DATA ANALYSIS AND PROCESSING TOOL FOR AUTOMATIC IDENTIFICATION SYSTEM (AIS) VESSEL DATA TO SUPPORT SIMULATION MODELING

Cory M. Rogers[a], Keith D. Hofseth[b], Richard M. Males[c]

[a]CDM, Carbondale, IL, USA
[b]U.S. Army Corps of Engineers, Institute for Water Resources, Alexandria, VA, USA
[c]RMM Technical Services, Inc.

[a]rogerscm@cdm.com, [b] keith.d.hofseth@usace.army.mil, [c]males@iac.net

ABSRACT

Automated Identification System (AIS) data provides ship tracking information for cargo vessels in coastal waters. Information includes vessel position, speed, draught, dimensions, and destination port. This information is updated on an hourly basis. The analysis of AIS data for individual vessels and for a given port provides a large amount of useful information on vessel routes and loading that can be used in conjunction with port simulation models. The Institute for Water Resources of the US Army Corps of Engineers is developing the Automatic Identification System - Data Analysis and Pre-Processor (AIS-DAPP) tool for visualization and analysis that can be used to support simulation modeling and planning activities. The AIS-DAPP applies genome sequencing practices to identify services from the data. It also works seamlessly with Google Earth™ for spatial visualization and incorporates analytical routines to determine common routes, speed within channels, and draught entering and leaving a port, all needed for simulation modeling.

Keywords: port simulation, automated identification system

1. INTRODUCTION

The U.S. Army Corps of Engineers is responsible for maintaining the navigable waterways of the United States, including ports and harbors. Improvements to the ports that are under the jurisdiction of the Corps must be economically justified based on analysis of the relative benefits to navigation (reduced transportation cost) and the cost of the improvements (P&G 1983).

The analysis of economic justification is done using a "without project" condition, i.e., the future configuration of the harbor if no improvement project is carried out, as compared to a "with project" condition, in which vessel traffic is projected based on the assumption of the new improvements being in place. Benefits resulting from such improvements are often dependent upon their effect on reducing transportation cost, for example deepening a channel allows deeper drafting vessels to service the port.

2. PROBLEM SETTING

Container vessels are not characterized by "there and back again" routes typical of bulk cargo vessels. Container vessels are deployed on liner service routes which typically have 5-15 ports. The sailing draft, and thus channel depth utilization, at each port can vary widely. This presents a difficult analysis for Corps' planners who are asked "what will be the economic impact of deepening the channel at port x"? As port X is only one of several ports on the service route the answer is highly dependent on the characteristics of the other ports on the service and the nature of the trade on that service. Figure 1 shows the Maersk Columbus Loop Service.

Figure 1: Maersk Columbus Loop Service

The Corps is currently evaluating the effects of deepening the channel at Savannah harbor. The analyst is asked to estimate incremental changes from deepening depths ranging from 2 to 6 feet. This requires the analyst to predict fleet deployment changes and changes in at the port vessel drafts for each increment of depth.

To aid the Corps analyst the IWR has developed an evolving set of statistical, forecasting and simulation tools for analysis of vessel movement in harbors. The AIS-DAPP provides the capability to visualize, analyze and synthesize historical AIS data for use in container port channel improvement studies and associated simulation modeling. Specifically, the AIS-DAPP was developed to aid the evaluation of channel improvements at U.S. container ports and to help populate IWR's HarborSym simulation model (Hofseth, Heisey, Males and Rogers, 2006).

3. LITERATURE REVIEW

The availability of AIS data is a major breakthrough for maritime transportation simulation systems. Shipping companies tend not to share information with the potential of rendering business practices that may give an edge to a competitor or expose what they may consider a trade secret (Greaves and Figliozzi, 2008). Previous work has been executed on land based vehicle position streams similar to the AIS data. Many of these studies have a component similar to an element of the author's current research efforts – Breaking the positional stream into movement segments bounded by stopping points. Stopping point identification from GPS streams is currently founded on a criteria expression parameterized on dwell time. Other criteria in the flagging expression using GPS data include heading changes and zero velocity (Mathew and Reddy, 2008). It should be made clear that these research papers involved GPS data streams of land based motor vehicles and that the trips and movement segments are much smaller than that of ocean going vessels. Identification of actual shipping routes, route usage rates, and vessel characteristics is very important to planners charged with optimizing maritime ports of call. The elements to be optimized can change based on stakeholder involvement and current constituent concerns. In a recent work, the European Space Agency released maps of actual shipping routes obtained through near real time ship identification by an orbital radar system overlaid with SO2 and NO2 maps (ESA, 2009). The vessel route information that was previously unavailable can show causality in environmental issues as well as provide for many other potential optimization targets. With respect to quality of human entered data in the AIS data stream, Bailey, et al. show a non-trivial error rate attributed to several factors including improper or lack of training (Bailey et al, 2008). This finding lead to the implementation of computed trip segmentation in the analytic tool.

4. CHARACTER OF AIS DATA

The IWR has acquired AIS data on worldwide container ship movements for some 3000 vessels during the period 2007-2008 from a commercial provider. This information is based on coastal stations, thus at sea movements out of range of land-based stations are not currently available, although at sea data is anticipated to be available for AIS data in the near future. The primary data consists of "ping" records (8.5 million in total) giving information on:

- Unique Vessel Identifier
- Movement Date and Time
- Latitude and Longitude
- Destination Port
- Ship Name
- Vessel Characteristics (Type, Length, Beam)
- Draught
- Speed

Sample data extracted from the large data set for a single vessel is shown in Table 1.

Table 1: Sample AIS Data (Single Vessel)

Sample AIS "Ping" Data (selected)					
Longitude	Latitude	Move Date	Destination	Draught	Speed
139.677515	35.4005667	1/5/07 3:55	Yokohama	12	0.1
139.677551	35.4005450	1/5/07 4:53	Yokohama	12	0.1
139.677528	35.4005700	1/5/075:55	Yokohama	12	0.1
139.677563	35.4005933	1/5/07 6:58	Yokohama	12	0.1
139.677658	35.4003567	1/5/07 7:59	Shanghai	10.4	1
139.758910	35.2709000	1/5/07 8:59	Shanghai	10.4	13.6
139.727796	35.0962800	1/5/07 9:48	Shanghai	10.4	18.5
121.734338	31.3160067	1/8/07 7:59	Shanghai	10.1	9.9
121.647738	31.3411417	1/8/07 8:52	Shanghai	10.1	0.1
121.647720	31.3411633	1/8/07 9:55	Shanghai	10.1	0
121.647766	31.3411367	1/8/07 10:55	Shanghai	10.1	0
121.647766	31.3411367	1/8/07 11:29	Shanghai	10.1	0
121.647723	31.3411050	1/8/07 12:52	Shanghai	10.1	0
121.647933	31.3411417	1/8/07 13:58	Shanghai	10.1	0
121.647840	31.3412000	1/8/07 14:42	Shanghai	10.1	0
121.648025	31.3415150	1/8/07 15:59	Yantian	10.3	0.3
121.726748	31.3169633	1/8/07 16:46	Yantian	10.3	11.7
121.987268	31.2479250	1/8/07 17:48	Yantian	10.3	13.4
122.058365	31.2367500	1/8/07 18:04	Yantian	10.3	13.8
118.318770	23.8947467	1/9/07 18:48	Yantian	10.3	24.3

5. AIS-DAPP OVERVIEW

The AIS-DAPP includes an underlying spatial relational database structure in the PostgreSQL environment. A user interface, developed in the C# programming language allows a user to select and filter data by: polygon; port; route; vessel; vessel class, and time. Analysis capabilities include: the ability to identify routes and services; development of statistical information on arrival and departure draughts; calculation of statistics on the amount of time spent by vessels at docks; determination of vessel speeds at various points along the route; and calculation of trip shares by vessel class and route. Visualization capabilities, using Google Earth and other tools as needed include: plots of individual vessel routes; representations of vessel class share of movements by route; visualization of representation of sailing draft by route and/or route segments; and plots of routes dominated by a selected vessel class.

An automated reporting tool has also been created to allow close to a "one-click" capability to generate information about the vessel traffic at a selected port, generating a vessel call list for use with the HarborSym model.

6. PROCESSING AND VISUALIZATION

A number of tools have proven useful in the processing and initial examination of AIS data. The PostgreSQL database (http://www.postgresql.org/) is a free, open source relational database server capable of handling extremely large data files. The PostGIS extensions to PostgreSQL provide a complete set of spatial analysis capabilities for the locational data (http://postgis.refractions.net/). This allows for the development of queries such as "Identify all the vessels that come within 100 miles of the location of the Port of

Yokohama", and allowing the database to serve directly as a data source for GIS platforms such as the Quantum GIS (http://www.qgis.org/). Google Earth™ (http://earth.google.com/) provides visualization capabilities for vessel data points and tracks, once information has been formatted into the required KML file format.

An example of AIS data visualization of vessel pings for a single vessel for a few days in January 2008 using Google Earth is shown in Figure 2, and larger scale movements for the same vessel, over the course of the two years of data, are shown in Figure 3.
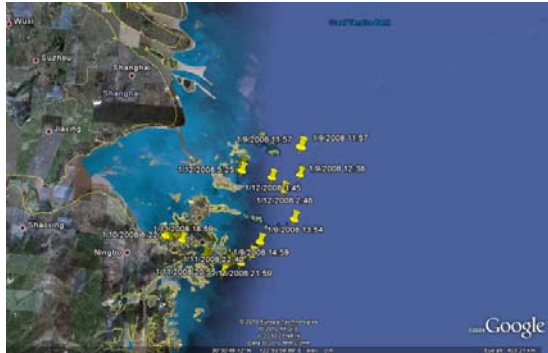

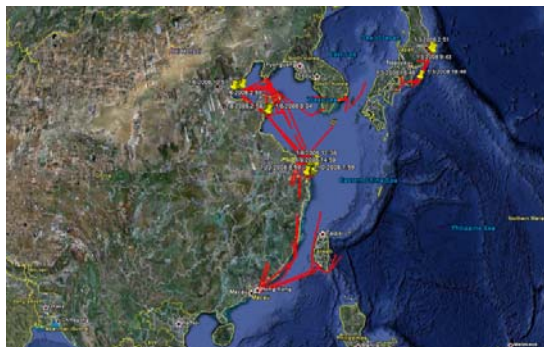Figure 2: Single Vessel Ping Visualization, Eastern China Sea

Figure 3 : Vessel Movements in Asia (single vessel, 2007-2008)

As seen in Table 1, each ping has an associated destination. A simple Python script can be used to extract the ports visited in order, as shown in Table 2.

7. TRIP SEGMENTION

AIS data, like most raw data sets includes a number of errors and omissions. Raw AIS data must be cleaned before it can be ported into the AIS-DAPP. An automated tool for cleaning the data was not developed. However, the IWR has developed detailed step-by-step instructions to transform the raw data to be compatible with the AIS-DAPP. This allows additional data that becomes available with the passage of time to be incorporated into the tool.

Table 2 : Port Visit Sample

Port Visits
Yokohama
Shanghai
Yantian
Balboa
Charleston
Norfolk
Newark
Bremerhaven
Felixstowe
Rotterdam
Le Havre
Newark
Charleston
Miami
Panama Canal
Balboa
Los Angeles

The first stage of AIS processing is to transform the positional data stream into a sequence of movements separated by points at which the vessel stops. The process of stream segmentation is a multi-step procedure consisting of stop identification, stop compression, and extraneous movement removal. To facilitate this procedure, additional data elements are introduced to the AIS data structure. In particular, three Boolean flags – Stop_pt, which is used to indicate an identified vessel stop, dest_arrival, which flags a stop as an arrival at some location, and dest_departure, which indicates the final stop point prior to a movement.

The algorithm employed for stopping point detection is a variant of the dwell time method used in Stopher's research (Stopher, et al, 2003). The data is ordered by the unique vessel ID and then by movement date/time. By default, all records are marked as non-stops. The ordered stream is then iterated over making a comparison between the current record and N subsequent records. The implemented algorithm identifies a stop as a movement ping record for a specified vessel in which neither latitude nor longitude has changed by 0.01 degrees and the vessel's velocity is less than 0.5. With a data time granularity of 60 minutes, a value of $N=1$ was selected to allow the detection of all vessel stops whose dwell time is greater than 1 hour 59 minutes. Examination of the raw data revealed a potential flaw in that the data gathered was not presented at a uniform sampling of 1 hour. The solution to the data gaps was to modify the algorithm to check for a standardized distance over time value. In the original test, the time element was ignored under the assumption of uniform samples. The modification determines degrees per hour moved and triggers if that value is under 0.01 degrees per hour. Analysis is currently underway to determine if the use of average velocity (degrees per hour) in addition to instantaneous reported velocity produces more accurate stop detection.

The second phase of the trip segmentation process is to identify the end points of consecutive stop records to delineate and form the actual stops. Once again, the movement records for each vessel are iterated over after being sorted by date/timestamp. If a record is marked as a stop and the previous record is not, the record is marked as an arrival. Departure records are identified by the opposite condition. Special consideration is given to the first and last records. If the first or last

movement record for a vessel is identified as a stop record, a special flag is raised indicating that the actual length of time the vessel at the location is not known. If the principal or ultimate movement records are non-stop records, no special processing is required. The final phase of the trip segmentation process is data compression. The records between paired arrival and departure records are deleted as are all records not identified as stops. The result of this process is a chronologically ordered list of stops that the vessel made.

8. PORT IDENTIFICATION

While the AIS data contains a field for a vessel's current destination, this data is potentially misleading because although the vessel is physically at a particular port, the moment the vessel begins to move the destination field reports its next destination. Instead a geo-spatial process is applied to determine the identity of the stop location. The system maintains a table of stopping locations; these locations are either of type port or anchorage.

Applicable to the process of port identification of a stop is the latitude and longitude of the port along with a port specific radius of influence. The radius of influence is set using the Movement and Port Explorer portion of the analysis tool seen in Figure 4.

Figure 4: Port Identification Radius

The Movement and Port Explorer plots selected AIS records and defined ports along with basic geographic boundaries to assist with the analyst's geographic orientation. For each new port, the analyst can see the clusters of candidate AIS records. Using the measuring tool, a radius of influence can be determined. The port information is directly updated using the tool. The establishment of the radius is a non-trivial task when there are multiple ports within similar distances from a cluster of AIS records. Both overlapping radii and those that do not sufficiently cover one or more clusters introduce questionable results from the port identification process. Figure 5 shows the ports of San Francisco and Oakland with the associated AIS cluster's denoted.

Figure 5: Radii for Ports of San Francisco and Oakland

The next step in the port identification process is the automated assignment of one or more existing ports to each AIS record when possible. This is a three phase process. The first phase is to create a Cartesian cross product between the AIS records and the defined ports where the AIS record is within the port's radius of influence. In addition, the distance from the port to the vessel is recorded. The result of this phase is a list AIS records with one or more port IDs sorted by vessel ID, MovementID, and distance to port. The list is then iterated over retaining only the first record of each set. From this list, the destination ID is set in the original segmented trip list selecting the closest port if more than one claims influence over the vessel stop. The second phase is to return to the Movement and Port Explorer with the dataset restricted to those AIS records without a valid destination ID. The analyst has the option to add new destinations with associated radii of influence to provide suitably identified destination information. These two operations are repeated until the analyst is satisfied with the coverage of identified destinations. The final phase of port identification is to remove stops that remain unidentified as they provide no relevant data.

9. REGIONS AND MULTI-LEVEL ANALYSIS

Level of detail is extremely important in exposing relevant trends in data to an analyst. Volumes of highly detailed data can mask trends that appear at higher levels. While the direct purpose of this analytical tool is to analyze vessel call trends, the nature of region definition directly affects the aggregation of the movement data and thus the level of trend visibility.

A region is defined as a collection of one or more distinct destinations or ports and the complete region list for this effort is shown in Table 3. It is important to note that a port or destination may be in, at most, one region for the aggregation to be valid. In the analysis performed, the desire was to explore vessel call trends between several geographic region in the United States and larger international regions.

Table 3: Region Names and Codes

Region Name	Region Short	Region Code
Africa	AFR	A
Black Sea	BLS	B
Caribbean / Gulf	CAR	C
Canada East Coast	CEC	D
Canada West Coast	CWC	E
East Coast South America	ECSA	F
Europe	EUR	G
US Gulf Coast	GCUS	H
Mediterranean	MED	I
Mid East	MID	J
North Asia	NAS	K
Oceana	OCE	L
Panama Canal	PAN	M
South Asia	SAC	N
South China	SC	O
Southeast Asia	SEA	P
Suez	SUEX	Q
US East Coast	ECUS	R
West Coast South America	WCSA	S
US West Coast	WCUS	T

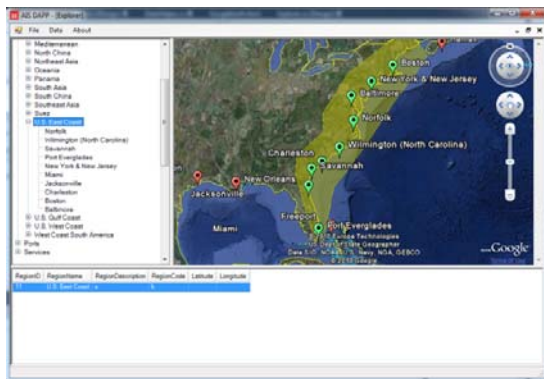Figure 6 shows the exploration view of the region defined as East Coast US.

Figure 6: View of East Coast U.S. Regions

Further processing to look for patterns and development of larger spatial scale service definitions (e.g. Asia – Panama Canal – East Coast US - Europe – East Coast US – Panama Canal – West Coast US – Asia) is then carried out, together with examination of port arrival and departure draughts as indicators of vessel loading and commodity import and export.

10. SERVICE IDENTIFICATION – GENOME SEQUENCING

Genome sequencing practices were applied in the processing of AIS data. Genome sequencing is quite simply, the process of determining the order of genetic base pairs that form an organism's genetic code. In genome sequencing the functional alphabet is a simple set of four nucleotides – adenine, guanine, thymine, and cytosine. These nucleotides are coded as 4 single letters – A, G, T, and C. This simple step of nucleotide encoding serves to both simplify the representation of a genome sequence without data loss and to enhance the ability to identify patterns. An example of a DNA sequence is "ATTCGCATT". This representation of the sequence is far more compact and "observable" than fully listing the nucleotide pairs. Quick inspection of this sequence yields that there are several repeated sequences – "AT", "ATT", and "TT". This trivial observation also shows that the sequence is bounded by the "ATT" sequence and that the other two repeated sequences are possibly related to the "ATT" marker. These simple observations serve to highlight a benefit of genome sequencing – pattern and positional observation in a potentially noisy field. An example of this, paraphrased from Genome News Network (J. Craig Venter Institute, 2004) – take a phrase, remove all punctuation, remove all white space, and add random letters between the original words:

ADRYYGPRAESENSULROVAWSDCRASIPLPULL ISSUNTGGEMELIORAWYTT

By identifying sequences or words that we know this becomes:

ADRYYG**PRAESENS**ULR**OVA**WSD**CRAS**IPL**PUL LISSUNT**GGE**MELIORA**WYTT

Extraction of these identified words yields a sentence that still makes little sense to most.

AD PRAESENS OVA CRAS PULLIS SUNT MELIORA

Once the language of the sentence is understood (Latin), the true meaning is determined; "A bird in the hand is worth two in the bush".

There are many similarities between the problem domain of genome mapping and transportation route/service identification. The port calls of a vessel are analogous to nucleotides. Vessel movement records can alternatively be thought of as DNA sequences or chromosomes. Minor aberrations in a vessel's activities as well as actual data errors introduce extraneous noise into the DNA sequence. Given the similarities in the problem domain, this analytical tool endeavors to leverage the techniques of genome mapping in order to extract similar benefits in pattern recognition, noise reduction, and eventually, linguistic understanding of shipping services.

The primary goal of the analysis is to observe patterns and trends in inter-region call lists and to produce a set of regional services through genome sequencing. Repeated patterns of inter-region vessel calls are known as services. While the produced list of services will encompass the entire data set, the

frequency of occurrence data collected will be used to show volume of movements within the genome.

The first component of this sequencing process is to obtain the DNA for each vessel in the processed AIS data. The fingerprint is in the form of a string of characters where each character represents a unique region the vessel visited in chronological order. This is done by ordering the call list for each vessel on arrival time and iterating over the resultant list. For each visit on the list, the associated region is decoded from the identified destination/port of the visit. If the decoding process yields more than one region code, an error is produced and the process stopped. If the decoding process does not find a region code, the visit is ignored. The most common result of the decoding process is one unique region code. This code is appended to the end of the string if the previous region code is not the same. This conditional appending serves to remove intra-region moves which would obscure the visibility of the inter-region movements. The result of this procedure is that each vessel in the active call list has a DNA sequence or chromosome encoding its inter-region movements over time for the analysis period.

The second component of the genome sequencing is the controlled decomposition of the chromosomes into candidate genes and gene complexes. In our problem domain a gene is a simple inter-region movement and gene complexes are a series of multiple inter-region movements. Both simple genes and gene complexes are candidates for identification as services. The data structure for this decomposition is a table with a string as the primary key which corresponds to a unique set of one or more inter-region movements as well as an integer column for accumulating the frequency of gene appearance in the dataset – the genome.

The chromosome is decomposed into candidate genes by extraction of substrings. The process starts with simple genes of length 2 and progresses up to the length of the chromosome string in full. As each gene string is extracted, the genome is checked for previous identification. If the gene string has not been previously identified, it is added to the genome with an initial occurrence count of 1. If the gene has been previously identified, the occurrence count is incremented.

The potential number of distinct genes and gene complexes in a chromosome of length n is bounded by the simple sum:

$$\sum_{i=1}^{n-1} i = \frac{n(n+1)}{2} - n = n\left(\frac{n+1}{2} - 1\right) \quad (1)$$

While this is the number of genes/gene complexes that must be examined for a given chromosome, the regional encoding will reduce the actual number of distinct entries in the genome.

The genome is then used as a source for service identification. The resultant table of services must exhibit two properties: it must be complete in that it must account for all identified inter-region movements, and the services must be mutually exclusive. Services are further defined to be cycle terminated. The rule of cycle termination defines that a service does not continue to another region after returning to a region that it has previously visited. An example of this is the following gene: "ABCAD". This gene is not a candidate for a service as it shows a visit to region D after returning to A. The genome mapping technique used will have also identified genes "ABCA" and "AD" which would pass the rule of cycle termination. All genes that fail the cycle termination rule are culled from the candidate genome table. The desired result property of completeness is not impacted by this culling due to the genome map containing the subsections of the culled genes that do pass the cycle termination rule. The second pass of gene culling is to ensure mutual exclusivity of vessel movement observations. For each gene remaining in the candidate service genome table, the process of genome mapping is reapplied with the results being stored in a temporary location. The original gene itself is removed from the temporary store. Each remaining value in the temporary store is found in the service candidate genome and the occurrence count reduced by the count value in the temporary store. The final step is to remove any gene from the service candidate genome whose occurrence count is zero as these genes are fully accounted for in more complex genes. The resulting genome is a set of services that are mutually exclusive with respect to inter-region vessel movements and fully covering the dataset of those movements. A completely mapped service between Africa and Europe is shown in Table 4.

Table 4 : An Africa to Europe Service

Port	Country	Region	Code
Casablanca	Morocco	AFR	A
Agadir	Morocco	AFR	A
Rotterdam	Netherlands	EUR	G
Bremerhaven	Germany	EUR	G
Helsingborg	Sweden	EUR	G
Antwerp	Belgium	EUR	G
Casablanca	Morocco	AFR	A

The actual chromosome found would be AGA, indicating a service that initiates in Africa, travels to Europe and returns to Africa. Because of the string nature of the genes, and our retention of the individual vessel's chromosomes, the analyst can find all vessels that share a common gene or, perhaps, vessels that have more than N occurrences of a specific gene.

11. CONCLUSION

The techniques described for identifying vessel trips, mapping of ports and illuminating vessel services have proved effective at uncovering valuable information from a dense yet incomplete data source. In particular, treating the vessel route as a genome sequence in order to programmatically uncover vessel services within the

data provided useful insights into the operations of shippers. With this information, a plethora of informative statistics can be easily calculated for further analysis, such as:

- Ocean speeds by vessel class and service
- Vessel speeds in reach at port
- Vessel arrival drafts by port and vessel class
- Vessel departure drafts by port and vessel class
- Imputed vessel capacity used by route, route segment and vessel class.
- Identification of total vessel capacity by service
- Service length
- Imputed service volume of traffic
- Port depths on service
- Share of route capacity by vessel class
- Evaluation of share over time
- Vessel class deployment by route characteristics
- Vessel deployment over time

All of these statistics are extremely useful for analysis and simulation. A list of vessels calling at a specific port with their arrival times and draughts is also produced from the system. This vessel call list is intended to be used as input to a monte-carlo simulation model for the economic evaluation of potential port improvements such as deepening a channel to allow larger vessels to call the port.

ACKNOWLEDGEMENTS

REFERENCES

P&G 1983 - Economic and Environmental Principles and Guidelines for Water and Related Land Resources Implementation Studies, March 1983

Hofseth, K.D., Heisey S., Males, R.M., Rogers, C. 2006. Development of Commodity Driven Vessel Movements for Economic Analysis of Port Improvements http://www.corpsnets.us/docs/HarborSym/06-NETS-P-02.pdf

Greaves, S. and Figliozzi, M., 2008. Urban Commercial Vehicle Tour Data Collection Using Passive GPS Technology: Issues And Potential Applications. Proceeding of the 87th Transportation Research Board Annual Meeting CD rom - Washington DC. USA

J. Craig Venter Institute 2004, http://www.genomenewsnetwork.org, (*accessed June 9, 2010*)

Mathew, S. and Reddy, YV, 2008. Traffic Study of Individual Behaviour in Tiruchirapalli City Using GPS, Available online at http://www.gisdevelopment.net/application/miscellaneous/mi08_268pf.htm (*accessed August 2010*)

ESA, 2009, ESA Map Reveals European Shipping Routes Like Never Before, available online at http://www.esa.int/esaEO/SEMBDI0OWUF_index_2.html (*accessed August 2010*)

Bailey, N., Ellis, N. and Sampson, H, 2008. Training and Technology Onboard Ship: How Seafarers Learned to Use the Shipboard Automatic Identification System (AIS). Seafarers International Research Centre, Cardiff University, UK

Stopher, R. and Jones, P. 2003, Transportation Survey Quality and Innovation, Volume 2001, Emerald Group Publishing, Bingley, UK

AUTHORS BIOGRAPHY

Cory M. Rogers
Mr. Rogers is a computer scientist currently serving as the National Practice Leader for Applications Development at Camp Dresser McKee, Inc. He has over 15 years of experience in the design and development of software applications to solve complex decision making problems. He has managed teams to develop successful software applications ranging from single-user desktop applications to web-enabled systems with a user base of thousands. Prior to joining Camp Dresser McKee, Inc, he served as the Director for Software Engineering for a small consulting firm and was a founding member of the Curriculum Advisory Board for the Information Systems Technology program at Southern Illinois University. He has worked on numerous decision support software models which are currently being used to perform crucial functions such as estimating flood damages caused by coastal storms, investigate the economic feasibility of improvements to the U.S. inland waterway and port infrastructure and manage financial risk at global food and beverage companies. Mr. Rogers holds a B.S. degree in Computer Science from Southern Illinois University.

Keith D. Hofseth
Mr. Hofseth has served as an economist at the Corp of Engineers for over 25 years; conducting studies in flood control, deep draft navigation, inland navigation and hydropower. He co-authored the first guidance for major rehabilitation analysis and was a pioneer in bringing risk and uncertainty into the Corp's planning process. Mr. Hofseth started with the Corps in the Seattle district and later served as the Chief Economist for the Alaska District. He is currently the Technical Director of the Navigation Economic Technologies

research program at the Institute for Water Resources and is a graduate of the Corps' Planners Associate Program. Mr. Hofseth holds a B.S. degree in Finance and Economics from the University of Montana and an M.A. degree in Applied Economics from Johns Hopkins University.

Richard (Dick) Males

Dr. Males has worked professionally in consulting since 1968, and is owner and principal of RMM Technical Services, Inc., an independent consulting firm providing technical and computer support activities to a wide range of government, business, and non-profit agency clients. Dr. Males holds a Ph.D. in Water Resource Systems Analysis from the Massachusetts Institute of Technology. Dr. Males' primary professional activities have been in the arenas of development and effective use of computer-based tools, with emphasis on water resource systems applications. Areas of work include: design and implementation of advanced planning methodologies; simulation modeling; risk-based analysis; multi-criterion decision-making techniques; data base, budgeting, tracking, and decision support systems; mathematical modeling; geographic information systems development and applications; computer graphics and CAD applications; and laboratory data systems. His recent work areas include object oriented simulation models for risk-based analysis of navigation (deep-draft harbors, inland waterways and locks); coastal shore protection modeling; Monte Carlo simulation models for repair/rehabilitation of hydropower and levees; and development of decision support systems.