Portland State University

PDXScholar

Geography Masters Research Papers

Geography

2014

The Cost of Hauling Timber: A Comparison of Raster- and Vector- Based Travel-Time Estimates in GIS

Sara M. Loreno
Portland State University

Follow this and additional works at: https://pdxscholar.library.pdx.edu/geog_masterpapers

Part of the [Geographic Information Sciences Commons](#), [Physical and Environmental Geography Commons](#), and the [Remote Sensing Commons](#)

Let us know how access to this document benefits you.

Recommended Citation

Loreno, Sara M., "The Cost of Hauling Timber: A Comparison of Raster- and Vector- Based Travel-Time Estimates in GIS" (2014). *Geography Masters Research Papers*. 7.
https://pdxscholar.library.pdx.edu/geog_masterpapers/7
10.15760/geogmaster.08

This Paper is brought to you for free and open access. It has been accepted for inclusion in Geography Masters Research Papers by an authorized administrator of PDXScholar. Please contact us if we can make this document more accessible: pdxscholar@pdx.edu.

**The Cost of Hauling Timber: A Comparison of Raster- and Vector- Based Travel-Time
Estimates in GIS**

by

Sara M. Loreno

A research paper submitted in partial fulfillment of the
requirements for the degree of

Master of Science
in
Geography

Committee:
Jiunn-Der Duh, Chair
David Banis
Martin Lafrenz

Portland State University
2014

Abstract

The cost of transporting forest products from harvest site to processing facility is a critical indicator of timber project feasibility. In order to accurately assess these timber haul costs, least-cost travel-time estimates among the site and the facility can be estimated using vector- or raster-based GIS methods. This study explores the applicability of both the raster- and vector-techniques within a variety of landscapes, including forested areas where comprehensive transportation data is limited. A comparison between the travel-time estimates derived from both methods was performed on three study sites using real-world data. The comparison also tested the effect of resolution on the reliability of raster-based calculations. The results indicate that travel-time estimates derived from a small resolution raster grid ($<25\text{m}$) are sufficiently similar to those derived from a customized vector layer, for both urban and forested areas. While the raster-based method did tend to slightly underestimate travel cost, harvest project managers are urged to utilize this method (in lieu of a customized vector layer) over Euclidean distance measures, for raster cost calculations were found to be a reliable and relatively accurate way to approximate travel-costs, without the extra time and requisite expertise of custom dataset creation.

Table of Contents

Abstract	1
List of Figures	4
List of Tables	5
Introduction.....	6
Research Question	6
Background	6
Spatial Methods for Measuring Transportation Costs	9
Method based on vector transportation data	9
Methods based on raster cost grids	12
Vector versus raster travel cost estimate methods	13
Case Study Methodology	15
Study Areas	15
Relative Accuracy Assessment Experiments	15
Uncorrected Dataset.....	16
Corrected Dataset.....	19
Travel-time Calculation Methodology	22
Origin and Destination Layers	23
Travel-time Calculations.....	24
Results.....	27
Travel-time Estimates	27
Discussion	32
Conclusion	38
References.....	39

List of Figures

Figure 1. Euclidean distance between two points as compared to network distance.....	7
Figure 2. Network travel model with edges (streets), junctions (intersections), turns and cost values.	9
Figure 3. Raster adjacency with a four node, eight node, and sixteen node connection.....	13
Figure 4. Map of study areas.....	16
Figure 5. Map of study areas with connecting highways. Map inset shows road detail of all merged roads from all data sources within study area 3.....	19
Figure 6. Example of an uncorrected road layer as compared to a corrected version.....	21
Figure 7. Representation of vector road layer converted into a 10 m, 25 m, 50 m, and 100 m raster cost grid.	26
Figure 8. Box plot and histogram comparisons of percentage difference in travel-time, relative to estimates from corrected data, for three study areas independently.	29
Figure 9. Histogram showing distribution of difference and percentage difference of calculated travel-times derived from the uncorrected vector, and 10m, 25m, 50m, and 100m raster datasets, relative to estimates from corrected data.	30
Figure 10. Histogram showing distribution of difference and percentage difference of calculated distances derived from the uncorrected vector, and 25m, 50m, and 100m raster datasets, relative to estimates from corrected data.	31
Figure 11. Roads on uncorrected network and the calculated shortest route between origin and destination, as compared to same route calculated on the corrected network.	32
Figure 12. Box plot and histogram comparisons of percentage difference in travel-time, relative to estimates from uncorrected data, for three study areas independently.	34

List of Tables

Table 1: Percentage of roads, in terms of length, on uncorrected network within various buffer distances of the corrected network.....	22
Table 2: Number of harvest sites and processing facilities utilized as origins and destinations for each study area.....	23
Table 3: Descriptive statistics for differences in travel-time calculations between corrected and other road datasets in minutes.....	28
Table 4: Results of Concordance correlation, Pearson’s correlation, and paired t-test, for the differences between corrected vector dataset estimates and estimates from the remaining five datasets.....	30

Introduction

Research Question

When calculating travel-time estimates between timber harvest sites and processing facilities in GIS, do raster-based methods provide sufficiently similar cost estimates as a customized vector network layer, and if so, what resolution of raster is most appropriate?

Background

The cost of hauling forest products, in various forms, from harvest site to processing facility constitutes a significant portion of the total cost of timber harvest projects. Depending on fuel costs and haul distance, the transportation of wood fiber can account for up to 50 percent of total harvest cost (McDonald et al. 2001). Length of haul distance alone is a significant predictor of harvesting behavior (Munsell 2011); beyond a certain distance profitability is limited and costs become directly proportional to haul lengths (Ashton et al. 2007). There are a variety of tools that forest managers use to evaluate the feasibility of proposed harvest activities, an example of such being the BioSum model (Fried and Christensen 2004) was developed. This software provides users a framework for evaluating silvicultural treatments by combining treatment costs, profits, and effectiveness, ultimately identifying the most appropriate treatment for any given stand. Haul costs, in the form of travel-time estimates, are a fundamental requirement of the BioSum model, as they constitute a significant portion of costs. In order to effectively utilize BioSum, it is necessary for users to accurately calculate travel-time estimates between harvest sites and processing facilities, often at a multi-state scale.

However, despite agreement that travel costs between locations are more complex than can be estimated by simple straight-line (Euclidean) distance measures (Figure 1), this representation is often used in forest management studies (e.g. Zhang et al. 2010; Coban and

Eker 2014) and as input in applications such as BioSum. This disparity between geographic information systems (GIS) capabilities and accepted practice in forest research likely arises because spatial data is often scarce and generally incomplete or outdated in regard to roads in forested areas. Without a comprehensive spatial dataset of forest and local roads in the area of interest, network based travel-time estimates depend on the creation of a custom road layer, typically compiled from multiple sources. This conflation of multiple vector layers involves first, the technical knowledge required to do so, and second, a significant time investment. While the resulting custom vector layer can be imperative to certain objectives, in the case of forest management, it is possible that a simpler and less time intensive method will produce sufficiently similar results. This study aims to explore the effectiveness and reliability of utilizing a raster cost grid, created from disparate vector road layers, for calculating travel-time between forest harvest sites and processing facilities, thereby eliminating the need for a customized road layer. The premise of using a raster approach to produce similar estimates of travel-time as a custom vector approach is that inconsistencies in the disparate vector road layers will be smoothed out, or erased, during the vector to raster conversion, resulting in more accurate estimates than could be achieved with uncorrected vector data.

Figure 1. Euclidean distance between two points (a) as compared to network distance (b).

While the relative benefits of raster and vector spatial data representations have been explored exhaustively in GIS and Geography literature, issues regarding travel-time estimation in remote areas have not been fully explored in forest management research. Considering the importance of accurate haul cost estimates to harvest project planning, an exploration of the data models and methods is warranted. In this study, travel-time estimates calculated using a custom network dataset will be compared to those derived from a raster cost grid (converted from multiple disparate vector datasets) in an effort to explore the effectiveness and feasibility of raster utilization in forest management applications. The main questions that will be addressed within this document include:

- What is the level of agreement between travel-times derived from raster and vector based models?
- What raster resolution, if any, is most appropriate for this application?
- Do landscape characteristics affect the reliability of results?
- What effect do underlying model assumptions have in travel-time outputs?

This paper contains a brief introduction to spatial data and models used to estimate travel costs, as well as a review of relevant literature, followed by a summary of a case study, its results, and a discussion of the implications for measuring timber haul costs.

Spatial Methods for Measuring Transportation Costs

Method based on vector transportation data

To estimate the minimum travel-cost between locations within a road network, a GIS network dataset can be created from a vector transportation layer. These networks model vehicle flow along street edges based on defined rules of connectivity. The three basic elements of a transportation network in a GIS are: (1) edge elements created from polylines in the vector source layer, (2) junction elements created by points in source layer, typically found at intersections, and (3) turn elements containing turn attributes that overlay two or more connected edges (Zeiler and Murphy 2010) (see Figure 2). In addition to these basic elements, various properties (i.e. elevation, traffic, and multi-model connectivity) can be assigned to network features. Once the model is constructed, cost measure attributes (such as travel-time) can be used by network solvers to find routes through a network that minimize the accumulation of that cost (Zeiler and Murphy 2010). Travel cost accumulations are calculated as a function of travel speed, edge distance, and turn delay. Network data models are commonly used in travel cost applications because real-world connectivity between locations is maintained: road segments are connected at road intersections. The accuracy of travel-cost estimates derived from network data models is highly dependent on the accuracy of the source vector representations and attributes.

Figure 2. Network travel model with edges (streets), junctions (intersections), turns and cost values.

Within a geographic information system (GIS), spatial data that has discrete boundaries, such as county borders and tax lots, are often stored in a vector data model. Vector data consist of points that can be joined to create lines or polygons; additional rules can be defined to determine how these vertices are joined. Transportation networks are most often represented using a vector data model, due to their line-based composition. Regional municipalities often create and maintain comprehensive spatial databases detailing the existing transportation network throughout their locality, including street attribute tables cataloging road type, speed limit, street name, and more. In addition to regional layers, transportation datasets are often created for larger areas, representative of an entire state or nation. In most cases, road features at this scale do not contain enough vertices to accurately delineate the sinuous nature of the feature, at which point data is said to become generalized, and positional accuracy is lost (Frizzelle et al. 2009). Positional accuracy, a measure of how accurately an object is positioned on the map with respect to its true position on the ground (Mowrer 1998), is typically much higher for large scale datasets (smaller areas, i.e. regions) than small scale representations (larger areas, i.e. state or national geographies). For this reason, large-scale spatial datasets are preferred for analytical purposes.

Spatial transportation datasets are utilized in a variety of ways by a variety of users. Common uses include address geocoding, route planning, facility service area calculation, and travel cost estimation. For many of these applications, it is possible that the area of interest will be larger than the extent of a regional scale transportation layer. Rather than use a more generalized dataset, various regional datasets can be combined, using the process of edge-matching, into a single layer covering the full study area. The merging of disparate datasets can be problematic because attribute data is often mismatched and rarely comparable (Frizzelle et al.

2009). However, this option is attractive because it utilizes high quality datasets in the areas where they are available, and requires limited GIS expertise.

In unpopulated and rural areas, however, regionally comprehensive transportation datasets are less likely to exist. In many cases ownership can be split between federal, state, and privately owned land. While it is possible to find agency-specific maps, these are often created at the state and/or multi-state level and are specific only to lands within agency ownership boundaries. In these cases and in contrast to urban areas, creating a positionally accurate and comprehensive dataset is more difficult than simply edge-matching regional datasets. Combining transportation layers with overlapping extents, particularly when dealing with a large area of interest, is a tedious and time-consuming process. Furthermore, the difficulty increases as the geometric distortion within each source layer increases. Conflation—the process of identifying control points on each layer and adjusting the data according to the correspondence between point pairs (Chen et al. 2004)—is often utilized in cases such as this to improve the features of one layer by combining the features of another. However, given that errors exist in geographic datasets and that conflation merges layers based on similarity thresholds, there is a degree of uncertainty as to the positional accuracy of conflated data (Adams et al. 2007). Additionally, conflation is a multi-step process demanding a high level of GIS competency. While automated approaches exist, none of the technical solutions have been able to fully incorporate the flexibility of human judgment (Schuurman et al. 2006), and human intervention remains a necessity. An alternative to conflation is the manual combination of various datasets by a GIS specialist, utilizing an orthorectified satellite or aerial images as a ground truth reference. While this manual effort is considerably more time consuming, there is less uncertainty as to the positional accuracy of the resulting vector layer.

Methods based on raster cost grids

Raster data models consist of a matrix of cells organized into rows and columns, with each cell in the grid containing a value, and are commonly used to represent thematic data (i.e. land cover or soils) or continuous data (i.e., temperature or elevation). A raster intended for use in cost-path analysis contains cell values representative of the time needed to traverse the cell (Delameter et al. 2012). Movement from one cell to an adjacent cell, in most GIS software packages, is defined by a node/link cell representation in which the center of each cell is a node and each node is connected to adjacent nodes by multiple links (ArcGIS Resource Center 2011). The number of links for which adjacent movement can occur depends on the software package. Typically, adjacency can be defined as a four node connection (cardinal directions), an eight node connection (cardinal and diagonal directions), or a sixteen node connection (cardinal and diagonal directions, plus “knight” movement) (see Figure 3). Each link, then, has an associated impedance derived from the cell cost and direction of movement (ArcGIS Resource Center 2011). Using this raster cost model, a cumulative distance grid can be created in which each cell is assigned the accumulative cost to the closest source cell. The main limitation of this procedure is that the defined connectivity between cells is only an approximation of the infinite number of movement options across the cost surface (Antikainen 2013).

Raster-based least-cost path algorithms are most often used to find routes across a continuous surface, not an edge based vector system, such as a transportation network. When linear features are converted into a raster representation their shape can be distorted, and feature width often becomes greatly inflated (Antikainen 2013). Furthermore, the jagged appearance of the vector network in the raster surface can cause overestimation of costs along a link (Choi 2014). Travel-time estimates derived from the raster cost model based on discrete vector

networks are highly dependent on not only on the accuracy of the source data, but also the resolution of the grid and number of links for which movement is allowed.

Figure 3. Raster adjacency with a four node (a), eight node (b), and sixteen node (c) connection.

Vector versus raster travel cost estimate methods

Studies that explicitly compare travel costs produced using vector data models with raster-based travel cost estimates are limited. While travel cost was not specifically considered, one of the most relevant studies was conducted by Sander et al. (2010). They examined the differences among distances calculated in three ways: Euclidean distances, vector-based road-network distances, and raster-based cost-weighted distances. Specifically, the distance between sample properties in the Twin Cities Metropolitan area and the closest open space area was determined using these three distance measures. Sander found that although distances were highly correlated between the three methods, the choice among the relatively common distance calculations could have potentially unintended consequences in subsequent analysis. Euclidean distances were found to significantly underestimate distance, while raster-based calculations were highly dependent on the resolution of the grid (Sander et al. 2010). This conclusion is in line with prior studies investigating the scaling properties of raster-based models: outputs are

strongly dependent on model resolution (e.g., Horritt and Bates 2001). In light of this caveat, raster resolution will be explicitly considered in this inquiry.

The field of health services research, while exploring levels of geographic access to care facilities, has produced several relevant investigations. For example, Delameter et al. (2012) examined the differences between raster and network data models in regard to population-based travel-time models. The study revealed that both network and raster-based methods of calculating travel-time among urban locations were heavily dependent upon an accurate representation of road location and travel speed. They also found that assigning raster cost values based on the *slowest* speed limit of all road segments within the cell boundaries resulted in a general overestimation of travel-time, and recommended the use of a different speed assignment rule in future studies. In the study undertaken here, raster cost values will be based on the *fastest* speed limit of roads within a cell to help reduce similar overestimation and test the effectiveness of this assignment rule

Case Study Methodology

Study Areas

Three study areas, varying in population density, street density, land cover, and land ownership, were identified within Oregon (Figure 4). The study areas range in size from 1,100 to 1,300 square kilometers. Area 1 (forested) is located in the upper northwest corner of the state and is comprised of mainly forested areas, specifically the Clatsop State forest. The majority of roads within Area 1 are forest roads, with very few highways or main arterial roads, and the study area contains no cities or towns. Area 2 (urban) is located in the southwest corner of Oregon and contains the cities of Medford—the fourth largest city in Oregon—and Grants Pass. Out of all selected study areas, Area 2 has the highest density of roads, least amount of forested area, and the greatest proportion of highways and major arterial roads. The final study area, Area 3 (mixed) contains the small town of La Pine, with a population of under 2,000 (United States Census Bureau 2013). The remainder of Area 3 is composed of a mixture of rural/agricultural land and forest (USDA Natural Resources Conservation Service 2014).

Relative Accuracy Assessment Experiments

In the case study, travel cost estimates were performed based on vector and raster methods. For the vector methods, two sets of vector transportation network datasets, uncorrected and corrected, were used. The datasets and methods performed in the experiments are described below.

Figure 4. Map of study areas

Uncorrected Dataset

The uncorrected transportation network dataset is a combination of multiple vector road layers from different sources. The four datasets selected for this study include the Bureau of Transportation (BLM) road data for Oregon/Washington, the Forest Service Topographic (FSTopo) Western USA transportation data, the Oregon Department of Forestry (ODF) road data, and the ESRI Streetmap North American road data. These layers were selected because when combined, they provided comprehensive coverage of roads located in rural and forested areas. For all sources the most recently released data (prior to July 2014) were utilized.

The BLM road dataset is maintained by the Oregon/Washington office of the Bureau of Land Management and is freely available to the public through internet download (Bureau of Land Management 2012). Visual comparison between the layer and remotely sensed imagery

revealed that this was by far the most comprehensive and most positionally accurate layer of the four datasets, however, only road segments falling on BLM land had an associated speed limit (miles per hour) attribute. In the case that a road segment did not have a speed limit attribute, a value of 25 mph was assigned, simply as a default value if the roads was not represented in any of the other three road layers. This process is explained later in this paper.

The Forest Service Topographic Western USA transportation data layer was created and is maintained by the USDA Forest Service. Like the BLM layer, the FSTopo dataset is available to the public free of charge, and can be downloaded directly from the internet (USDA Forest Service 2013). This layer was created from data derived from the U.S. Geological Survey, 7.5-minute topographic map series, with additional enhancements as necessary. While the quality of the FSTopo layer was similar to that of the BLM dataset, this dataset was less comprehensive than the former. The representation of roads in forested areas was much more complete than in urban or even agricultural areas. All roads in the FSTopo layer had an assigned speed limit attribute.

The Oregon Department of Transportation (ODOT) periodically releases an updated Oregon Transportation Network spatial data layer, easily acquired from the Oregon Spatial Data Library (2014) by the general public for free. The ODOT layer is compiled from numerous sources of data throughout the state. In this study, only data sourced from the Oregon Department of Forestry (ODF) were utilized. Data from other agencies were eliminated because roads sources from agencies other than ODF were present in the other datasets, along with a corresponding speed limit. Unfortunately, a speed limit attribute was not included for any of the roads in the ODOT dataset; because the roads were all sinuous forest roads, they were assigned a common speed limit of 20 mph. Due to the Oregon forest-specific nature of the data source

these data were especially useful in Area 1 (forested), where often ODF roads were not represented by any of the other three layers used in this study. Unfortunately, a speed limit attribute was not included for any of the roads in this dataset; because the roads were all spurious forest roads, they were assigned a common speed limit of 20 mph.

Lastly, the ESRI Streetmap North American road layer is available for download only with an ESRI ArcInfo license. This dataset was visibly less spatially accurate than the other layers, due to generalization typical of a dataset of a national scale. Many of the local and rural roads symbolized within this layer had unacceptable levels of positional error when compared to remotely sensed imagery. These gross generalizations were most evident on residential and forest roads, and for this reason road segments with an assigned Arterial Classification Code (ACC) of five—neighborhood and community access roads—were omitted. Unlike the other layers, the ESRI dataset was not clipped to the study area boundaries. In order to evaluate travel-times of longer trips, it was necessary to include origin points outside of the study areas. For this reason, ESRI roads (with an $ACC < 5$, that is, major roads include interstate and state highways and arterials) were included for the entire state of Oregon, and serve to connect all three study areas (see Figure 5).

Figure 5. Map of study areas with connecting highways. Map inset shows road detail of all merged roads from all data sources within study area 3. (Data source: ESRI 2014, ODOT 2014, USDA Forest Service 2013, Bureau of Land Management 2012)

After the layers were identified and acquired, they were merged into a single dataset that represented all roads existing within the selected study areas (whereas individually each data layer had major, and systematic, omissions). This layer will hereafter be referred to as the “uncorrected” dataset. Although comprehensive, because the uncorrected dataset was created from four different sources, roads are often represented by more than one line segment. In rare cases, as many as four lines symbolized a single road. Additionally, due to the varying levels of positional accuracy and generalization between data sources, some road segments reflect the geographic position of a street with a high level of accuracy, others do not.

Corrected Dataset

The uncorrected dataset compiled previously was manually examined and cleaned to create a more accurate representation of the existing road network, hence the “corrected” dataset

(see Figure 6). The cleaning process involved three tasks: (1) updating the speed limit attributes of the most positionally accurate road segments to reflect the *highest* assigned value of all segments that represent that road, (2) deleting duplicate and extraneous road segments, and (3) adjusting vertices to ensure connectivity between road segments. These tasks are further explained below.

The first step was to identify, out of all the line segments symbolizing a certain road, which was the most positionally accurate. To do this, the uncorrected vector dataset was manually compared to remotely-sensed imagery from the National Agriculture Imagery Program (NAIP). Specifically, a USFS GIS technician examined the uncorrected dataset and identified, for each road visible in the imagery, the vector line segment that most accurately represents it. Before deleting the inaccurate and duplicate road segments, however, it was necessary to assign the highest value speed limit from all representations of a road in the NAIP imagery to the identified “best fit” line segment. Only after the speed limit attribute was updated were duplicate and spatially inaccurate lines deleted from the corrected dataset. Finally, the GIS analyst examined the connectivity of the corrected vector lines. Deleting inaccurate line segments often resulted in dangling lines, or intersections in which the vertices didn’t meet precisely. These errors were corrected by moving individual vertices when necessary, to ensure that line segments connected with each other when appropriate. It took 80 hours to complete the cleaning process for the data within the study areas.

Figure 6. Example of an uncorrected road layer (left) as compared to a corrected version (right). Colors represent speed limit attributes.

To verify the road layer corrected by the USFS GIS analyst, the author independently completed the cleaning process on randomly selected areas that cover 20% of the three study areas. The resulting corrections were then compared to the results produced from the GIS analyst to ensure data quality. There were no discrepancies between the two datasets, indicating that data correction was completed in a predictable and reproducible manner. Table 1 details the percentage of road features (in terms of length) represented in the uncorrected network, within various buffer distances of the corrected layer, for each study area. These values were derived by first applying a buffer (of the indicted size) to all road features in the corrected layer. The uncorrected dataset was then overlaid on that buffered layer, and the percentage of total road length in the uncorrected layer that fell within buffer boundaries was calculated. This percentage value represents the positional agreement between two linear vector layers. A value of 100% indicates that no positional displacement larger than the buffer distance is observed between the two layers. Additionally, total length of roads in the corrected and uncorrected datasets, and the difference between the two (the length of road removed during cleaning), is also detailed for each study area. Out of the three study areas included in the analysis, Area 3 required the greatest

amount of editing (with 3,232 km of line removed from the original dataset), while Area 1 required the least (919 kilometers removed). The data indicate that the spatial (positional) displacement between the corrected and the uncorrected road layers is the most in Area 3 and the least in Area 1.

Table 1. Percentage of road length on uncorrected network within various buffer distances of the corrected network and total length (km) of roads in the corrected and uncorrected datasets, and the total length of lines removed during the correcting process.

	Buffer Size				Length of Roads in km		
	5 (m)	10 (m)	20 (m)	30 (m)	corrected	uncorrected	removed
Area 1	98.00%	98.40%	98.80%	99.00%	3,383,164.4	4,302,063.9	918.9
Area 2	88.40%	94.20%	98.40%	98.90%	3,990,185.4	5,259,686.0	1,269.5
Area 3	85.60%	93.00%	97.40%	98.60%	5,105,067.5	8,336,917.0	3,231.8

Travel-time Calculation Methodology

To assess the relative accuracy of calculated travel-time estimates from (often remote) timber harvest sites to the closest appropriate timber or bioenergy processing facility, both vector- and raster-based methods were used. Three processing facilities, as well as 36 (or more, depending on the study area) harvest sites, were identified as origins and destinations for each study area. An additional 4 processing facilities, located within a 100 mile buffer of the study areas, were also included for each of the three study areas (Table 2). Travel-time estimates were then calculated between processing facilities and harvest sites using network analysis methodology for both the corrected and uncorrected vector layers. Additionally, the uncorrected vector layer was converted into several cost rasters of varying degrees of resolution, including 10, 25, 50, and 100 meters, and a least cost-distance path analysis was used to calculate travel-times based on these raster layers. The end result is travel-time estimates (in minutes) from

processing facilities to the harvest sites within and surrounding that study area, for all methods. Estimates from the uncorrected and raster datasets were compared to those derived from the corrected dataset to determine if any of the alternative methods are adequate substitutes for the customized, corrected layer.

Table 2. Number of harvest sites and processing facilities utilized as origins and destinations for each study area.

	harvest sites	processing facilities (within area)	processing facilities (outside area)	processing facilities (total)
Area 1	41	3	4	7
Area 2	36	4	4	8
Area 3	40	3	4	7
Total	117	10	12	22

Origin and Destination Layers

Locations of actual timber and bioenergy processing facilities were used as the origin layer for this study. Using facility location data, a spatial layer of all currently operating processing facilities within Oregon was created. This layer was then clipped to include only the sites within study areas boundaries, plus an additional 4 sites for each area, all of which fell within a 100 mile buffer of area borders. This resulted in Areas 1 and 3 each having a total of 7 processing facility options, and Area 2 containing 8. It should be noted that travel-time is assumed to be the same when going from origin to destination, as when traveling from destination to origin. For this reason, processing facilities are used as origins, simply because there are fewer of them, simplifying data management slightly.

The harvest site destinations were represented by permanent inventory plots used in the USDA Forest Service Forest Inventory and Analysis (FIA) program. Since the 1920s the FIA program has collected extensive inventories of over 125,000 field plots systematically located across all forest lands in the United States (O’Connell et al. 2014). These plot locations were

determined based on a national hexagonal grid system, and are supplemented regionally by additional, “intensified” plot locations surrounding them. For this study, only plots with a status defined as being “forested” were considered. This filtered layer was then clipped to the study area boundaries. After removing plots that were outside the study areas, Areas 1 and 3 both contained 41 destination locations, and Area 2 contained 34.

Plot locations in forested areas are usually not directly accessible by road. In order for travel-time calculations to be performed, it was necessary for plots to be located on the road network. For this reason plot centroids were moved to the closest road segment on the final corrected network dataset. The corrected dataset was used for this operation because while all roads in this layer exist in the uncorrected version, the opposite was not true. Had the uncorrected data layer been used, plots could then reside on a road segment deleted from the corrected version, an undesirable result considering destination locations needed be identical among the different source layers used for travel-time calculations. Processing facility locations were also moved to the nearest road. While the majority these points were already on or near a road, it was important to ensure they fell on one of the line segments in the corrected dataset. The distance that each processing facility and plot was moved was recorded.

Travel-time Calculations

Network datasets were created from both the corrected and uncorrected road layers. During network creation, several cleaning operations and topology rules were applied to both datasets. All dangling road segments with an end node within 10 meters of another road were trimmed or extended to meet the intersecting road. The ArcGIS Integrate tool was then used to insert common coordinate vertices for roads falling within a 2 meter resolution of each other to ensure connectivity within the dataset. The integrate tool also created vertices wherever two line

segments crossed. Estimates of travel-time from the processing facility locations (origins) to the harvest sites (destinations) were calculated for both the corrected and uncorrected vector datasets using ArcGIS Network Analyst OD Cost Matrix tool. Line segments were assigned a travel-time value based on line segment length and speed limit attribute value. Travel-time (in minutes) was specified as the unit of accumulation.

In addition to travel-time estimates, travel distance (in meters) was also calculated for all origin-destination pairs, for all study areas. The same methodology used in travel-time estimation was used to calculate distance; however, length in meters was specified as the unit of accumulation, rather than minutes. This measure was to serve as a more direct comparison between the corrected and uncorrected networks, without the potential skew caused by the assigned speed limit parameter.

Raster datasets were created by converting the uncorrected vector layer into raster grids with resolutions of 10, 25, 50, and 100 meters (Figure 7). Cell values were assigned based on speed limit; miles per hour units were converted to seconds per kilometer to ensure compatibility with the cost algorithm and road grid cell units, whose horizontal unit was in meters. In the circumstance that two road segments with different speed limit assignments fell within a single cell, the *greater* value was used, with the assumption that drivers would choose the faster road when given an option.

Figure 7. Representation of vector road layer (a) converted into a 10 m (b), 25 m (c), 50 m (d), and 100 m (e) raster cost grid.

With the cost rasters used as impedance values, a least-cost distance raster was created for each processing facility in each study area. This step was completed using the ArcGIS Cost Distance tool, which utilizes an eight node connection model. The resulting output raster contained the least-cost distance (or minimum accumulative cost distance) of each cell to the input processing facility location. For example, seven least-cost distance rasters were created for each of the seven processing facilities in Area 1. Harvest site locations were then overlaid onto each of these output rasters and the resulting value was assigned as the travel-time estimate for that pairing. Raster-based travel-time estimates were generated for all four raster cost grids: 10, 20, 50, and 100 meters.

Results

Travel-time Estimates

Travel-time estimates derived from the corrected vector dataset were assumed to be the most accurate out of the six datasets (i.e., corrected, uncorrected, 10m, 25m, 50m, and 100m), due to the extensive manual data correction and removal of over 5,000 km of extraneous and inaccurate road representations from the original uncorrected dataset. Descriptive statistics comparing these estimates to those from the other five data layers for all study areas, independently and combined, can be found in Table 3. Overall, the 10m resolution raster grid produced estimates that most closely resembled those from the corrected layer. The average difference between estimates from the two layers was 4.69 minutes, with differences ranging from .02 to 31.32 minutes. Unsurprisingly, estimates calculated using the 100m resolution raster dataset were least similar, with an average difference in travel cost estimates of 11.46 minutes, and differences ranging from .01 to 43.67 minutes.

Given the range of travel-times between the processing facility and harvest site pairs (7 – 232 minutes), difference in estimated travel-time among datasets was also calculated as a percentage of the total travel-time. To arrive at this value, the difference between travel-time estimates derived from the corrected dataset and travel-times derived from other five datasets were divided by the corrected dataset estimate to produce the percentage difference. These values were calculated for each study area independently. Figure 8 compares the distribution of differences, and percentage differences between datasets for the three study areas.

The 10m resolution grid produced travel costs most similar to corrected costs, with 87% falling within 25% of total corrected travel-times in Area 1, 92% in Area 2, and 85% in Area 3. Travel cost estimates derived from the 100m resolution grid deviated most from the corrected

estimates, with only 64% of travel costs, across all study areas, within 25% of corrected travel estimates.

Table 3. Descriptive statistics for differences in travel-time calculations between corrected and other road datasets in minutes for all study areas combined, as well as each study area individually. All differences are positive, indicating travel-times from corrected dataset are higher than alternative calculations. Percentage differences are provided in parentheses. n = 855 total origin-destination pairs.

All Study Areas										
	mean		median		SD		minimum		maximum	
Uncorrected	8.60	(13.18)	6.96	(2.27)	7.59	(14.81)	0.00	(0.00)	34.87	(49.09)
10m	4.69	(09.06)	3.76	(7.71)	7.39	(12.24)	0.02	(0.02)	31.32	(45.94)
25m	6.97	(12.73)	5.46	(10.94)	7.30	(11.69)	0.01	(0.01)	32.09	(57.42)
50m	8.54	(15.46)	6.91	(14.99)	7.40	(12.09)	0.02	(0.01)	34.36	(58.35)
100m	11.46	(20.10)	10.67	(19.65)	8.06	(13.38)	0.01	(0.03)	43.67	(60.89)

Study Area 1										
	mean		median		SD		minimum		maximum	
Uncorrected	9.90	(14.50)	8.70	(15.60)	7.90	(17.66)	0.00	(0.00)	34.87	(49.09)
10m	7.00	(11.40)	5.70	(10.00)	7.03	(10.81)	0.03	(0.05)	31.32	(45.94)
25m	9.20	(14.80)	8.00	(13.80)	7.38	(11.13)	0.02	(0.05)	32.09	(57.42)
50m	10.70	(17.50)	9.50	(16.80)	7.55	(11.17)	0.08	(0.13)	33.09	(58.34)
100m	14.00	(23.00)	13.00	(23.00)	8.11	(11.90)	0.42	(0.34)	36.81	(60.88)

Study Area 2										
	mean		median		SD		minimum		maximum	
Uncorrected	6.43	(9.44)	3.98	(7.10)	8.61	(12.96)	0.03	(0.21)	30.72	(47.69)
10m	1.90	(3.49)	0.46	(0.76)	8.36	(12.73)	0.05	(0.05)	26.57	(42.03)
25m	5.16	(9.08)	2.53	(4.69)	8.11	(11.58)	0.02	(0.03)	31.17	(47.93)
50m	7.00	(12.16)	4.41	(8.32)	8.31	(12.15)	0.02	(0.01)	34.36	(48.68)
100m	9.88	(16.37)	7.40	(12.63)	9.27	(13.54)	0.00	(0.00)	43.67	(53.36)

Study Area 3										
	mean		median		SD		minimum		maximum	
Uncorrected	9.20	(16.90)	8.29	(17.09)	4.80	(10.77)	0.56	(0.26)	20.65	(42.12)
10m	4.59	(12.19)	4.67	(12.73)	5.30	(11.41)	0.03	(0.02)	17.50	(38.45)
25m	5.84	(13.95)	5.73	(15.41)	5.05	(11.65)	0.01	(0.01)	17.91	(40.03)
50m	7.18	(16.31)	6.83	(17.53)	4.91	(12.54)	0.18	(0.11)	18.43	(47.13)
100m	9.61	(20.29)	10.12	(22.23)	4.97	(14.16)	0.74	(0.33)	24.28	(53.51)

Figure 8. Box plot and histogram comparisons of percentage difference in travel-time, relative to estimates from corrected data, for three study areas independently.

Agreement between the travel costs estimated by different methods (i.e., uncorrected, 10m, 25m, 50m, and 100m) and by the corrected dataset was further explored by Concordance and Pearson correlations. A paired t-test was conducted to identify the significance of differences (Table 4). While travel cost estimates from all datasets were found to be highly correlated to estimates from the corrected dataset, differences were found to be significant. Travel-time estimates showing the strongest agreement with corrected calculations were produced using the 10m raster ($p_c > .987, p < .00001$), 25m raster ($p_c > .982, p < .00001$), and uncorrected vector dataset ($p_c > .977, p < .00001$). However, despite the strong relationship, travel-times calculated

using the corrected dataset were generally longer than those derived from the other layers (Figure 9).

Table 4. Results of Concordance correlation, Pearson's correlation, and paired t-test, for the differences between corrected vector dataset estimates and estimates from the remaining five datasets.

	Concordance Coefficient (ρ_c)	Pearson Correlation Coefficient (r)	paired t-test (t)
Uncorrected	0.977*	0.990*	29.67*
10m	0.987*	0.991*	16.60*
25m	0.982*	0.991*	24.99*
50m	0.978*	0.991*	30.21*
100m	0.966*	0.989*	37.21*

* p -value < .0000

Figure 9. Histogram showing distribution of difference and percentage difference of calculated travel-times derived from the uncorrected vector, and 10m, 25m, 50m, and 100m raster datasets, relative to estimates from corrected data.

In addition to travel-time, travel distances between harvest sites and processing facilities were calculated using each layer to separate the influence of travel speed assignment on predicted travel time. When a cell contains multiple road features the fastest travel speed was assigned to generate the travel-time for that cell. The calculation of travel distance, because it relies only on feature length, not assigned speed, provides a different context for evaluating the

alternatives. Unlike the travel cost comparisons, travel distance derived from the uncorrected vector layer showed the closest agreement to the corrected layer. Raster distance estimates differed substantially from corrected estimates, with no clear trend in deviation (see Figure 10). Due to data corruption and processing time restraints, distance estimates derived from the 10m raster are not included in these graphs.

Figure 10. Histogram showing distribution of difference and percentage difference of calculated distances derived from the uncorrected vector, and 25m, 50m, and 100m raster datasets, relative to estimates from corrected data.

Discussion

The results of the analysis show that travel cost estimates derived from all five datasets tested in this study were highly correlated to those from the corrected dataset albeit consistently shorter than the estimates derived from the corrected dataset. In the instance of the uncorrected dataset, these underestimations makes sense. Roads that do not exist, and were removed in the creation of the corrected dataset, effectively create shortcuts in travel route that often reduce travel-time estimates significantly. Figure 11 compares a travel route from processing facility to harvest site as calculated on the uncorrected network, to that same trip calculated on the corrected network. Upon removal of an erroneous feature, the route is forced along a more sinuous and indirect road, increasing travel-time.

Figure 11. Roads on uncorrected network (a) and the calculated shortest route between origin and destination (b), as compared to same route (d) calculated on the corrected network (c).

Considering that the raster datasets were created from the uncorrected vector dataset, it would be safe to assume travel-time estimates derived from these layers would show similar amounts of underestimation. However, the small resolution rasters showed closer agreement to the corrected dataset, than did the vector it was created from (a 9% average difference using the 10m raster, as compared to a 13% average difference using the uncorrected vector). By comparing the travel-time estimates from the raster to travel-time estimates from the uncorrected dataset, it is evident that the four raster datasets deviate from their reference dataset in different ways, depending on the resolution of the grid. The 100m, and to a lesser extent the 50m, underestimate travel-time as compared to the uncorrected dataset (Figure 12). This is presumably due to the coarse nature of the grid, the large area of the cells create connections that don't exist in the vector representation, reducing travel times. However, the smaller resolution rasters overestimate travel-time when compared to their reference layer. During conversion from vector to raster, the sinuous nature of the roads, easily represented in vector form, is exaggerated by the small blocky nature of the fine scale rasters, resulting in travel-time overestimation. This overestimation may inadvertently compensate for the underestimation of the uncorrected dataset, resulting in travel-times that more closely match the corrected dataset.

Figure 12. Box plot and histogram comparisons of percentage difference in travel-time, relative to estimates from uncorrected data, for three study areas independently.

When deciding on the appropriate methodology for travel cost estimation, there are several points to consider. First, it is important to determine the level of accuracy required to meet the study objective. In the example of forest management, while it is imperative that reasonable time estimates for travel between harvest site and processing facility are available to help managers project costs, often these times are rounded to a more coarse increment than an individual minute (e.g. 15 min increments or a distance-based cost estimate), . In this case, based on previous studies (Sander et al. 2010) and the results from this study, using a raster-based methodology provides travel-time estimates more accurate than Euclidean based methods,

without the requisite time and expertise needed to create a custom vector road network. For many other objectives such as emergency response time, however, a higher level of accuracy is necessary, for which a cost distance raster methodology might not be appropriate. As is evident by the travel-time comparisons presented here, raster-based calculations generally underestimate cost when compared to those derived from the corrected vector dataset. This bias can be minimized by utilizing a finer resolution raster; as cell size is reduced, topological correspondence between the original vector data and the converted raster data improves. However, a bias is present even at a 10m resolution.

Second, the quality road data available for the area of interest is a key consideration. As mentioned previously, regional governments and municipalities often maintain, and make publically available, vector road datasets of very high quality that contain minimal amounts of positional error. When determining whether a publically available road layer is appropriate for use, the accompanying metadata can provide some insight; often the scale and positional accuracy of the layer is provided therein. Additionally, when assessing a potential road dataset, it is important to consider the date the layer was created, and/or last updated. The more recent the data, the more likely newer roads will be represented. If a high quality dataset exists for the entirety of the study area (or several high quality layers can be edge-matched to provide full coverage), such a dataset is very similar—in terms of advantages of use—as a custom dataset. However, if no comprehensive road data layer is available, as is often the case when dealing with forest roads, and it is necessary to combine several layers of overlapping extent, use of a small resolution cost raster appears to offer a sufficient alternative to the creation of a custom layer.

A third consideration in deciding the appropriate methodology for travel-time calculation is the extent of, and landscape within, the study area. Network-based calculations require more

memory than their raster counterparts; it is often the case that calculation of network travel-times within a large extent (such as a multi-state analysis) will result in an error due to lack of memory resources on the user machine. A solution is to break the network into several sections and run each individually; however, the process is time intensive and is only effective if origin and destination locations are in close enough proximity to fit within the subsetting data. Raster cost distance calculations do not typically require as much memory, and therefore don't require segmentation in most cases. The landscape of the study area can also make a difference in the effectiveness of a raster in travel-time estimation. The grid-like road network of area 2 (urban) allowed for relatively accurate travel-time estimates produced by the raster-based method, (average deviation of 1.9 minutes, or 3.5%, difference). Comparatively, the sinuous roads of area 1 (forested) produced less accurate estimates, with an average deviation of 7.0 minutes (11.4%) difference.

It is important to note that while fine-scale cost rasters produced travel-time estimates that were reasonable when compared to a corrected dataset, raster distance estimates were found to be less reliable. Based on analysis results, vector-based methods appear to be the better choice for this application. In the example of forest management, travel-time estimation is often calculated by simply assigning a per mile cost to distance estimates (McDonald 2001). While this estimation method is less accurate than using a travel-time methodology, if utilized, a vector layer (even if non-customized) appears to be a better choice than a raster-based approach.

A final consideration in the process of calculating travel-times—and choosing a method for how to do so—is the accurate assignment of travel speed to road features. In this study, travel speed attributes were based on posted speed limit information found within the publically available vector datasets used in this study. It is assumed that these speed limit designations

generally do an adequate job of representing the reality of vehicle travel on any given road.

However, many factors that influence travel speed were not represented here, including slope, road material, traffic, and turn delays. In further research with the raster methodology, it is also necessary to consider what decision rule will be used when converting from vector to raster cost grid. Cell values are typically based on the travel speeds of road features falling within a cell extent; however, there are a few different options for which value will be assigned. In this study the fastest travel speed was used, resulting in a general underestimation of travel cost. Past research has shown that assignment based on the slowest travel speed resulted in a general overestimation of travel cost (Delameter et al. 2012). Another, potentially more appropriate, decision rule would be to assign raster cost based on mean travel speed of roads within the cell.

Conclusion

When calculating travel cost estimates among locations, there are many benefits to using a customized road network. Unlike many publically available pre-packaged datasets—which are often outdated or intended for use at a small-scale—a custom dataset can help limit the positional error and maintain the temporal accuracy of represented road features. Travel cost estimates derived from these custom network layers are typically reliable and their use in subsequent analysis easily justified. However, the creation of such a dataset is a time-intensive process that requires a high level of spatial analytic expertise, and is therefore not always an option when time and technical knowledge is limited. For this reason, it is important to identify alternative methodologies that produce similar time cost estimates as those derived from a custom network dataset.

This case study explored the suitability of utilizing various resolutions of raster cost grids to calculate travel-time estimates between forest harvest sites and timber processing facilities—an estimate vital to harvest project planning. Based on the comparison of travel-time estimates between an uncorrected vector layer, and 10m, 25m, 50m and 100m raster datasets, it appears the estimates calculated using a small resolution raster grid (10m) are in closest agreement to those derived from the corrected dataset. However, when calculating travel distance, rather than time, a vector dataset (even uncorrected) was found to be the most appropriate choice. Therefore, when determining which method to employ to calculate timber haul costs, it is recommended the one of the aforementioned methods be used instead of Euclidean distance, in order to produce more reliable and accurate estimates.

References

- Adams, B., L. Li, M. Raubal, and M.F. Goodchild. 2007. A General Framework for Conflation. University of California, Santa Barbara, CA.
<http://www.cs.ucsb.edu/~badams/GiscienceConflationAbstract.pdf> (last accessed 17 June 2014).
- Andersson, G., P. Flisberg, B. Liden, and M. Ronnqvist. 2008. RuttOpt – A Decision Support System for Routing of Logging Trucks. *Canadian Journal of Forest Research* 38: 1784-1796.
- Antikainen, H. 2013. Comparison of Different Strategies for Determining Raster-Based Least-Cost Paths with a Minimum Amount of Distortion. *Transactions in GIS* 17(1): 96-108.
- ArcGIS Resource Center. 2011. *Desktop 10: Cost Distance Tool*.
http://help.arcgis.com/en/arcgisdesktop/10.0/help/index.html#/How_the_cost_distance_to_ols_work/009z00000025000000/ (last accessed 17 June 2014).
- Ashton, S., R.Schroeder. 2007. Cost Factors in Harvesting Woody Biomass. In *Sustainable Forestry for Bioenergy and Bio-based Products: Trainers Curriculum Notebook*, eds. W. Hubbard, L. Biles, C. Mayfield, and S. Ashton, 153-156. Athens, GA: Southern Forest Research Partnership, Inc.
- Bureau of Land Management. 2012. *GIS Data for Oregon and Washington*.
<http://www.blm.gov/or/gis/data.php> (last accessed 25 May 2014).
- Chen, C.C., C.A. Knoblock, C. Shahabi, Y. Chiang, and S. Thakkar. 2004. Automatically and Accurately Conflating Orthoimagery and Street Maps. *GIS'04* Washington, DC. November 12-13.
- Choi, Y., J. Um, and M. Park. 2014. Finding Least-Cost Paths Across a Continuous Raster Surface with Discrete Vector Networks. *Cartography and Geographic Information Science* 41(1): 75-85.
- Coben, H.O., and M. Eker. 2014. A Hierarchical Approach to Estimate Spatially Available Potential of Primary Forest Residues for Bioenergy. *BioResources* 9: 4076-4093.
- Delameter, P.L., J.P. Messina, A.M. Shortridge, and S.C. Grady. 2012. Measuring Geographic Access to Health Care: Raster and Network-Based Methods. *International Journal of Health and Geographics* 11(1): 15.

- Fried, J.S., and G. Christensen. 2004. FIA BioSum: A Tool to Evaluate Financial Costs, Opportunities, and Effectiveness of Fuel Treatments. *Western Forester*.
http://jeremyfried.net/pubs/biosum_western_forester_2004.pdf.
- Frizzelle, B.G., K.R. Evenson, D.A. Rodriguez, and B.A. Laraia. 2009. The Importance of Accurate Road Data for Spatial Applications in Public Health: Customizing a Road Network. *International Journal of Health Geographics* 8: 24.
- Horritt, M.S., and P.D. Bates. 2001. Effects of Spatial Resolution on a Raster Based Model of Flood Flow. *Journal of Hydrology* 253: 239-249.
- McDonald, T.P., S.E. Taylor, R.B. Rummer, and J. Valenzuela. 2001. Information Needs for Increasing Log Transport Efficiency. *Precision Forestry Symposium*, University of Washington: Seattle, WA. June 17-20.
- Mowrer, H.T. 1998. Accuracy (Re)assurance: Selling uncertainty to the uncertain. In *Spatial Accuracy Assessment: Land Information Uncertainty in Natural Resources*, ed. K. Lowell and A. Jaton, 3-10. Chelsea, MI: Ann Arbor Press.
- Munsell, J.F., S.M. Barrett, and M.C. Bolding. 2011. An Exploratory Study of Biomass Harvesting among Logging Firms in Virginia and North Carolina. *Forest Science* 57(5): 427-434.
- O'Connell, B.M., E.B. LaPoint, J.A. Turner, T. Ridley, S.A. Pugh, A.M. Wilson, K.L. Waddell, and B.L. Conkling. *The Forest Inventory and Analysis Database: Database Description and User Guide Version 6.0 for Phase 2*. U.S Department of Agriculture Forest Inventory and Analysis, Washington D.C., 624 p.
- Oregon Spatial Data Library. 2014. *Transportation Data Framework Data*.
<http://spatialdata.oregonexplorer.info/geoportal/catalog/main/home.page#> (last accessed 25 May 2014).
- Sander, H.A., D. Ghosh, D.V. Riper, and S.M. Manson. 2010. How Do You Measure Distance in Spatial Models? An Example Using Open-Space Valuation. *Environment and Planning B: Planning and Design* 37(5): 874-894.
- Schuurman, N., D. Grund, M. Hayes, and S. Dragicevic. 2006. Spatial/Temporal Mismatch: A Conflation Protocol for Canada Census Spatial Files. *The Canadian Geographer* 50: 74-84.

- United States Census Bureau. 2013. *Population Estimates*.
<http://www.census.gov/popest/data/cities/totals/2012/SUB-EST2012.html> (last accessed 9 June 2014).
- USDA Natural Resources Conservation Service. 2014. *Geospatial Data Gateway: NAIP Imagery*. <http://datagateway.nrcs.usda.gov/> (last accessed 10 July 2014).
- USDA Forest Service. 2013. *Vector Data Gateway: Forest Service Topographic Map Data*.
<http://svinetfc4.fs.fed.us/vector/index.php> (last accessed 25 May 2014).
- Zeiler, M., and J. Murphy. 2010. *Modeling Our World: The ESRI Guide to Geodatabase Design*. Redlands, CA: Esri Press.
- Zhang, F., D. Johnson, and J. Sutherland. 2010. GIS-based Approach of Identification of the Optimal Pulpwood-to-Biofuel Facility Location in Michigan's Upper Peninsula. *POMS 21st Annual Conference* Vancouver, Canada. May 7-10, 2010.