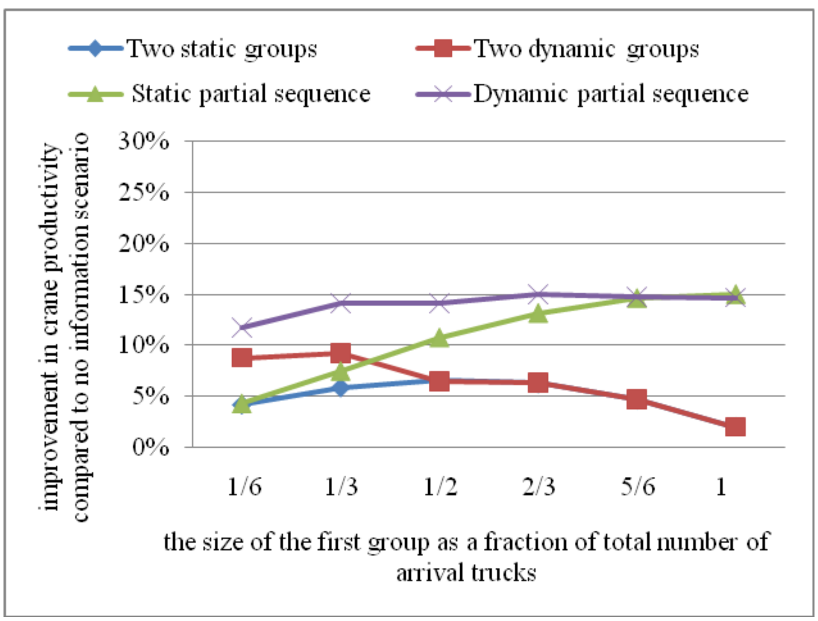

Figure 23. Improvements in crane productivity under various first truck arrival group sizes

The performance analysis under various truck arrival rates illustrates that reductions in truck turn time achieved from any level of information quality grow exponentially with the truck arrival rate. Figure 25 provides an example of truck transaction time reductions under various arrival rates for a block with 40 bays, 6 rows, and 5 containers in each stack. It is assumed that arriving trucks retrieving containers from the same bay are assigned into two groups, with the first group accounting for 1/3 of the total number of arriving trucks. When the truck arrival rate is approaching the crane service rate, a 35% reduction in transaction time can be realized from only knowing truck arrival groups. Therefore, the truck information is more valuable for the system operating near capacity, and a small amount of truck information can be very effective in reducing truck delay. Figure 25 also demonstrates the consistent effect of truck information quality on truck transaction time under different truck arrival rates. In general, information for two static truck groups can generate almost 1/2 of the truck time saving achieved from complete sequence; dynamic group information is more valuable than knowing 1/3 of truck arrival sequence and can result in an additional 2%-4% time saving; dynamic partial sequence information can provide almost the same amount of benefit as complete sequence information. Therefore, better information quality can further reduce truck transaction time but the complete sequence is not required.

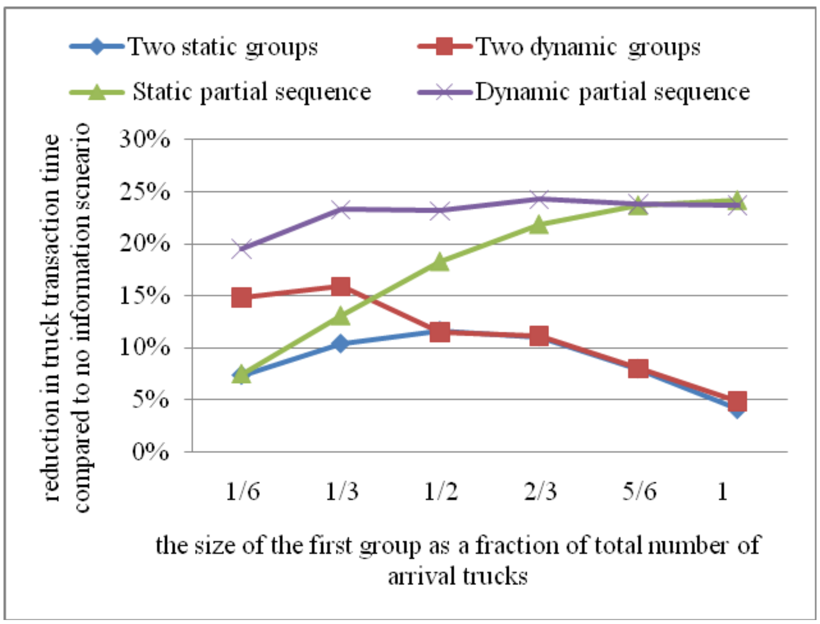

Figure 24. Reduction in truck transaction time under various first truck arrival group sizes

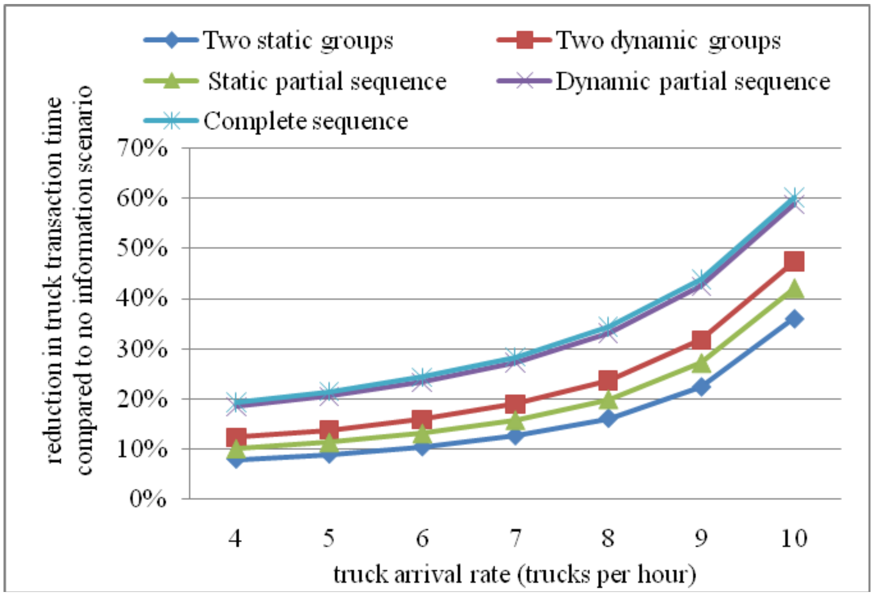

Figure 25. Percentage savings in truck transaction time under various arrival rates.

The performance analysis results with different block configurations demonstrate that stack height is the most important design factor regarding the effectiveness of truck information. Figure 26 provides an example of the truck transaction time reduction under various block configurations for a block with a total of 1200 containers and six rows in width, and truck arrival rate as 6 per hour. It is assumed that arriving trucks retrieving containers from the same bay are assigned into two groups, with the first group accounting for 1/3 of the total number of arrival trucks. It can be observed from Figure 26 that the truck information generates larger benefit for the block configuration with higher stacks and fewer bays given the same level of information quality. In addition, better information quality can bring additional benefit for the block configuration with higher stacks and fewer bays; however, its value decreases with the stack height. Static group information is sufficient for system improvement for the block configuration with shorter stacks and more bays. Figure 27 illustrates truck transaction time change under the block configuration with initial stack height as five. It can be observed that given the same level of information quality, the information provides larger benefit for the block configuration with more rows and fewer bays. In addition, the magnitude of benefit grows steadily with better information quality for any combination of row numbers and bay numbers. The comparison between Figure 26 and Figure 27 shows that stack height has more impact on the effectiveness of utilizing arrival information than other block configuration factors.

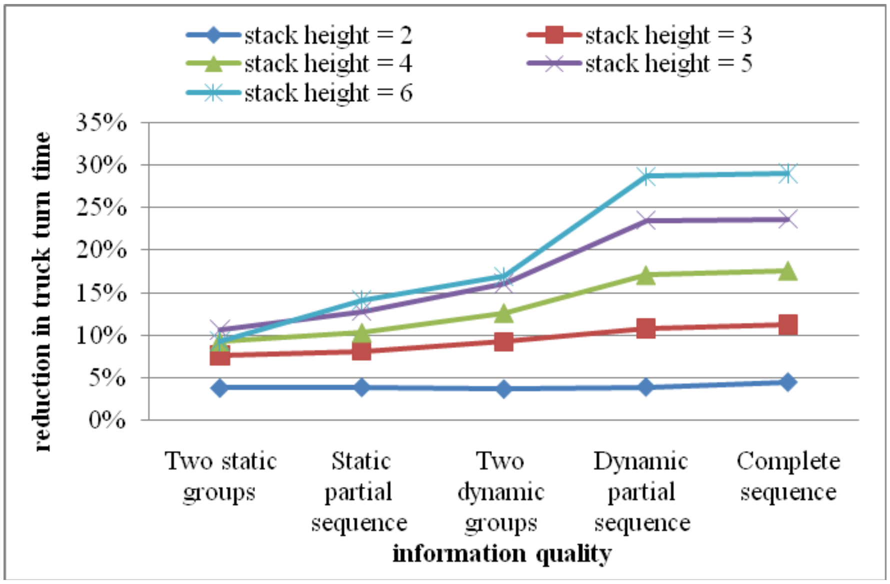

Figure 26. Truck transaction time under various configurations of stack height and bay numbers

Reliability analysis of port drayage network and truck arrival time prediction

The drayage network reliability was examined by calculating the coefficient of variation (COV) of truck trips terminating at the Ports of Los Angeles and Long Beach. The impact of the trip origin's distance from the port on truck trip reliability is illustrated in Table 9. It can be observed that the coefficient of variation does not vary much across the network and thus the origin's distance from the port has little impact on the travel reliability. Notice that the coefficient of variation (COV) is a relative measure of travel reliability. As the travel time or distance from the port increases, so does the average trip time, and therefore the denominator. The variability in terms of number of MINUTES may increase, while the COV may not. The trip travel reliability during different times of day is shown in Figure 28. Figure 28 illustrates that the coefficient of variation is higher during AM and PM time periods but much lower during midday, night and weekend. Therefore, the drayage network is less reliable during morning and evening peak hours compared to other times of the day.

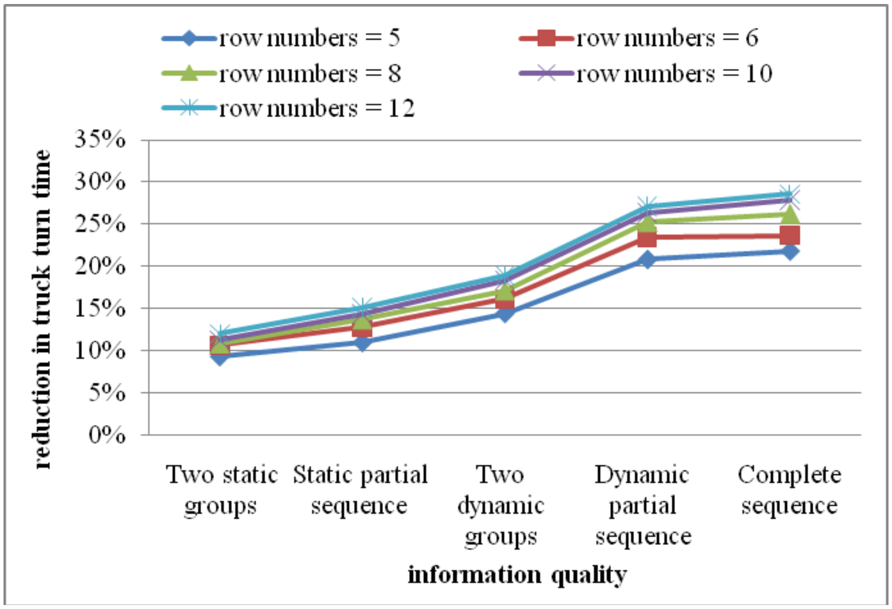

Figure 27. Truck transaction time under various configurations of row numbers and bay numbers

Table 9. Impact of Euclidean distance of trip origins from the port on the mean COV

Euclidean distance of trip origin from port (miles)	0-9	10-19	20-29	30-39	40-49	50-59
Mean Coefficient of Variation	0.151	0.141	0.195	0.139	0.167	0.114

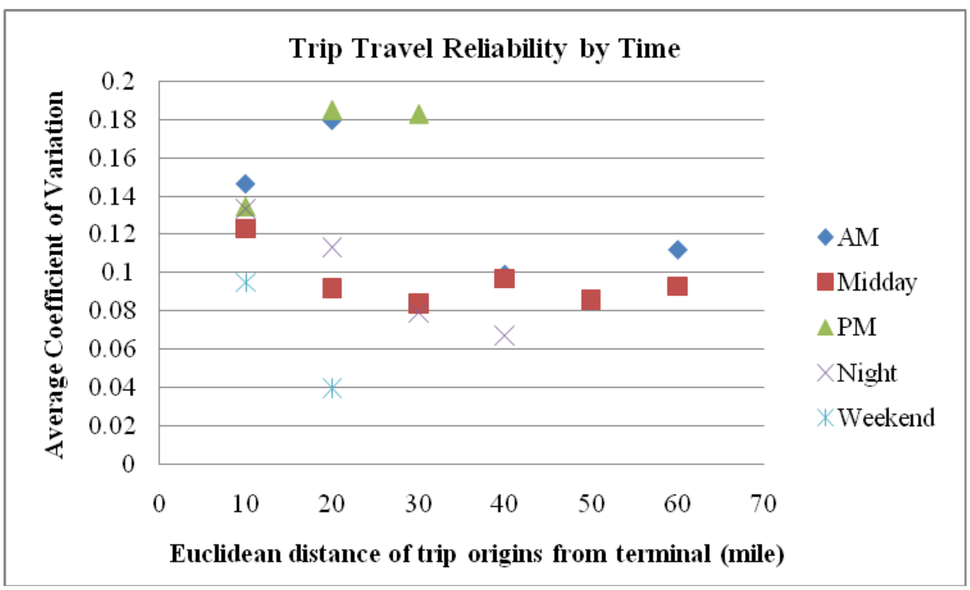

Figure 28. Coefficient of variation averaged over spatial extent during different time periods

To forecast truck arrival times at terminals, two months of truck data was used for model development and the last month's data reserved for model validation. Table 10 shows the 95% confidence interval of travel time estimated for truck trips terminating at the intersection of I-710 S and Ocean Blvd. The second to fifth columns of Table 10 represent the predicted confidence interval during different times of the day by considering all the historical trucks routes, and the last column shows the prediction accuracy by comparing the actual trip time with the predicted confidence interval. Each row corresponds to an OD pair (origin is marked with a star, and the destination is marked with a circle in the Figure 29). It can be observed from Table 10 that the predicted confidence interval is quite accurate in estimating the truck arrival time window at the terminals, since the accuracy rate of travel time prediction for any OD pair is no less than 84% and the average accuracy rate is as high as 94%. In addition, Table 10 and Figure 29 indicate that when the trip origin is close to the port, our proposed prediction method could provide a quite tight estimation of truck arrival time window. For example, with regard to the trips originating from the intersection of I-710 S and I-405, the width of the predicted travel time interval is only 4 minutes. If the trip origin is farther away from the port, this prediction method is still able to provide a tight estimation of arrival time window for those trucks departing during midday and night. This can be explained by the network reliability characteristics. As presented in above paragraph, the drayage network is more reliable during midday and night; therefore, the travel time variability is also smaller during midday and night, resulting in a narrower confidence interval of travel time. Therefore, our proposed prediction method is very effective in estimating the truck arrival time window and can be utilized by terminal operators to acquire better knowledge about truck arrival information.

The model validation shows that the estimated 95% confidence interval is sufficiently accurate to be used in predicting group arrival time windows of trucks. The methodology demonstrated here will be useful for terminal operators to estimate the truck arrival time window if the truck departure time and origin information is known in advance. In addition, more accurate prediction can be obtained if the terminal operator knows the truck routing information. However, such information is not necessary and the prediction can be made based on historical data. Also, such travel time interval estimation is more useful for terminal operators if the truck departs its origin during midday or night because a narrower arrival time window will be achieved, which could translate into truck arrival information with higher quality, such as arrival sequence instead of arrival groups.

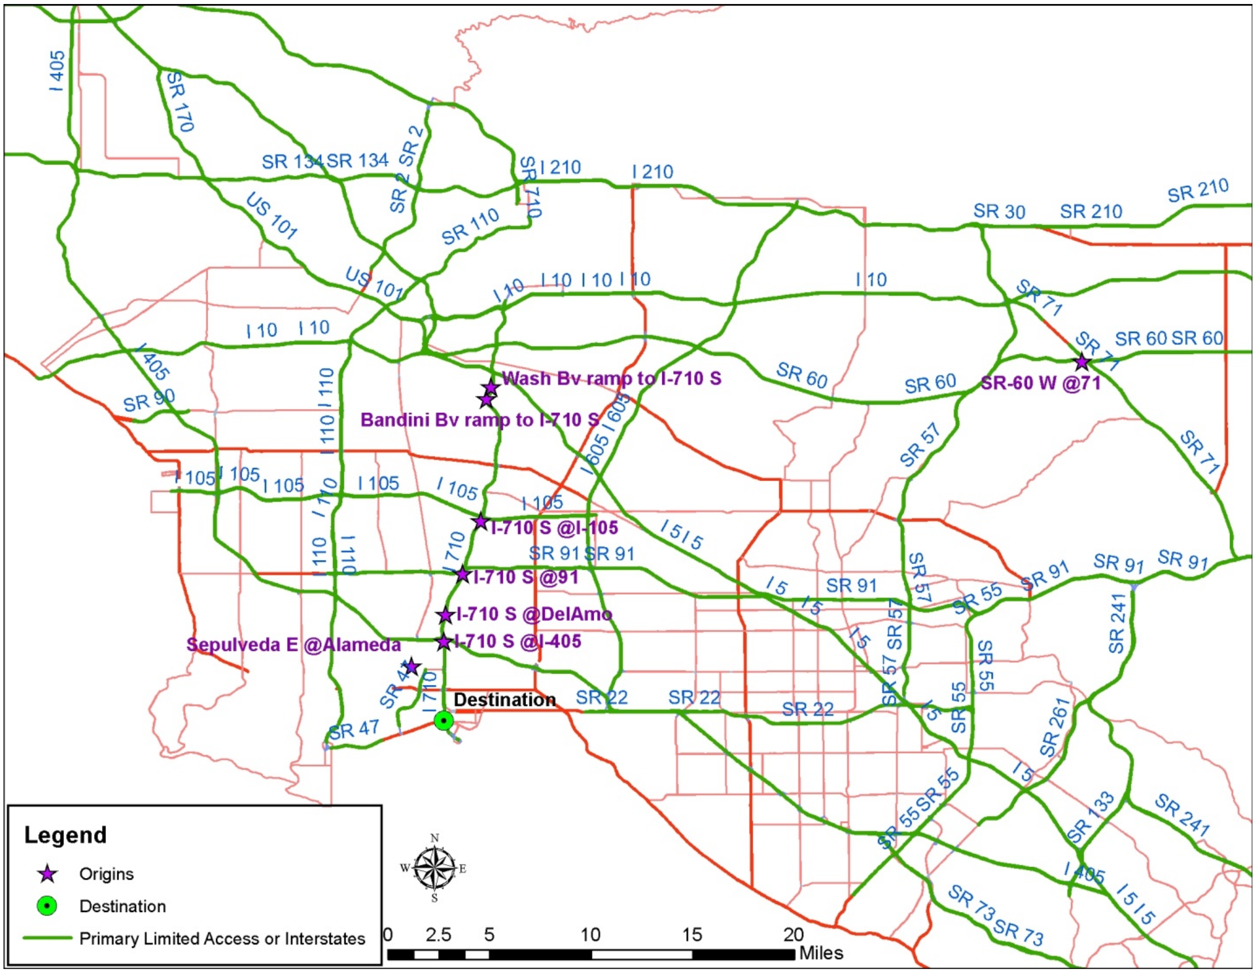

Figure 29. Distribution of estimated travel time interval across temporal and spatial extent of the drayage network

Table 10. Predictive accuracy of travel time to port, by trip origin

Origin	AM	Midday	PM	Night	Predictive Accuracy
Sepulveda E @Alameda	[7, 11]	[7, 11]	[8, 12]	[6, 10]	92.6%
I-710 S @I-405	[4, 6]	[4, 6]	[4, 6]	[4, 6]	100.0%
I-710 S @Del Amo	[5, 9]	[5, 9]	[5, 9]	[5, 9]	90.8%
I-710 S @SR-91	[7, 13]	[7, 11]	[7, 11]	[7, 10]	100.0%
I-710 S @I-105	[10, 17]	[10, 14]	[10, 15]	[9, 13]	92.9%
Bandini ramp to I-710 S	[17, 27]	[15, 27]	[19, 27]	[14, 27]	97.9%
Wash Bv ramp to I-710 S	[18, 34]	[16, 34]	[18, 34]	[15, 34]	92.7%
SR-60 W @SR-71	[51, 103]	[39, 74]	[44, 78]	[41, 55]	84.0%

Conclusion

In this research we evaluated the impact of truck arrival information on the drayage truck/container terminal system. A simple rule for using truck information is adopted to reduce container re-handles, and a hybrid approach of computer simulation and queuing theory is used to model the interaction between the yard crane and arriving trucks. The model is designed to evaluate how strategic factors, such as the level of truck information quality and container block design, affect system improvements achieved from utilizing truck information. These results can identify terminals likely to experience significant benefits, and inform the design of a data sharing system. For very detailed estimates of improvements at a particular terminal, a micro-simulation model should be developed that captures the unique terminal configuration, flow rates, and processing times.

Our research results demonstrate that truck arrival information is effective for improving crane productivity and reducing truck transaction time. Group information alone can effectively improve system performance; updating information in real time lowers the information requirement and provides significant benefit at small amount of information. In fact, real-time partial sequence information can generate about the same benefit as the complete arrival sequence, even if the partial sequence is for just 1/3 of total number of trucks. Complete sequence information is not required to maximize the benefit.

The results also shed light on the relationship between benefits and block configuration. For those terminals with limited yard space and high stacking, truck information is more effective for system improvement and better information quality is useful for further enhancing the magnitude of benefit. For those terminals with more yard space, the static truck group information can moderately improve system efficiency. Truck information is especially valuable for the system operating near capacity.

The work also illustrates that historical GPS data can be used to quite accurately predict truck arrival windows at the terminals and provide more knowledge about truck arrival information for terminal operators. Tight estimation of truck arrival time can be derived from the GPS data for truck trips departing during midday or night, which could be translated into more accurate truck information, such as truck arrival sequence instead of just arrival groups. This information could be used by terminal operators to further improve the drayage truck/ container terminal interface.

Recommendations and Further Work

We set out to answer the question of whether or not advanced arrival information regarding trucks could be used to reduce both terminal work and truck wait time. Our results show this is clearly true:

1. Small amounts of information can be used to reduce rehandling work in the terminal and therefore truck turn around time.

   Complete truck sequences are not required. In fact, any amount of information that can be obtained, can improve terminal operations and therefore truck delay time. If one rehandle is reduced, this is an improvement. While more information allows for more significant benefits, even information regarding a small number of groups can provide benefits.

2. Given the road network around the San Pedro Bay ports, and current traffic conditions, travel times can be estimated sufficiently accurately to implement such a system.

   Current travel times are sufficiently predictable to implement such a system.

3. Both the truck and terminal reap the rewards of sharing information.

The terminal's effort is paid off in reduced terminal operating cost, increased velocity, and increased capacity. The truck's effort it paid off in reduced wait times, and therefore reduced costs, reduced turn times, and increased capacity. Regionally, there would be benefits for traffic congestion and air quality. The terminal does not need to incur cost for the benefit of others. The same is true for trucks.

While information sharing can be undertaken between individual terminals and trucking companies (as it is currently implemented at the Port of Seattle), we recommend the establishment of a third party who would be responsible for receiving and transmitting data. The cost of such an organization would be born in part by the trucking companies and terminal operators, but also by the port authority. Much like Pierpass, this entity would be responsible for data confidentiality and would serve all parties involved.

Truck arrival information would need to be conveyed to yard crane operators. Currently, these crane operators receive information regarding desired containers, in many cases, on visual displays in the crane. Operators could be provided with specific instructions as to where containers should be rehandled by the terminal management software. Terminal management software would determine the best storage location in concert with other decisions, after having received the arrival information from the third party data manager.

Whether the cranes are operated by individuals or automatically is immaterial. Reducing wasted effort is reducing wasted effort. Once in place, adjustments to terminal resources could be made based on terminal specific preferences for extra capacity and resilience.

References

Aydın, C., 2006. Improving rehandling strategies for container retrieval process. Master Thesis, Graduate School of Engineering and Natural Sciences, Sabancı University.

Ben Akiva M., Kroes, E., De Ruiter, J., Smulders, S., Whittaker, J., Lindveld, K., 1992. Real-time Prediction of Traffic Congestion. Vehicle Navigation and Information Systems: Conference Record of Papers, New York, NY, Institute of Electrical and Electronics Engineers, 557–562.

Chien, S., Kuchipudi C.M., 2003. Dynamic Travel Time Prediction with Real-time and Historical Data, Journal of Transportation Engineering, 129(6), 608-616.

Giuliano, G., O'Brien, T., 2007. Reducing port-related truck emissions: The terminal gate appointment system at the Ports of Los Angeles and Long Beach. Transportation Research Part D 12, 460-473.

Jones E. G., Walton C. M., 2002. Managing containers in a marine terminal: Assessing information needs. Transportation Research Record: Journal of the Transportation Research Board No. 1782, 92-99.

Huynh, N., Walton, C.M., 2008. Robust scheduling of truck arrivals at marine container terminals. Journal of Transportation Engineering 134(8), 347–353.

Kim, K. H., Hong, G. P., 2006. A heuristic rule for relocating blocks. Computers & Operations Research 33, 940-954.

Kim, K. H., Kim, H. B. The optimal sizing of the storage space and handling facilities for import containers. Transportation Research Part B: Methodological 36, 821–835.

Kim, K.H., Lee, K.M., and Hwang, H., 2003. Sequencing delivery and receiving operations for yard cranes in port container terminals. International Journal of Production Economics 84(3), 283–292.

Park, D., Rilett, L.R., Han, G., 1999. Spectral basis neural networks for real-time travel time forecasting. Journal of Transportation Engineering 125 (6), 515–523.

Rice, J., Zwet, E.V., 2004. A simple and effective method for predicting travel times on freeways. IEEE Transactions on Intelligent Transportation Systems 5(3): 200-207.

6—Outreach and Commercialization

From the outset, this project was designed to have a strong applied component, with four private sector participants partnering with two universities. End-users and industry decision makers were consulted extensively, on the Steering Committee, and in the course of the research. Consequently the research effort is well known in key areas of federal, state and local government, the ports, rail and trucking industries.

Outreach

Outreach on the MeTrIS vision commenced in 2005, two years prior to the start of this research project, with a series of meetings with state and local agencies, private sector investigative teams, and potential industry partners. This effort produced a balanced team of investigators and cost-sharing partners.

In the course of the research, concepts and interim results were presented orally and via posters to more than 30 audiences, including scientific meetings (National Urban Freight conference, Esri Space-Time conference), several Transportation Research Board meetings, the Harbor Association of Industry and Commerce, the Harbor Trucking Association, California Maritime Leadership Symposium, consulting firms and area agencies including California Department of Transportation, Southern California Association of Governments, the office of the Mayor of Los Angeles and the Port of Seattle.

Steering Committee

A Steering Committee was constituted in 2007, consisting of senior experts in a broad variety of associated fields:

- Kerry Cartwright, Port of Los Angeles
- Don Cooke, Tele Atlas North America
- Deborah Estrin, University of California, Los Angeles
- Jim Fitzgerald, BNSF Railway
- David Maguire, Esri
- Frank Quon, California Department of Transportation
- Randy Rogers, USDOT Maritime Administration (MARAD)
- Linda Styrk, Port of Seattle

Two meetings were held in the course of the project. Proceedings of those meetings are described in detail in quarterly reports. Some key recommendations of the committee are worth repeating briefly as they played a significant role in the direction of the project:

- Air quality issues are a valuable area of study, but they are a significant body of research in themselves and can detract focus from the project.
- A study of commercialization potential should be undertaken from the outset.

- The primary value of the study lies in data, data, data.
- Consider expansion of the program to other ports, even in the course of the study period.

Web site

Two outreach web sites have been maintained in the course of the project. The National Center for Geographic Information and Analysis (NCGIA) at UCSB is host to the academically oriented web site, with a focus on research reports (*www.ncgia.ucsb.edu/ncrst*). Separately, DGRC maintains a site with a focus on technology and the industry (*www.metris.us*). Other private web pages hold fleet performance reports for participating trucking firms.

In addition, the consortium has hosted the gateway for the entire NCRST research effort, at *www.ncgia.org*. Research reports of other consortia are preserved at this site, to provide continuity to the user community, and legacy and context for new consortia.

Events

The consortium hosted two Steering Committee meetings, and visits by the Agreement Officer Technical Representative, Caesar Singh, and the Administrator of the Research and Innovative Technical Administration, Peter Appel.

Commercialization

Several methodologies employed or developed by the project have potential for commercialization in terms of identified need and benefit to end-users, though the path to profitability is not easy or assured:

- Specifications of the vehicle tracking system introduced three years ago are still not matched in the mainstream marketplace, in terms of data quality and price. Many pieces of this system are already commercial. There is relatively little potential for further commercialization in isolation from other project components, because of the rapidly changing nature of the underlying hardware.
- Conflation techniques have been well received academically and have already been independently funded by another federal agency.
- Models for empty container management and port-truck synchronization are strong candidates for commercialization. The recommended measures require changes to practice and investment in information system modifications, and they affect revenues in the short term. While considerable benefits are forecast, implementation is challenging.
- Models for truck-terminal synchronization can be applied in a relatively short timeframe. Again, there are institutional challenges to be overcome, in the absence of which there is no commercial promise.

A credible commercialization agenda has to consider (a) potential uses of the technologies: need for information, appropriate delivery of that information (through raw data, visualization, analyses and models), business development opportunities and support, (b) markets for the products, value, pricing, expense and revenue forecasts, and investment opportunities, (c) production, staffing, management structures and business partnerships.

Ultimately commercial success depends on a large number of factors, from the substance of the offering to the quality of leadership and promotion, and the timing of product availability with respect to market needs and competing products and services.

Activities most appropriate for pursuit under this federally funded project are in the realm of item [a]. Items [b] and [c] fall under business planning, and are not pursued in this study. It is expected that at least some consortium members will advance their respective creations to these steps in the coming months or years.

The remainder of this section focuses on the components of [a] that are not addressed elsewhere in this report, in particular the business development environment at the ports, examples of commercial successes and the challenges faced by some past projects.

Team Expertise

Two team members were particularly active developing the commercialization component of the research. John Glanville of Athenaeum Capital Partners has extensive experience in the development, management and funding of technology startups in the information and port technology arenas. William Lyte, representing the California Marine and Intermodal Transportation Systems Advisory Committee (CALMITSAC), is a source of considerable knowledge and experience on the port industry and local business development initiatives.

Examples from the Industry

The San Pedro ports are in many respects an international showcase of port technology, and encourage development in four principal areas: emissions, congestion/logistics and throughput, security and renewable energy.

Several research efforts have been undertaken in cooperative ventures among private and public agencies and local universities. A variety of business models have been adopted, that may serve as models for commercialization of our technologies.

POLA/POLB Clean Air Action Plan (CAAP)

In 2006, the ports of Long Beach and Los Angeles created and approved the San Pedro Bay Ports Clean Air Action Plan. The CAAP, as it is known, provides the overall strategy for dramatically reducing air pollution emissions from port-related cargo movement. CAAP's primary goal was to dramatically reduce emissions and their associated health risks for the Southern California region while allowing port development to continue. The CAAP specified control measures for every type of operating system at the ports. This includes ships, tugboats, yard equipment, drayage trucks, and trains. The overall program cost is approximately $2 billion, divided between the ports themselves and their tenants. The original CAAP was focused on the near-term, five-year planning window from fiscal year 2006 through 2011. The 2010 CAAP Update is a new, improved version of the CAAP, providing near-term planning through 2014 and establishing long-term goals.

Port Technology Advancement Program (TAP)

The TAP was jointly funded by the Ports of Los Angeles and Long Beach, under the Clean Air Action Plan (CAAP), mentioned above. It supports the regulatory approval of promising air emissions reduction technologies. Both ports have been contributing $1.5 million per year since 2006.

There are four fundamental areas in which the program focused its initial work:

- Specific control measure requirements
- "Green-Container" Transport Systems
- Emerging Technology Testing
- Emissions Inventory Improvements

The primary focus for this program will be to reduce emissions of DPM, NOx and SOx, consistent with the Clean Air Action Plan, and ultimately Green House Gas (GHG) emissions and ultrafine particles as well.

The TAP provides partial funding for many technology research projects, usually with industry or regulatory agency partners. The projects below have, with one exception, been financially supported by the TAP program.

- *Ocean-going Vessels, APL Singapore Vessel Retrofit*. Under the Technology Advancement Program, the ports have completed participation in a three-year demonstration of emission reduction technologies aboard the container ship APL Singapore. The APL Singapore, which can carry the equivalent of 5,100 20-foot containers, travels monthly to the San Pedro Bay and Oakland Ports from ports in China, Japan, Korea and Taiwan. Two emission control technologies were demonstrated - and water-emulsified bunker fuel using an innovative onboard water-in-fuel emulsifier, and the use of Slide Valves in lieu of mini sac fuel injectors for the vessel main engines.
- *Harbor Craft, Foss Maritime Diesel/Electric Tug*. Foss Maritime, a tugboat operations firm, achieved several significant milestones in the development of the World's first diesel electric hybrid tugboat. Christened the Carolyn Dorothy, the FOSS Green Assist™ hybrid tug now offers performance comparable to a conventional Dolphin Class tugboat, but with an anticipated exhaust emissions and fuel consumption reduction up to 44 percent lower than a conventional vessel.
- *Cargo Handling Equipment, LNG Yard Hostler*. In a parallel technology development program to the TAP, Sound Energy Solutions (SES) and International Transportation Service, Inc. (ITS) jointly developed a one-year demonstration project to test the operations of three LNG-powered yard hostlers at the ITS container terminal in the Port of Long Beach, California. Yard hostlers (also known as yard tractors, terminal tractors, or utility tractor rigs) are common at port terminals, rail yards, and distribution centers. Their function is to move containers around the facility. At a port, containers are loaded off a ship onto a bobtail rig that is pulled by the yard hostler to an intermodal point or to a storage facility. Yard hostlers often sit idling as they wait in queues to pick up or drop off their loads.

  Currently, there are approximately 3,000 heavy-duty LNG-powered trucks, buses, and vehicles operating throughout California and other parts of the nation. LNG engines produce 93% lower nitrogen oxide (NOx) and particulate matter (PM) emissions, 100 percent less sulfur dioxide (SOx) emissions, and 20% fewer greenhouse gas emissions than their traditional diesel counterparts.
- *Cargo Handling Equipment, Hybrid Yard Hostler*. As a follow on to the SES/ITS demonstration of LNG in yard tractors operating at the ports, the TAP investigated the feasibility and commercial viability of using advanced technology drive systems in

cargo handling equipment. The ports' TAP, in partnership with the US EPA's West Coast Collaborative, worked together to develop and test hybrid technology yard tractors for use at container terminals. The ports of Long Beach and Los Angeles partnered with CALSTART, a non-profit company that focuses on advancing cleaner technologies, to manage a project demonstrating three (3) diesel-hybrid yard hostlers at the Long Beach Container Terminal (LBCT). The U.S. EPA also provided funds for the design and development of the hybrid drive system. Vehicle emissions and performance were evaluated relative to diesel yard hostlers, and a cost benefit assessment performed to determine the financial viability for hybrid yard hostlers when used in a marine terminal role.

- *Cargo Handling Equipment, Vycon REGEN System*. The TAP program funded VYCON Energy Corporation to apply their advanced REGEN flywheel energy storage system to container handling cranes. Vycon is an innovator in the design and manufacture of technologically advanced flywheel energy storage systems which enable a highly reliable, cost-effective and "green" energy storage solution for a variety of applications. The REGEN flywheel systems, used in container cargo handling crane applications, store power generated by the lowering of a container from a ship, and use it in the next lift of the container. This reduces power and energy costs to port operators as well as provide a reduction in green house gases.
- *Locomotives, Pacific Harbor Line Locomotive Retrofit*. This TAP-funded project demonstrated the effectiveness and durability of DPFs as a strategy to reduce diesel particulate matter from switch locomotives operating at the ports. Under this project, a MobiClean™ active regeneration DPF was installed on a Pacific Harbor Line switch locomotive. Pacific Harbor Line is the exclusive provider of rail switching services at the ports of Long Beach and Los Angeles. For this TAP project, a Tier 2-compliant PHL switch locomotive was retrofitted with a MobiClean™ Active Regeneration Diesel Particulate Filter. This project was the first demonstration of a DPF in the U.S. on a switch locomotive with a four cycle engine. This technology application is reducing PM by approximately 90 percent. Citizens living near the Ports of Los Angeles and Long Beach, as well as along the Alameda Corridor are benefiting from these emission reductions.
- *Trucks, Balqon Corporation Electric Class 8 Truck*. The Balqon E-30 Electric Terminal Tractor was built as a demonstration vehicle, which was co-funded by the Port of Los Angeles and South Coast AQMD, and designed specifically for drayage operations. Developed by the Balqon Corporation as a Port of Los Angeles initiative, the prototype E-30 all-electric terminal tractor successfully completed cargo terminal tests during 2008. As a result, the Los Angeles Board of Harbor Commissioners approved the purchase of 20 Balqon electric trucks as part of the Port's "Green Terminal" program.
- *Trucks, Westport GX LNG Engine Development*. Westport Innovations (Westport), developer of the High Pressure/Direct Injection (HPDI) liquefied natural as (LNG) fuel system technology, is developing an LNG 15-liter heavy-duty truck engine hat will meet the 2010 on-road NOx emission standard of 0.2 grams per brake horsepower-hour (g/bhp-hr). The 400- and 450-horsepower rated heavy-duty engines are based on the 15-liter Cummins ISX diesel engine platform and are designed to satisfy the performance requirements of class 8 tractors that provide drayage service at the ports.

Port Technology Support Organizations

In support of the technology demonstration and deployment programs referenced above, three programs were established by the Los Angeles port private and university leaders. These are the following:

- *Port Tech L.A.* William Lyte of the NCRST project team helped to lead the startup of a port technology incubator at the Port of Los Angeles. This incubator, Port Tech L.A., is now in operation, providing advisory services to the early stage technology companies of the Ports of Los Angeles and Long Beach. The Port Tech L.A. incubator, and overall technology initiative, is modeled upon an earlier activity which Mr. Lyte led in Pasadena, California, in association with the California Institute of Technology and NASA Jet Propulsion Laboratory. This was the establishment of the L.A. County Business Technology Center, a highly successful technology incubator, chaired by William Lyte from its startup. The BTC, which was the model for Port Tech L.A., has generated more than 50 companies with collective invested capital of $100+ million, and generation of 1.000+ jobs.
- *Harbor Association "Tech Mixer" Series.* In association with Port Tech L.A., the Harbor Association of Industry and Commerce, in which Mr. Lyte chairs a technology committee, hosts quarterly "tech mixers" addressing port technology issues. The intent of the overall Tech Mixer program is to link approximately $4 billion of recently approved capital projects with the innovative port technologies which are required for their construction. Recent Tech Mixers have featured advanced technologies for reduction of pollution from railroads, trucks, ships. The program has also showcased renewable energy systems for ports, and approaches to treatment of ballast water from ships, which can carry foreign organisms. The NCRST program was showcased in a December, 2009 "tech mixer, in association with a panel from NASA JPL.
- *California State University Regional Technology Center*. In addition, Mr. Lyte and colleagues in the port technology sector helped to establish a California State University Regional Technology Center. This virtual incubator worked closely with CSULB's Engineering School to identify promising technologies for port and intermodal applications. The Regional Technology Center is now a gateway for promising entrepreneurs from the CSU Long Beach community to meet with the port industrial sector.

MeTrIS as a Commercial Offering

The term MeTrIS, a service mark of DGRC, refers to (a) a vision as outlined in Chapter 1, (b) a suite of hardware, data and methodologies from tracking to data analysis, and (c) a broad set of research and commercial activities that implement the vision in general or specialized contexts. The process models for empty container management and synchronization, developed for port-related truck operations in this project, are an example of a contextual layer.

The commercial prospects of the technology can be considered at three levels:

Level 1. In the course of this research project, to provide an incentive to motor carriers to participate in the study, DGRC implemented a rudimentary reporting service. While larger carriers already had their own tracking systems and had no need for routine vehicle location reports, small firms found the logs valuable. Offering equivalent services for a fee, in competition with others, and achieving financial sustainability, is where the commercial challenge lies in future. Supporting fleet management software components have to be developed, e.g. routing and scheduling, and management of performance-based driver salaries. Comprehensive business planning, in terms of the elements listed on page 75, is an essential part of this business development process.

Level 2. MeTrIS is much more than a tracking service. It has to evolve into a metropolitan aggregator of data, rather than a provider of tracking services in an increasingly crowded market. The broader business challenge is ultimately to establish a customer base—trucking firms as well as port authorities, marine and rail terminal operators and other stakeholders—who have benefits to gain and a concomitant willingness to pay, for services that use the operations of the drayage fleet (a) to measure port and drayage productivity, (b) to identify congestion bottlenecks, and (c) to relieve the bottlenecks.

Level 3. The highest level challenge following from this project is to establish the institutional relationships and financial incentives that would permit say a trucking firm to forego revenues hauling empty containers, so as to achieve environmental benefits (in the broadest sense: congestion, consumption and emissions) that accrue to the community and the industry. This is not entirely implausible; it just takes a wider and more complex support system to deliver it. The examples cited in the previous section, of businesses promoting environmentally sustainable practices, indicate that some support may be found from cities and other agencies such as air resource management boards. As a business proposition, there are risks, the returns are not immediately obvious, and it would take creative schemes such as advertising or trading in carbon credits to craft this activity set into a profitable business operation.

There are questions concerning the organizational form and business model most appropriate to MeTrIS. The Traffic Management Center (TMC), best known in Intelligent Transportation Systems (ITS), is one model. It could be privately or publicly owned. As a large data repository with a constant flow of analytical challenges, MeTrIS may also lend itself to a university or laboratory environment. Third, there is the model of a public utility. A problem with these models is that they are rooted in traditional physical concepts of instrumentation, staffing and support. MeTrIS is in fact remarkably undemanding in physical and organizational infrastructure, and there is no urgency to adopt any of the above models. The imperatives for organizational expansion will depend on the scope of services.

At the time of writing this report, in December 2010, DGRC had developed a preliminary business model and launched a commercial tracking service. About half the firms that partnered in this study subscribed immediately, indicating exceptionally strong support in the industry for the goals this study set out to achieve, and reflecting enormous goodwill over the motivation and achievements of the project to date. DGRC is also in discussion with a consortium of port freight stakeholders to provide a regular flow of performance data on truck interactions with marine terminals.

Conclusions, Recommendations and Further Work

RITA's goal in funding this project was to have technology deployed in ports to mitigate freight congestion. The project ran an active outreach and recruitment program, deployed the technology immediately in the course of the research, and has already started to commercialize it, past "Level 1" as described above, to Level 2. Successful commercialization unmistakably establishes the value and sustainability of the technology.

At Levels 1 and 2, customer costs are low, hence the risk factor in adopting a new technology is acceptable, even if the benefits are lower than expected. The greatest public benefits—and risks—lie at Level 3. This is not an area that is easily developed commercially. Benefits are not to any single private sector customer, but to communities and a variety of public agencies mandated to protect the interests of those communities. Success requires participation and goodwill from all stakeholders, hence the risk factor is much higher. Extensive consultation and even negotiation are required, and businesses have to be persuaded and incentivized to take actions that would otherwise not be in their short-term interest. This is an obvious area where funding by RITA, perhaps in cooperation with other federal agencies (e.g. MARAD) and state and local agencies (e.g. Caltrans), could enable an activity that promises long-term public benefits.

From a commercialization standpoint, it would be valuable for RITA to survey ports and intermodal industry groups to identify operational areas requiring new technology. This would then point the way for targeted research. RITA could develop outreach programs to these local port sectors within which successful projects could be showcased to further commercialization success. There are synergies between the utilization of the MeTrIS and other RITA project technologies, and the reauthorization of the federal Transportation Bill. For example, technologies developed with RITA funding should be particularly applicable to the needs of the national goods movement sector, and can be highlighted as such in the Transportation Bill.

On the matter of outreach, a final point concerns the web site *ncrst.org*, that serves as a public portal to this consortium as well as other NCRST sites funded by RITA. The consortium took over the domain and content administration of *ncrst.org* in 2009 (which had been static since about 2001), and updated it to reflect new awardees. The site now points to web pages of current NCRST institutions, if such sites exist, or hosts the final research reports of consortia that have not continued to maintain their own sites. This ensures that research results continue to be made available to the transportation community. There is considerable value in synthesizing the research of the RITA-funded consortia. The common web site helps to forge a common identity, but it needs to be supplemented by technical exchanges, synthesis publications, and outreach events of common interest. Remote sensing and GPS have evolved into mainstream research areas in transportation over the past decade. The NCRST program played no small role in this process, and this work stands to gain considerably by being orchestrated and showcased by RITA.

7—Conclusions and Recommendations

A large port is a complex, competitive and at times fierce business environment. The magnitude of trade passing through its gates, the time-sensitivity of goods delivery and vessel handling, and the mix of private and public interests and administrative structures—public port authorities, private marine terminal operators (often foreign-owned), owner-operator and employee truckers and unionized longshoremen—together create a setting in which it is highly challenging to bring about change.

This project did not set out to take on that challenge to *deliver* change. Instead it sought to create or to package sensing technologies as appropriate, to document the dynamics of the goods movement industry with respect to port operations, to identify the most pressing congestion concerns, to identify specific cases in which congestion could be mitigated, to evaluate the benefits of specific mitigation measures, and to support the decision-making process by evaluating multiple alternative scenarios.

By necessity, the technical definitions and proposals simplified and abstracted the problems, and skirted some institutional issues. But in many respects, and to the extent possible, the models did consider implementation issues, for example evaluating different levels of subscription. Difficult processes of decision support and negotiation remain to be undertaken to bring about the transformations we envisage.

Accomplishments

The consortium brought together academics in geography and engineering, business-oriented consultants and industry partners. The accomplishments of the team accordingly varied, ranging from methodological inventions, discoveries and academic papers to services, contact networks, goodwill and awareness. Most notably:

- *Port data model.* A data model was created to capture elements of interest to port operations. The model builds on an existing industry-standard, UNETRANS, developing the port context layer in detail.
- *Conflation*. A set of methods was developed in the general area of conflation, which is the problem of relating two geographic data sets to each other, e.g. GPS and underlying geospatial layers. The National Geospatial-Intelligence Agency funded further development of the techniques.
- *Instrumentation*. About 250 trucks from a dozen motor carriers were instrumented with GPS receiver-transmitters, providing a detailed real-time data flow on port truck movements.
- *Mapping and Analysis*. Numerous information products were created, representing the pattern of freight flows out of the ports and to intermodal yards and warehouses throughout the LA basin. Our measurement of wait time outside marine terminals was of particular interest to the industry, and at least one terminal expanded its land bridge service considerably in response to MeTrIS queue time observations.
- *Empties Management Model.* Simulation and optimization models were created to plan ideal locations for empty storage yards in the LA basin. The scheme avoids 4,500

truck entries into port terminals, greatly reducing congestion, and reduces CO_2 emissions by 50 tons per day.

- *Port Synchronization Model*. Simulation and optimization models were created to study the impact of GPS-based arrival information made available to marine terminal operators in advance of truck arrivals at terminal gates. The models predict an improvement of 15% in the speed of container stack handling.
- *Commercialization*. Potential for deployment and commercialization of tracking and optimization technologies was actively explored throughout the project, and members of the research team launched a commercial service to maintain the continuity of the data flows established under the project.

Challenges

The potential benefits of the measures proposed are enormous, detailed in the statistics quoted above. It is useful to consider where these proposals, despite their obvious merits, might face roadblocks. Some of the challenges posed to the research team are:

- A Virtual Container Yard (VCY) was implemented earlier in 2006, and proved unsuccessful. The ports expended effort and expense to implement the it. While there has been no objective analysis of the reasons for failure, port officials are reluctant to consider similar proposals. There are important differences between the ESY and the VCY, but a common factor is that motor carriers lose the revenue they might make moving empties. They need to be compensated for this. To succeed, the ESY must be supported by a system of incentives and disincentives: per diem charges on containers must be reduced, and a disincentive charge levied on the transportation of empties from portside terminals, with corresponding incentives for turning around empties outside the port.
- Marine terminals are concerned about the re-tooling required to synchronize terminal operations with truck arrivals, the training of crane operators to process the new information (on real-time truck arrivals) effectively, and the amount of work involved in stack sorting. One marine terminal executive claimed that benefits would accrue to truckers, not to terminals. Of course, this is untrue. A 15% improvement in stack handling velocity reduces queues and expands capacity and the potential for new business.
- There is fear that the longshoremen's union will oppose technological innovation, whether or not automation costs jobs. Our conversation with a senior labor union executive does not support this. The downturn of 2009 exacted a severe toll on longshoremen jobs, and unions are concerned that failure to achieve efficiencies will drive business to east coast ports through the Panama Canal, or to other west coast ports eager for the business. The union executive welcomed the prospect of MeTrIS technologies.

Recommendations

1. The first and most important recommendation arising from this study echoes the views of federal, state, local and industry officials familiar with MeTrIS. They strongly urge its continuation, to keep up the flow of data on the operational health of the San Pedro

ports. Clearly the program cannot be perpetuated at federal expense. Consortium members have already taken steps to ramp the offerings into a commercial service, and motor carriers have responded enthusiastically.

2. Marine terminals and motor carriers should work together, in cooperation with ports and beneficial cargo owners, to agree on appropriate metrics of "turn time" and strategies to reduce it. In the month following the end of this project period, this is precisely what transpired. A Turn Time Stakeholders Group (TTSG) was constituted, with representation from these organizations. Consortium members are in consultation with TTSG to provide a regular stream of data that can measure the effectiveness of remedial methods such as extended working hours at port terminals.

3. The range of data gathered by MeTrIS should expand, to include at least rudimentary payload data, i.e. bobtail, chassis, empty or full container. This provides a better understanding of the movement of chassis and empties around the basin, and can feed better models to minimize unproductive travel.

4. MeTrIS should document and estimate travel speed on port-bound freeways, to predict the precise arrival time of trucks at marine terminals, enabling implementation of truck-terminal synchronization. The truck driver's intentions need to be documented: the fact that a truck is heading in the direction of the port does not mean that it intends to arrive at a particular terminal or to collect a particular load. This uncertainty needs to be resolved to improve the value of the data stream.

5. A dialogue, perhaps in the form of a conference, should be held to address the potential of strategic planning measures advocated in this research, specifically the establishment of ESYs and incentives to use them:

   - Determine optimal location and capacity for ESYs, using scenario evaluation and decision support tools such as those developed for this study.
   - Institution of disincentives for removal of containers from marine terminals, coupled with incentives for using ESYs. Incentives and disincentives should be revenue-neutral.
   - Steamship lines should increase inventory of containers, and drop per-diem charges, to encourage efficient handling and short-term storage. This constitutes a cost to the lines, and appropriate incentives should be offered.

   The dialogue would identify the concerns of all stakeholders regarding each of these and other issues, assess the support for these measures, and arrive at strategies for implementing them.

6. Separately, MeTrIS needs to engage marine terminals and motor carriers, perhaps with the involvement of TTSG, to implement the synchronization proposals advanced by this research effort. Benefits of the proposals are large, and costs are minimal, consisting principally of minor enhancements to information systems already in place at many terminals.

7. Given the magnitude of logistical challenges at the San Pedro ports, serious consideration should be given to radical changes in the process of container pickup. One model is the "taxi service" analogy in which containers are stacked by delivery priority, trucks are members of a single resource pool, and a truck picks up the next

available container and delivers it to its destination or to an intermediate off-port flip yard. This mechanism is already in use in the form of “land bridges,” and is under consideration in the case of the Clean Trucks Coalition. Some terminals are contemplating or in the process of constructing container racks, which are an alternative to stacks. These changes may obviate the synchronization benefits proposed in this research, partially or entirely, but they may well be the only way San Pedro can cope with future traffic demands.

8. RITA should follow through on the role it has played in funding the development and delivery of these solutions in the San Pedro ports. It can assist in the promotion of the solutions, and in cooperation with MARAD and FHWA, it can identify additional measures to facilitate freight movement in the vicinity of this critical national facility.

Next Steps

This project was completed well under budget, with roughly 8% of project funds or $150,000 unspent at the time of closure. Further work is contemplated, to engage the port community more closely in the tuning of models and decision support systems, and to advance all of the above recommendations to practice. The timetable and the nature of the mission differ in some respects from those of the current research project, and it is appropriate that these tasks be the subject of a new and separate engagement.