

TECHNICAL PAPER

THE 2006 UPDATE OF THE CHRISTCHURCH COMMERCIAL VEHICLE MODEL

Authors:

Julie Ballantyne, BEng, MCILT
Senior Associate, Traffic Design Group
Contact: julie.ballantyne@tdg.co.nz

Andy Carr, BSc, MSc, MBA, MIPENZ, CPEng, IntPE(NZ)
Senior Associate, Traffic Design Group
Contact: andy.carr@tdg.co.nz

Presenter:

Andy Carr

ABSTRACT:

This paper discusses the proces of updating the Christchurch Commercial Vehicle Model (CCVM) to a 2006 base. In an approach considered innovative, GPS tracking data from sample Operators was used in conjunction with interview data from roadside surveys to develop the base year demand matrix. The observed demand matrix is the basis from which the three stage predictive model was subsequently estimated. This paper covers the data collected and the processing required for construction of a model.

OVERVIEW

The Christchurch Commercial Vehicle Model (CCVM) is a predictive three stage transportation model of commercial vehicle trips in greater Christchurch. The model, which was commissioned by Environment Canterbury, was built between 2006 and 2009 and has a 2006 base year. It replaces the previous model which was developed in 1990.

“Commercial vehicles” include medium and heavy vehicles, the definition of which is in Table A2.1 of Land Transport NZ’s Economic Evaluation Manual (EEM). Medium commercial vehicles are defined as two axle heavy trucks without a trailer over 3.5 tonnes gross laden weight. Heavy commercial vehicles include: rigid trucks with or without a trailer; articulated vehicles with three or four axles in total; trucks and trailers; and articulated vehicles with or without trailers with five or more axles in total.

As commercial vehicles constitute one component of traffic on the road network, the CCVM was devised as an integral component of the Christchurch Transportation Model (CTM), which forecasts person travel by private vehicle, public transport, and active modes. The integration of the CCVM and the CTM ensures all vehicles using the road network are incorporated in the planning process.

The main stages in building the Commercial Vehicle Model were:

- Data collection, including identification of suitable data sources;
- Development of matrices of observed travel patterns; and
- Estimation of the mathematical equations which constitute the model.

OBJECTIVES AND CONTEXT

The primary objective of constructing the Christchurch Transportation Model is to produce a robust analytical tool to reveal transportation issues and evaluate policies and proposals in the greater Christchurch area relating to the long-term strategic planning of the transport network. This will be a key contributor to an integrated, safe and sustainable approach to transportation planning for the future. To do this, the model must be capable of assessing the impact of strategic transport policies and facilitate monitoring and measuring of appropriate transport statistics as well as forming the basis for the assessment of major strategic infrastructure projects and land use developments.

The model covers (most of) Christchurch City, and parts of Waimakariri and Selwyn Districts. The Ashley River forms the northern boundary, with the townships of Kaiapoi, Rangiora and Woodend included. In the south, Rolleston, Burnham, Lincoln, and Tai Tapu lie within the model area, with Telegraph Road marking the south western corner. The model extends west as far as Aylesbury on West Coast Road and east to include the settlement at Diamond Harbour. The model area was defined to be consistent with that of the Greater Christchurch Urban Development Strategy (UDS).

The model has a base year of 2006, which was defined to be consistent with the National Census. The time periods for which traffic volumes on the road network are estimated are the AM peak (7am-9am), Interpeak (9am-4pm), PM peak (4pm-6pm) and Overnight (6pm-7am). Daily flow estimates are produced by summing the four period assignments.

OBSERVED DATA

Prior to embarking on building the model, the first step was to build a matrix of observed commercial vehicle trips. This observed dataset is fundamental in that it is the basis from which the mathematical equations that constitute the model are derived. The data required includes travel pattern data, which is generally collected on a sample basis but must include the origin and destination of each trip by definition, and spot count data for independent verification of the model's ability to replicate observed trips.

The preliminary phase of the data collection exercise involved investigating possible data sources, considering cost implications and any potential statistical biases that could skew the dataset. Examples of biased datasets include collecting more information from companies undertaking long distance haulage compared with those involved in the shorter, metropolitan trips, without any adjustment for the differential sampling. This would produce a model over-representing longer trips and under-representing shorter movements. The following describes the data collected and the post-processing.

Travel Pattern Data

The two main sources of travel pattern data for commercial vehicles within, and through, the model area were:

- Roadside interview surveys (RSI); and
- GPS polling data.

Roadside Interview Surveys

Twenty-two sites were surveyed in early 2007 using the 'traditional' roadside interview survey technique for all vehicle types. Roadside interview surveys involve actively intercepting drivers during their journey, with interviewers asking drivers for information on their vehicle trip, specifically, where they had come from, where they were going, and the purpose of their activities at those locations. The interviewer also records additional information relating to the time of the survey, the vehicle type and occupancy.

In practice, this is achieved by establishing a coned-off survey bay plus a bypass lane for the survey direction. On the driver's side of the survey bay, another area is coned-off for the interviewers to stand. Vehicles travelling in the opposite direction are not intercepted. The survey bays are established to accommodate a minimum of four cars at any time, with some capable of holding six vehicles. While the length of the bay varies depending on site-specific conditions and expected traffic volumes, bays were of sufficient length and taper to enable heavy vehicles to be surveyed. Appropriately skilled traffic management personnel flag vehicles into the survey bay. After the bay is full, traffic is then directed to use the bypass lane. Sampled vehicles are asked to participate in the brief survey, and following the interview, they exit the bay and merge with traffic using the bypass lane. Once the survey bay has emptied, new vehicles are directed into the bay and the cycle repeated. In light traffic conditions, vehicles were typically delayed for approximately two minutes, based on a one minute duration for the actual survey.

The overall process was tested using pilot surveys, which were conducted in February 2007 at two sites. The main fieldwork followed in March and April 2007, with continuous surveying from 7am to 7pm on a weekday at each site, excluding school holidays and public holidays.

Each of the twenty-two sites was surveyed for one weekday. Manual classified counts were conducted for both directions of travel during the interview surveys, recording vehicles by type in 15 minute increments. There were a total of just under 30,000 interviews, of which just under 6,000 represented travel by either medium or heavy commercial vehicles.

The sites for the roadside interview surveys were selected as all roads crossing three internal screenlines plus the external screenline. This ensured that all trips crossing any of the roadside interview screenlines were intercepted and sampled. As an example, the Waimakariri River was a screenline and trips crossing the river were identified by interviewing immediately south of the motorway on Main North Road and at a similar north-south location on Old Main North Road/Marshland Road. The sectors formed by the roadside interview survey sites are shown in the following table and illustrated graphically in Figure 1. These sectors are key in the process of combining all the data collected, which will be described later.

Table 1: Sectors Defined by RSI Sites

SECTOR	DESCRIPTION
1	Waimakariri (defined as within Model Area)
2	South of Waimakariri, North of QEII Drive
3	Sumner, Lyttelton and Governors Bay
4	Rest (Greater ChCh south of QEII, Selwyn in Model Area)
5	External to Model Area, Waimakariri
6	External to Model Area, Selwyn

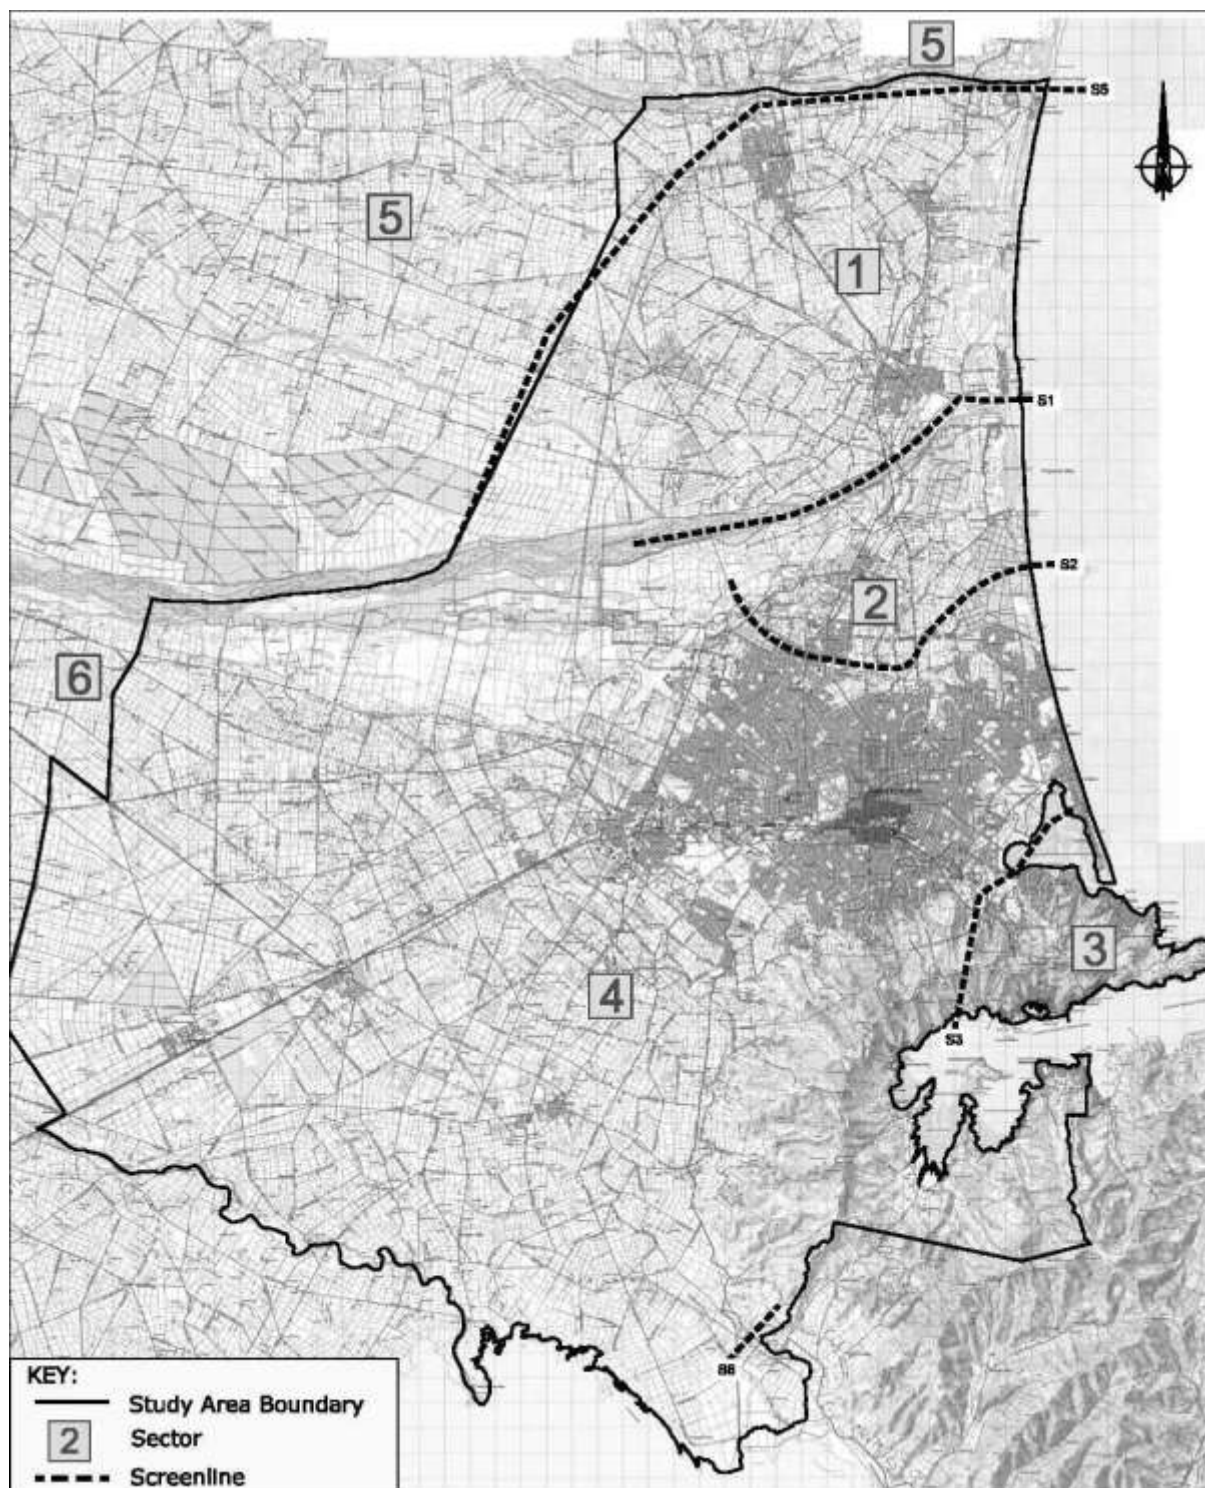


Figure 1: Sectors Defined by RSI Sites

Expansion from the sample to the total traffic volume was undertaken using manual classified counts undertaken concurrently with and at the same location as the roadside interviews. The expansion process considered vehicle type and time periods. While initially medium and heavy vehicles were expanded separately, this resulted in too many instances where there was a count but no interview. Medium and heavy vehicles therefore had to be combined to undertake the expansion.

GPS Data

While the roadside interview surveys provided a robust sample of travel pattern data, the data on commercial vehicles was limited (by nature of the traffic patterns), and a more targeted approach was required. Utilising GPS data to collect travel patterns of commercial vehicle trips was therefore pursued as a cost effective mechanism to collect the essential detailed travel patterns of commercial vehicles sampled.

Initially a pilot study was commissioned by Environment Canterbury where a mobile GPS unit was placed in a truck cab (the unit being of the same type as might be used by a tramper). This showed that GPS data had great potential for unlocking large quantities of useful data but that using mobile units was not feasible due to the cost of post-processing the data collected, and the number of such units that would be required to collect a suitable sample.

Shortly after this, a meeting occurred with a representative of a commercial GPS provider. During the meeting it transpired that all operators using their system received a summary of their fleet operations, being a condensed version of the data that was transmitted from each cab unit to a central location. It was surmised from this that there was the potential to intercept the raw data and process it in a manner which could be used inform journey origins and destinations.

A phone survey with follow-up questionnaire was sent to transport operators to ascertain the level of use of GPS systems and to identify companies willing to participate in the study. Meetings were subsequently held with representatives of companies using GPS systems who indicated some willingness to participate in further negotiations over use of their data. One of the issues was that the GPS data included the vehicle speed and there were concerns in providing Environment Canterbury with this information. The solution was to involve a trusted third party to access, process and clean the GPS data before it was made available for use. Five companies eventually agreed to provide GPS data for their fleet, although as confidentiality agreements were signed, details of the participating operators cannot be revealed. This process took several months to complete.

The participating companies provided access to their fleet's GPS polling data which was a continuous record of vehicle movements at either two or five minute intervals (system dependent). The data also identified each individual vehicle, and the time and date of travel. The amount of data provided varied by company, but the maximum was six month's worth of travel data and the minimum was three weeks.

The GPS data at this disaggregate level required significant processing to convert the individual pollings to trip records. This was undertaken by a third party, who formed trip records by aggregating the two or five minute pollings. The challenge in this exercise was determining which poll locations identified either the start or end of a trip. This was accomplished by applying the following general rules:

- If the fleet turned-off the ignition during deliveries, this was used to define the trip start and trip end; and
- For those that did not turn off their ignitions for deliveries, a trip start was defined as the first record of movement over 20 kph after a period of idling. A trip end was then defined as the converse.

Spurious GPS pollings and data outside the scope of the study were also removed. These included:

- GPS polling that indicated no travel between ignition on/off i.e. idling;
- Trips undertaken wholly within a transport yard (same origin and destination or very low speeds);
- Trips completely outside of the model area (e.g. wholly within the North Island); and
- Weekend travel.

The process of identifying a “trip” was complicated and required several iterations to refine, as follows.

The first dataset provided was plotted graphically using GIS. Interrogation and analysis of this first dataset revealed a potentially fatal flaw. Vehicle trips were starting at locations that were not the same as where the previous trip had ended. Basically the data implied that a vehicle ended a trip then appeared at a different location to start the next trip. This proved relatively easy to resolve as it was found to be a simple record processing error. The processing was revised and reapplied such that the start location of every trip was the same as the end of the previous trip, rather than reading the next record to be the trip start.

The second dataset correctly showed travel by vehicle to consecutive destinations. Another serious problem was discovered, however, with trip starts and ends found in extremely unlikely locations. The most significant of which was the number of trip ends on the approaches to busy intersections. While the possibility of on-street deliveries adjacent to intersections was considered to account for a proportion of these, the number of observations suggested that this was unlikely to be the main reason. Further investigation revealed that the GPS systems did not register movement at vehicle speeds less than 20 kph, so that vehicles queuing at traffic signals were being incorrectly identified in the data processing as a trip end. The solution was to apply GIS buffer zones around all major intersections and relink the GPS pollings so that any trip starts or ends within the buffer zones were ignored. It is accepted that in applying this amended rule, any on-street deliveries in the immediate vicinity of intersections will have been incorrectly removed from the dataset. But the alternative of retaining such a high number of trip ends within the road reserve was considered to give rise to a far more significant problem.

Other spurious start or end locations were also found in the dataset and were addressed at the same time as the GIS buffering for trip ends within intersections. These spurious locations were generally a result of operators not turning off the ignition of the vehicle until the end of the day, which meant identifying the start and end of each trip was more difficult. The types of spurious trip ends identified included:

- locations where trip ends were not allowed, such as on motorways; and
- locations where trip ends were impossible, such as in the ocean or too far from the road network.

These questionable trip ends were removed from the dataset and complete trips reformed by linking the previous and subsequent trip ends.

The final part of the process was aggregating the individual trip start and end locations to model zones. This ensured confidentiality was maintained as the actual address of each trip was not retained in the processed dataset.

Traffic Count Data

A total of 418 automatic traffic counts were collected for roads crossing pre-specified screenlines, being a combination of existing data from the various road controlling authorities and newly-collected data. There were 23 screenlines defined in total across the model area, which were established to check the validity of the model's ability to replicate observed traffic counts.

All of the automatic traffic counts were directional, classified (i.e. by vehicle type), by hour of the day, and represented consecutive days. An average weekday flow (Monday to Friday) was estimated for each count site by hour of the day, vehicle type and direction. This averaging considered tube failures that resulted in data loss, so that some averages were calculated over less than five days. Some counts were collected prior to 2006, which is the base year of the model. These counts were factored to represent 2006. All counts were seasonally adjusted to represent the neutral month of October.

DEVELOPMENT OF OBSERVED TRAVEL PATTERNS

Data Processing

The roadside interview survey data was combined to produce a matrix of observed trips by medium or heavy commercial vehicles. The RSI data processing included:

- For each of the 22 sites, factoring the expanded interview dataset to the nearest automatic traffic counts (initial expansion was to manual classified counts);
- Combining the sites to form a single matrix of observed travel. The key to this was identifying and removing vehicle trips that would pass through more than one of the survey sites. This ensured longer trips were not over-represented in the survey database;
- Estimating the reverse (i.e. non-survey) direction of travel. This not only entailed reversing origins and destinations, but also time periods so that vehicles observed southbound in a morning peak were inverted to represent northbound trips in the evening peak. As well as reversing the trip record, the trips were then factored to the automatic traffic counts for the non-survey direction by vehicle type and peak period; and
- Adjusting the combined matrix so that the trips crossing each screenline were correct. Although each site originally had the correct number of trips (by definition), the adjustment to remove double (and triple) counting of certain movements introduced discrepancies between the processed roadside interview survey database and the traffic counts.

The final step was accomplished using a manual factoring process rather than applying automatic matrix estimation techniques. The manual factoring enabled a hierarchy of adjustments to be imposed, so that, for example, trips across the Waimakariri River were corrected first while trips across the QEII Drive screenline further to the south were corrected

in the next step, with any error at this screenline allocated to trips from south of the Waimakariri River (as cross-Waimakariri trips had already been adjusted in the previous step). This ensured the Waimakariri River screenline, which was captured by surveying at two sites on the same day, took precedence over the QEII Drive screenline, which was surveyed on four different days and had eight sites in total forming the screenline.

The resulting expanded matrix of daily travel included just under 13,000 commercial vehicle trips, which are shown in the following table in the sector system described in Table 1.

Table 2: Daily RSI Commercial Vehicle Trips by Sector

		Destination Sector						Total
		1	2	3	4	5	6	
Origin Sector	1	0	137	44	876	332	54	1,443
	2	187	0	55	1,115	80	51	1,487
	3	52	33	0	1,037	42	71	1,236
	4	752	1,007	1,012	692	553	1,393	5,409
	5	401	125	13	489	15	58	1,100
	6	126	72	26	1,632	112	3	1,970
	Total	1,518	1,375	1,149	5,841	1,134	1,628	12,645

The GPS dataset represented a matrix of travel on weekdays for five companies and up to six months worth of observations.

This included a total of 97,000 trips by medium or heavy commercial vehicles. No adjustment for multiple observations was required as the sampling basis was by company and hence there was no possibility for any trip to be double counted. Adjustment from the six months of data to a single average weekday was required, however, and this was accomplished by dividing by the number of days' worth of travel data for each vehicle. Following the adjustment to an average weekday, the GPS dataset included a total of 1,137 daily trips by medium or heavy commercial vehicles. This still represented the sample only, however, so adjustment was required to factor the sample to the total population.

The fleet size in greater Christchurch was estimated to be in the order of 2,500 vehicles compared with 130 vehicles sampled through the GPS survey. This represents an average expansion factor of 19 to convert the sample to the total.

This simplistic factoring, however, does not consider how representative each of the five companies is of the total freight industry, nor does it account for any geographic or trip length differences unique to any of the companies. There was no solution to this issue, without conducting comprehensive industry-wide surveys (which were outside of the project budget and timescale). However this area represents a logical future refinement of the model.

An alternate approach for the expansion of the GPS dataset was to compare the GPS sample data to the roadside interview survey data for similar movements across screenlines. The daily trips from the two data sources are shown in the following table at four screenlines, with the implied expansion factor.

Table 3: Comparison of RSI and Sample GPS Data at RSI Screenlines

RSI SCREENLINE	RSI (EXPANDED)	GPS (SAMPLE)	EXPANSION FACTOR
Waimakariri River	1,795	54	33
North of QEII Drive	2,753	80	34
Sumner/Lyttelton	1,149	138	8
Externals (North & South)	2,762	106	26
Total	8,459	379	22
Total excluding Waimakariri (double counted)	6,664	324	21

The difference between the estimated expansion factors for travel to the Port (Sumner/Lyttelton screenline) and across the Waimakariri River does suggest that the trips in the GPS dataset were not equally sampled throughout the full geographic coverage of the model and that the sampled companies made more trips to the Port area than to Waimakariri relative to total commercial vehicle movements. Nevertheless, in the absence of other data, the GPS data must be used and an appropriate expansion factor calculated and applied. An overall expansion factor from this approach was calculated excluding the Waimakariri screenline, which incorporates some duplication of traffic movements across the QEII Drive screenline. This indicated an average factor of 21 to convert the GPS sample dataset to the total commercial vehicle movements in Christchurch.

An average factor from the two approaches, the fleet size based calculation which suggested a factor of 19 and the comparison with the RSI data which indicated a factor of 21, was adopted. A factor of 20 was therefore applied to the trips in the sampled GPS dataset to produce 23,000 daily trips representing the total fleet within the model area. The expanded GPS trips by sector are shown in the following table.

Table 4: Sampled Daily GPS Commercial Vehicle Trips by Sector

		Destination Sector						Total
		1	2	3	4	5	6	
Origin Sector	1	20	36	6	80	24	2	168
	2	30	200	86	628	128	16	1,088
	3	4	18	274	2,714	2	22	3,034
	4	70	622	2,632	11,161	782	960	16,226
	5	32	182	2	768	24	10	1,018
	6	12	76	42	926	10	142	1,208
	Total	168	1,134	3,042	16,276	970	1,152	22,742

The next stage was to combine the roadside interview survey and GPS datasets and develop a single matrix of travel by commercial vehicles. The methodology of combining the two datasets had to consider the differential sampling rates, which meant that the GPS and RSI data could not simply be added. As the basis of the GPS sampling was difficult to quantify

but the sampling of the RSI's was known, the RSI data was used in preference to the GPS data for movements across screenlines where roadside interview surveys were conducted.

The screenlines were used to develop a sector system so that any inter-sector movements were observed through the roadside interview surveys by definition while intra-sector movements were not. The GPS and RSI datasets were then spliced with sector-to-sector movements using the RSI data and intra-sector movements using the GPS data. This produced an initial estimate of base year (2006) travel by commercial vehicles of 24,000 daily trips, as shown in the following table.

Table 5: Combined RSI and GPS Data – Daily Commercial Vehicle Trips - Initial

		Destination Sector						Total
		1	2	3	4	5	6	
Origin Sector	1	20	137	44	876	332	54	1,462
	2	187	200	55	1,115	80	51	1,687
	3	52	33	274	1,037	42	71	1,510
	4	752	1,007	1,012	11,161	553	1,393	15,878
	5	401	125	13	489	15	58	1,100
	6	126	72	26	1,632	112	3	1,970
	Total	1,538	1,575	1,423	16,309	1,134	1,628	23,607

The shaded cells indicate the intra-sector movements where the GPS data was used; the remainder adopted the roadside interview survey data. Half of the total of this matrix is associated with trips within sector 4 which is the majority of Christchurch (south of QEII Drive, north of Sumner and Lyttelton Port). This reinforces the requirement to source and incorporate the GPS data as trips within Christchurch were only partially observed through the roadside interview surveys.

Matrix Refinement

The initial estimate of daily travel (24,000 trips) was checked by assigning this demand to the road network and comparing the result with observed traffic counts across screenlines. This comparison suggested that the actual number of commercial vehicle trips was significantly higher than the initial estimate of travel developed from the RSI and GPS datasets. Before a model could be built, it was essential to adjust the matrix of observed trips, which was achieved by applying matrix estimation techniques.

While matrix estimation is a powerful tool, it does not result in a single, unique solution – multiple solutions are possible. The more inputs to the process, the more robust the final solution is likely to be. The inputs to the daily trip matrix estimation process for the CCVM were:

- Initial trip matrix from GPS/RSI surveys;
- Traffic counts;
- Estimated trip generations by zone;
- Routeing through the network; and
- An estimate of confidence in each dataset.

Although the traffic counts, as described in section 3.2, were processed to distinguish medium from heavy vehicles, total commercial vehicles were used from this point onwards as the GPS data did not distinguish vehicle type.

For each model zone, trip ends were estimated and input to the matrix estimation. This was critical in ensuring large numbers of trips, implying excessive trip rates, were not allocated to any one zone in deriving the refined demand matrix. There were two components to the trip end estimate, commercial vehicle trips to residential areas and commercial vehicle trips to commercial/industrial areas.

Both elements were calculated by applying approximate trip rates to the relevant variable. For residential areas, a daily trip rate per household was applied. The household numbers were extracted from the 2006 Census while the residential trip rates were derived from a survey conducted specifically for this purpose. For industrial/commercial areas, daily trip rates were estimated for industrial, office and retail areas based on both site area (hectares) and floor area (m²). These trip rates were derived from data collected during previous surveys that were commissioned by Environment Canterbury.

Christchurch City Council and Waimakariri District Council provided comprehensive databases with site area, floor area and type of land use on an individual site basis. From this, every site was allocated into the appropriate category (industrial, office, retail) and the daily trip rate applied to the site area and floor area. The minimum value of the site or floor area trips was adopted, with a few exceptions that are discussed later. The daily trips for each site were then summed to zone level. For Selwyn District, a more simplistic approach was adopted due to the lack of detailed floor or site area data, with commercial area assumed for Town Centres and the above rates applied. Independently surveyed intersection turning counts were used to assess the traffic generation of the Rolleston industrial area.

For some areas, the estimated trip ends were less than the observations in the initial RSI/GPS matrix. This is a result of either the sampling to produce the initial matrix or the macro-scale calculation of trip rates resulting in underestimates for some individual sites. Any notable differences were reviewed and adjusted on an individual zone basis using professional judgement. In some cases, the maximum of the site or floor area trip rate was used instead of the minimum, while for other zones, an average trip rate was adopted.

Trip ends at external zones were set equal to the traffic counts.

Each dataset input to the matrix estimation process was allocated a confidence value. These are an important control in the matrix estimation process as they identify the relative accuracy of each dataset. The higher the confidence level, the more likely the data reflects observed and hence the changes imposed by the matrix estimation procedure are smaller. Traffic counts were allocated the highest confidence, while the initial origin-destination data sampled through the RSI/GPS surveys was allocated a much smaller confidence.

During the matrix estimation process, a number of issues with the observed data were uncovered and resolved. The main concern was that the GPS dataset was very sparse, with large numbers of observations for some zones in terms of origin-destination movements, but very few observations for other zones. This, of course, was a direct consequence of only sampling five companies. The movements regularly made by these companies were well represented in the initial matrix, but other areas were under-sampled. The issue was

addressed by seeding the matrix with data from other adjacent zones with similar land use characteristics, which smoothed out the estimated matrix and ensured a single observation was not increased to represent all daily travel for these problematic zones.

Following the application of the matrix estimation procedure, the initial estimate of 24,000 daily trips was increased to 38,000 daily commercial vehicle trips, as shown below by sector.

Table 6: Combined RSI and GPS Data – Daily Commercial Vehicle Trips - Refined

		Destination Sector						Total
		1	2	3	4	5	6	
Origin Sector	1	226	141	20	960	429	68	1,843
	2	189	330	21	1,016	123	63	1,741
	3	19	24	241	1,120	10	55	1,470
	4	855	1,004	1,183	24,589	444	1,428	29,502
	5	481	140	4	407	13	54	1,099
	6	112	62	40	1,656	109	2	1,981
	Total	1,881	1,701	1,510	29,747	1,129	1,669	37,637

The most significant difference comparing the refined daily matrix of commercial vehicle travel with the initial matrix developed from survey data are the trips within sector 4, which represents the majority of Christchurch. The initial matrix contained just over 11,000 trips whereas the adjusted matrix suggests that there are almost 25,000 daily commercial vehicle trips within Christchurch. This difference is a direct result of the estimate of expanding the sample of the GPS data to the total population. It also reinforces the need to collect a larger sample for future updates of the model and to further investigate the basis of the sampling.

This level of demand could be shown to reproduce observed traffic counts at screenlines throughout the City. It was finally concluded to be a suitable level of demand and trip making from which the synthetic three stage model could be developed.

CONCLUSION

In conclusion, the validity of using GPS data as a key source of travel information for commercial vehicle movements was confirmed through this exercise. The need for more information on the composition of the market in terms of short versus long trips was identified, which would in turn enable a more tailored expansion process to be applied in increasing the sample data to the total population. A larger sample of GPS movements would also be required for any future model updates, particularly in the absence of roadside interview data.