Sam Schwartz Engineering, D.P.C.
322 Eighth Avenue, 5th Floor
New York, NY 10001
phone: (212) 598-9010
samschwartz.com

Sam Schwartz
Transportation Consultants

Memorandum

To: NYC Department of Sanitation
From: *Sam Schwartz*
Date: June 2, 2016
Re: Private Carting VMT Analysis

Introduction

The overarching goal of the truck routing analysis was to quantify the extent of the annual vehicle miles traveled (VMT) by the private carting industry today and to compare that with a hypothetical exclusive zone collection system. The analysis focused on the 90 carters that pick up recyclable and putrescible waste. Data on the existing customers, carters, and the trucks routes were provided by Business Integrity Commission (BIC) and analyzed using ArcGIS software, Excel, and Python.

The general approach to the study was to map each carter's reported truck routes in order to estimate annual VMT of the private carting industry. The BIC customer registry was then used to create 11 hypothetical zones covering the five boroughs, where zones consisted of one or more Community Districts and were roughly equally sized according to the number of businesses within each. New, hypothetical truck routes were then mapped in each zone to estimate the annual VMT of a zoned system. These results were then compared to the existing conditions to estimate the change in VMT resulting from a zoned system.

Data Sets

There were two main data sets that were provided by BIC for this analysis: the customer registry and the private carter collection route data set.

BIC Customer Registry, December 2014

The customer registry is a complete list of every customer (account) that is served by a private carter in the city. It has over 130,000 records, noting each business name, type, address, and which carter they are served by, in addition to characteristics not germane to this part of the study. Some records also listed either the volume or weight of the waste picked up. This data set is self-reported by the carters to BIC.

This data set was filtered and cleaned to show the accounts that are served by the carters that primarily haul recyclable and putrescible waste. This consolidated version of the customer registry had nearly 108,000 records. The BIC customer registry was geocoded to map each account, which was used to understand the spatial distribution of private carter customers to create the foundation of the zoned system. The customer registry was also used to rank the 90 carters by number of accounts served. This information was used to select a subset of the overall data set to map in the routing analysis and to create a multiplier to adjust the subset to represent the full 90 carters.

Private Carter Collection Route Data Set

The private carter data was collected by BIC from each carter, and provides information on transfer station and garage locations, lists each account served by each truck route, and the order in which the accounts were served. BIC provided each carter with a set of instructions and a form in Excel to provide this information. Carters were requested to submit four weeks of collection data as a representative sample for the year:

1. July 6-12, 2014
2. September 7-13, 2014
3. December 21-27, 2014
4. January 25-31, 2015

Pursuant to the administrative code, companies licensed by BIC must maintain route sheets, which are subject to review by BIC. The administrative code does not specify how a company must maintain these files and this was the first time BIC has requested the route information from the carters.

While the routing data generally had the same format, each carter filled out the forms differently (and some were not consistent even among their own routes), and some provided one sheet for all routes, while others provided a single sheet for each route. This equated to over 4,600 individual Excel files for the 90 carters, with inconsistent formatting and other various issues related to self-reported data.

The private carter data set was used to map and quantify the VMT of the existing private carting system. For the subset of data that was analyzed, each route and its list of accounts were mapped. The geocoded accounts that were included in the Existing Analysis were then used in the Zoned Analysis to estimate the change in VMT.

Results

It was found that the private carting industry currently contributes approximately 23.1 million VMT annually to the streets of New York City and the adjacent counties, where many garages and transfer stations are located. Many existing routes are geographically dispersed, often serving several neighborhoods across multiple boroughs. Routes from the same and different carters often overlap along key routes and neighborhood streets, creating duplicative services across the city. For many routes, garages and transfer stations are far from the core service area of the route.

It was also found that a Zoned System could reduce the VMT of the private carting industry in the range of 50%-70%. This is primarily due to the fact that Zoned System's routes could be expected to be much more efficient than the existing system due to the geographic constraints created by exclusive zones.

Under a Zoned System, every neighborhood in the city would likely see a decrease in truck traffic, VMT, and the related emissions, however the areas with the highest concentrations of VMT will see the greatest benefits. These areas are in many cases the neighborhoods of greatest concern from an environmental justice standpoint (i.e. those with lower-income, disadvantage populations with worse health outcomes). The Manhattan CBD, which has the highest concentration of accounts in the city, would also see dramatic benefits.

Contents

Approach

Each carter's reported truck routes was mapped in order to estimate annual VMT of the existing private carting industry. This was then compared to a hypothetical zoned collection system to see how VMT would change. Four main analyses were conducted:

Existing
1. Existing Conditions Analysis
2. Optimized Existing Conditions Analysis

Zoned
3. Random Analysis
4. Clustered Analysis

The details of the approach and analyses are described in detail in this section. Figure 1 shows the overall work flow for the routing analyses, which is referenced in the following sections.

Figure 1: Work flow for VMT analysis

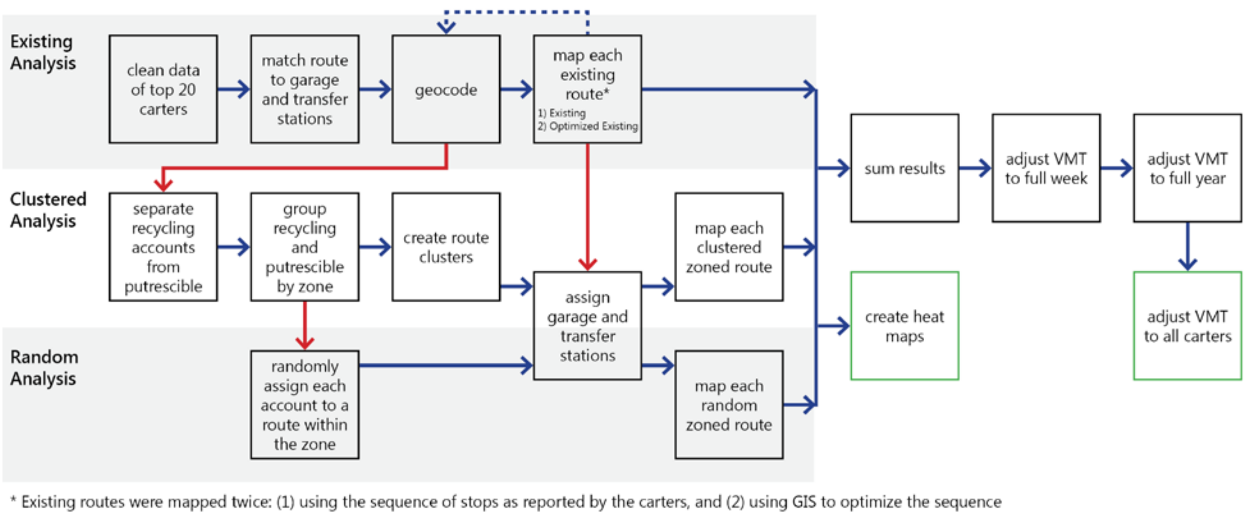

Existing Conditions Approach

In order to understand the annual VMT of the private carting industry the self-reported data set was organized, cleaned, and imported into ArcGIS so that each truck route could be mapped. However, due to the large size and varying formats of the private carter collection route data set, a subset of this data set was used for the existing conditions analysis.

The customer registry was used to rank each carter by the number of accounts served and the top twenty carters were selected for analysis. These top twenty carters represent 81% of the accounts served in the city and are listed in Table 1. A multiplier of 1.23 was used to adjust the total VMT of the analysis from 81% to represent the full set of carters. Based on our understanding from the market analysis, this approach conservatively estimates the VMT because the larger carters typically have more efficient routes as compared with the smaller carters, whose routes generally tend to be longer and more erratic on a per stop basis.

Table 1: Carter rank by accounts served according to the customer registry

Rank	Carter	Accounts Served*	Percent of Accounts	Cumulative Percent of Accounts	Number of Excel Files+
1	Carter 1	16,036	14.9%	14.9%	4
2	Carter 2	10,633	9.9%	24.7%	4
3	Carter 3	9,238	8.6%	33.3%	471
4	Carter 4	7,416	6.9%	40.2%	397
5	Carter 5	6,484	6.0%	46.2%	1
6	Carter 6	4,257	3.9%	50.2%	8
7	Carter 7	4,115	3.8%	54.0%	4
8	Carter 8	3,827	3.6%	57.5%	110
9	Carter 9	3,410	3.2%	60.7%	5
10	Carter 10	2,891	2.7%	63.4%	120
11	Carter 11	2,612	2.4%	65.8%	4
12	Carter 12	2,596	2.4%	68.2%	1
13	Carter 13	2,337	2.2%	70.4%	236
14	Carter 14	2,068	1.9%	72.3%	1
15	Carter 15	1,837	1.7%	74.0%	144
16	Carter 16	1,756	1.6%	75.6%	321
17	Carter 17	1,696	1.6%	77.2%	172
18	Carter 18	1,637	1.5%	78.7%	118
19	Carter 19	1,340	1.2%	80.0%	1
20	Carter 20	1,306	1.2%	81.2%	28
(70)	All Others	20,307	18.8%	100.0%	2,516
	Total	107,799			4,665

* According to the customer registry
+ Refers to the private carter collection route data set

BIC requested that the private carters submit four weeks of collection data as a representative sample for the entire year. In order to reduce the volume of data that would need to be cleaned, organized, and analyzed, a baseline day was selected to be mapped, which would then be adjusted to a full week. This obviates the need to analyze each day of each week of data for all twenty carters.

The Wednesday from the week in September was selected as the baseline case because it did not conflict with any major holidays (including the U.N. General Assembly), was during the school year, and was within the period of the year that New York City Department of Transportation allows traffic studies to be conducted. The week in July was not used because it was during the summer when school is out of session and activity in the city tends to be lower due to vacations. The week in December was not chosen because it coincides with the Christmas and New Year holidays. The week in January was not chosen because it was meant to represent a time period when it snowed. In other words, the September data represented the most typical week of the four weeks collected. Wednesday was chosen as the baseline day because it is in the middle of the work week and therefore the day that would least likely be impacted by people taking long weekends.

Weighted averages were calculated for each day of the week in order to adjust the baseline September-Wednesday case to a full week. Using the self-reported routing data, the number of customers served on each day of the week was determined and compared to the baseline Wednesday case. These daily totals were then divided by the customer total for Wednesday (9/10/14) to produce a ‘day multiplier’ for each day and carter. These multipliers were then multiplied by the percentage of total customers served by a given carter over the whole week in September (9/7 - 9/13). Summing the weighted multipliers of the carters for each weekday yielded a final weight for each weekday as compared to Wednesday (9/10/14). The results of this analysis are shown in Table 2.

Table 2: Weighted average of customers served for each day as compared with baseline case (9/10/2014)

Monday	Tuesday	Wednesday	Thursday	Friday	Saturday	Sunday
77.2%	102.1%	100.0%	103.3%	105.3%	74.7%	19.8%

Due to the irregularities in the reported data, not all carters could be used to determine the weights. Seven carters were selected to calculate the weights, including Carters 1 through 5, Carter 7, and Carter 19. These seven were selected because they represent over half of all accounts (51.3% according the customer registry), and for the format that their data was reported in which allowed for this analysis.

A similar methodology was used to calculate seasonal weightings from the four weeks of data provided. However, it was determined that seasonal weighting would not be used in the final analysis because the weights were close to the September data (ranging from 93% to 99%) and it was not clear from the data how to apply the weights seasonally to account for variations in the other 48 weeks of the year.

Route Creation

ArcGIS Network Analyst was used to map the routes as reported by the carters. Once the existing routes were mapped it became clear that for many routes the order in which it was reported that each account was served did not always seem to makes sense from a geographic efficiency standpoint. To account for this and other irregularities in the data, two approaches to routing the existing conditions were developed:

> **Existing Conditions Approach**: Routes were mapped as reported in the private carter collection route data set.
>
> **Optimized Existing Conditions Approach**: The same reported routes were mapped, however the sequence of the stops served was determined by GIS in order to optimize the route. For example, a route with 200 stops would still serve those same accounts but the order in which they are served would be rearranged to minimize VMT. This approach eliminates all factors that determine a route (side of street, type of trash, type of business, method of setout, time of day, carter's resources, etc.) and creates the most VMT efficient route.

The Optimized Existing approach was developed to create an idealized existing private carting system. This is useful as a comparison to the Zoned Clustered Approach (explained in the next section), which is an idealized zoned system. To be clear, the method used in this analysis to optimize the existing routes using GIS – while instructive for a comparison to a hypothetical zoned system – is not a realistic solution to reducing the VMT produced by the private carting industry in New York City because it ignores every factor that carters use in the decision-making process to create routes aside from VMT efficiency.

Zoned System Approach

Zone Creation

A hypothetical system was created in order to understand how VMT might change in a future zone-based system. A framework for the zoned system was created in conjunction with the Department of Sanitation (DSNY). The framework was designed to be a simple and logical means of testing the effects of a zoned system on VMT and is not intended to be a recommendation of how such a system should be designed. As such, it was established that for the purposes of this study the system would have the following characteristics:

- *Exclusive zones*: individual zones would be served by only one carter or entity
- *Roughly the same size*: Zones would roughly contain the same number of accounts according to the customer registry.
- *Geographic constraints*: Zones would follow community district boundaries where possible and would not cross borough boundaries.
- *Garages and transfer stations:* Each zone would be serviced by a single garage. The transfer stations that would service each zone would be based on existing patterns.
- *Routes:* Zoned routes would all be roughly the same size in terms of the number of accounts served and would have only one tip.
- *Number of zones*: There would be around ten zones, as dictated by the other constraints.

As a first step in creating the zones the customer registry was geocoded in ArcGIS in order to determine the number of accounts in each of the 59 community districts. Community districts were then grouped into zones using the framework outlined above. The resulting effort created eleven zones that are each within a single borough: one in Staten Island, two each in Bronx, Queens, and Brooklyn, and four in Manhattan. The results of the zone creation are summarized in Table 3 and shown in Figure 2.

Table 3: Each zone is made up of community districts from a single borough

Zone	Boro	Accounts*	CD1	CD2	CD3	CD4	CD5	CD6	CD7	CD8	CD9
1	SI	5,993	SI01	SI02	SI03						
2	MN	10,320	MN01	MN02	MN03						
3	MN	14,100	MN05								
4	MN	8,747	MN04	MN07	MN09	MN12					
5	MN	8,204	MN06	MN08	MN11	MN10					
6	BX	7,350	BX01	BX02	BX03	BX04	BX09	BX10			
7	BX	7,244	BX05	BX06	BX07	BX08	BX11	BX12			
8	QN	15,324	QN01	QN02	QN03	QN04	QN05	QN06			
9	QN	13,660	QN07	QN08	QN09	QN10	QN11	QN12	QN13	QN14	
10	BK	15,134	BK01	BK02	BK03	BK04	BK05	BK06	BK08	BK09	BK16
11	BK	15,551	BK07	BK10	BK11	BK12	BK13	BK14	BK15	BK17	BK18

* The total number of accounts differs from the total reported in Table 1 because an earlier version of the Customer Registry was used in the determination of the zone boundaries.

Figure 2: Community district boundaries overlaid with the zones created for this study

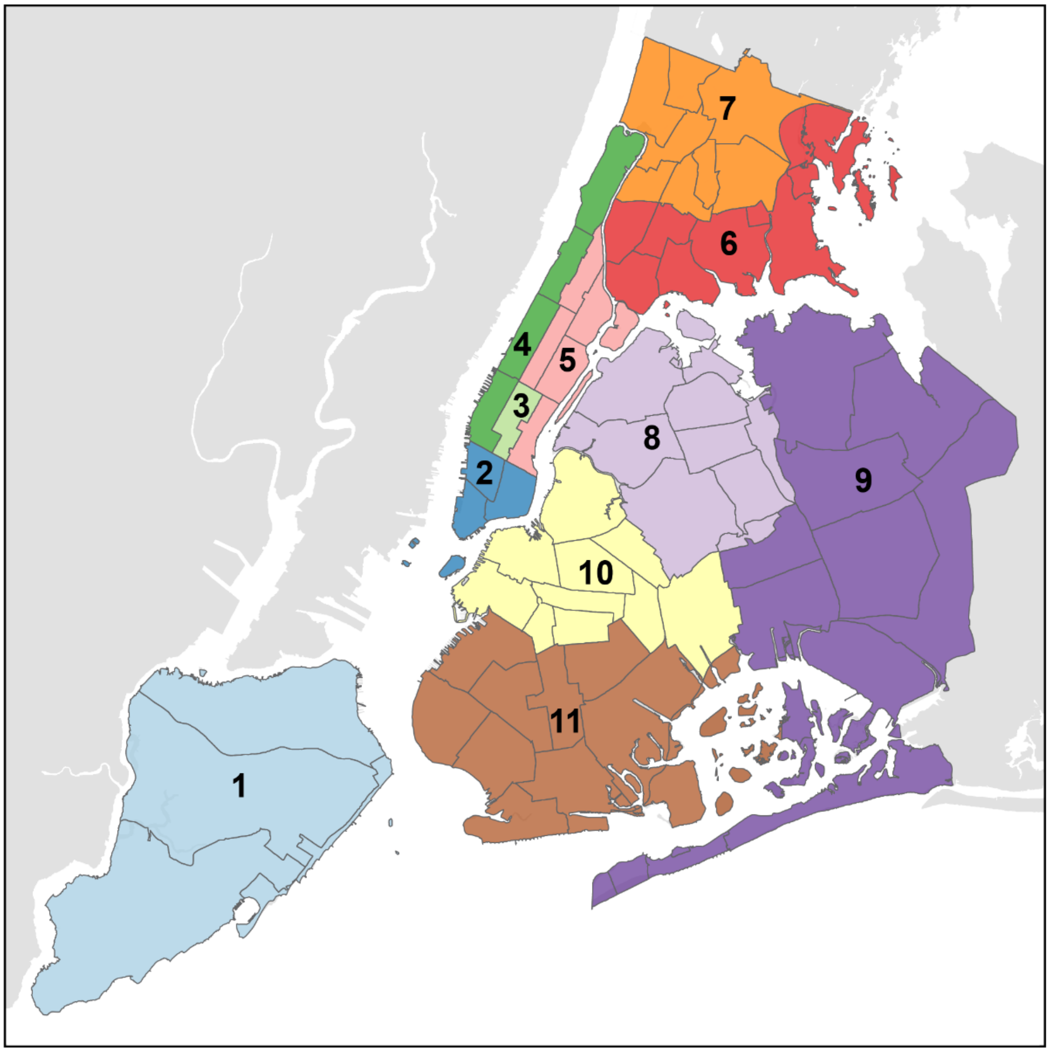

While the creation of the zones was based on the spatial distribution of accounts in the customer registry, it is important to note that the existing route data from the top twenty carters were used for the Zoned Analysis. The top twenty carters reported in their routing data that they served approximately 67,000 accounts on Wednesday, 9/10/2014. These same 67,000 accounts were then analyzed in the Zoned Analysis. A summary of how the two data sets were used appears in Table 4.

Table 4: Comparison of the Customer Registry and the Private Carter Data Set

	BIC Customer Registry	Private Carter Collection Route Data Set
Approximate number of accounts	108,000	67,000
This data was used to determine…	Which community districts were grouped together to form a zone	The VMT of the current system
	The adjustment factor to scale the top twenty carter subset to the full data set	The daily multiplier to adjust the analysis to the full week
		The accounts mapped in the Zoned Analysis
Size discrepancy is due to…	Includes all businesses in the city served by a private carter	Only includes businesses served by the top twenty carters
	May not reflect current activity of the business or the carter	Only includes businesses served on Wednesday 9/10/2014

Using the businesses listed in the private carter collection route data set for both the Existing and Zoned Analyses allows for the same businesses to be used in both analyses and therefore provides an “apples to apples” comparison between the existing and zoned systems.

Route Creation

Creating the routes for the zoned system required first that all the accounts and routes of the top twenty carters be mapped. These same businesses were then assigned to their respective zones and split into two groups: recyclable waste and putrescible waste. If the waste type was mixed or not reported it was put in the putrescible category.

It was assumed for the Zoned Analysis that the zoned routes would each serve roughly the same number of accounts. An average of 200 to 250 accounts per route was calculated for the largest carters in order to create a baseline for the hypothetical route length for the zoned system.

Two approaches were used to create routes for the hypothetical zones:

> **Random Approach**: The recycling and putrescible accounts in each zone were randomly assigned to a route such that each route in the zone was between 200 and 250 accounts. This approach more closely approximates the geographic spread of existing routes but is constrained within the zone. A given zone would still have the same number of routes as in the Clustered Approach, but this approach accounts for waste type directly and other complicating factors (time of pickup, volume of waste, side of street, etc.) indirectly by introducing randomness to the route which artificially makes the routes less efficient than in the Clustered Approach.

> **Clustered Approach**: Routes were created by clustering accounts that were near each other into a single route of between 200 and 250 accounts. Both recycling and putrescible routes were created on a zone by zone basis. This approach assumes that businesses near each other could all be served on the same route and is similar to how DSNY picks up residential waste. This approach was created as a parallel to the Optimized Existing Approach.

A sample comparison of the Existing Analyses with the Zoned Analyses is illustrated in Figure 3. The existing route is not constrained by a zone and the order in which accounts are served is set by the hauler, creating a geographically dispersed route. The optimized existing uses the same accounts as in the existing route, but the route is optimized which creates less overlap. The clustered route covers one discrete area in Zone 6 while the random route is spread throughout the zone.

Figure 3: A comparison of an existing and random sample route, and an optimized existing and clustered sample route in Zone 6

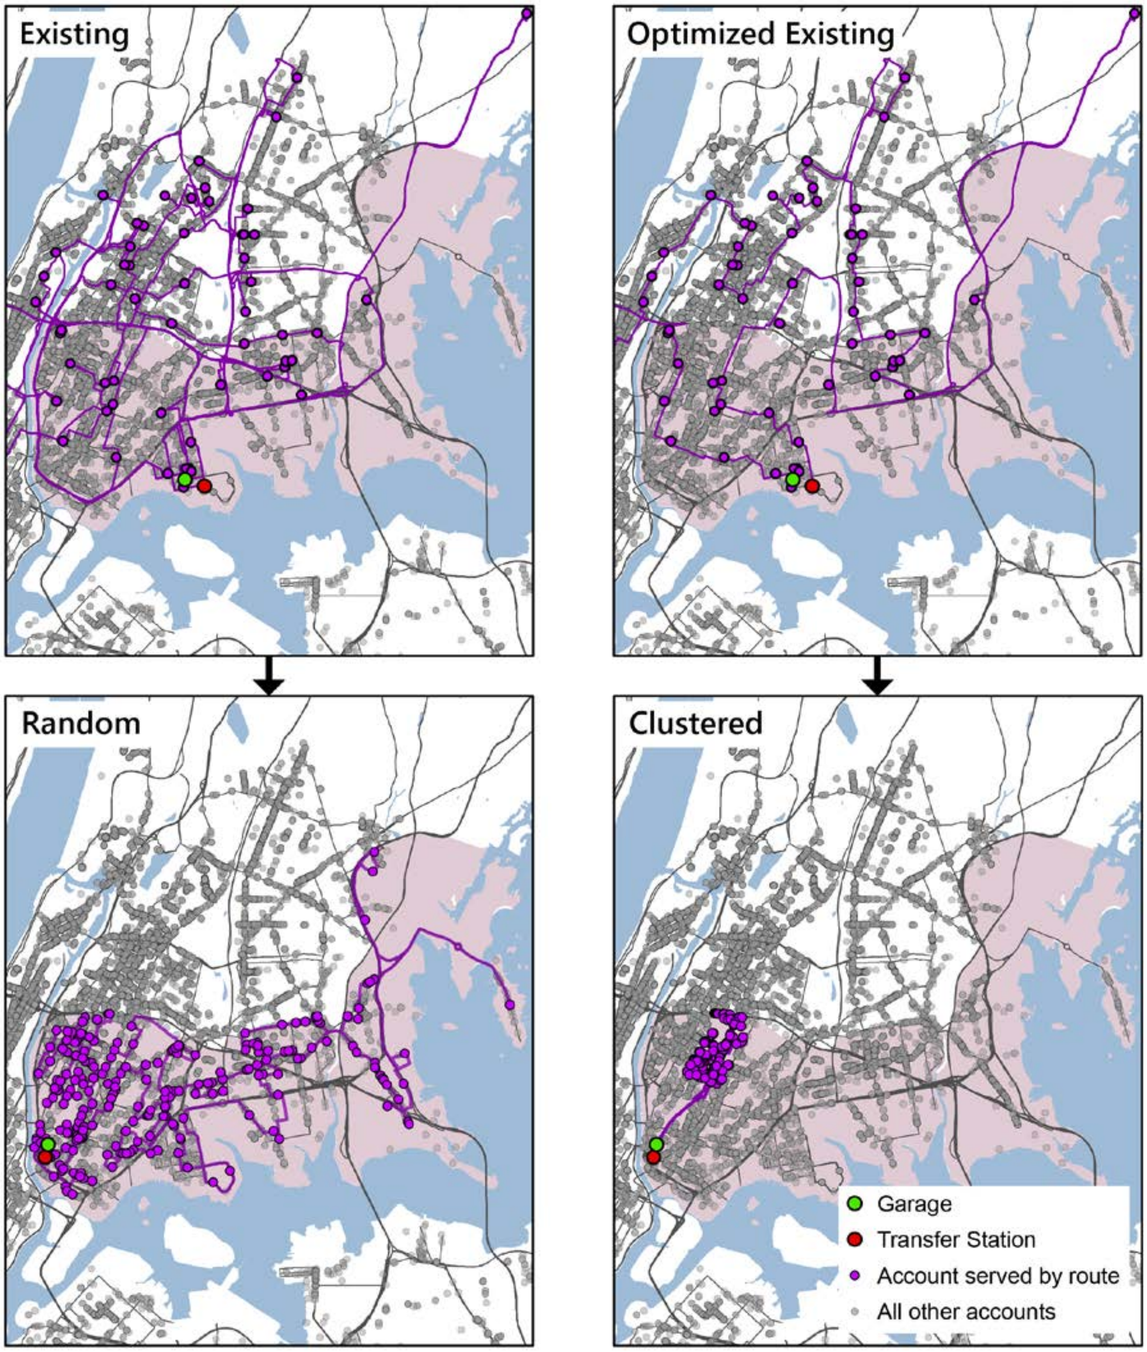

The Clustered Approach offers an idealized version of a zoned system, while the Random Approach offers a more conservative method. In zones that cover larger areas, the difference between the Clustered Approach and the Random Approach will become larger, representing a range of possible VMT reductions. Taken together with the Existing and Optimized Existing Analyses they offer a range of the likely VMT reductions in implementing a zoned system in New York City.

Heat Maps

Heat maps of the density of VMT were made for each of the analyses. These density heat maps are graphical representations (i.e. raster image) of the number of vehicle miles traveled per square mile. Heat maps can be visually inspected to understand where the VMT density is high (i.e. where the impact of overlapping and inefficient trucks routes is the highest) and low. In addition, one heat map can be 'subtracted' from another (i.e. map algebra) to show the difference in densities between the two. This is useful to visualize the difference in density of VMT between the Existing Condition and the Random Analysis, and between the Optimized Existing Condition and the Clustered Analysis. Six heat maps were created in total:

1. Existing Analysis — Figure 8
2. Optimized Existing Analysis — Figure 9
3. Random Analysis — Figure 10
4. Clustered Analysis — Figure 11
5. The difference between the Existing and Random Analyses — Figure 12
6. The difference between the Optimized Existing and the Clustered Analyses — Figure 13

A kernel density analysis was used in ArcGIS to create the heat maps. A cell size of 1,320 feet (0.25 miles) and a search radius of 5,280 feet (1 mile) were used. The same extent was used to create each raster so that map algebra could be used to create heat maps 5 and 6 in the list above.

Data Cleaning and Organizing

The initial data sets from the top twenty carters came in a variety of formats. A combination of Python coding language, Excel, and ArcGIS was used to clean and organize the data. The final product of the data organization was twenty geodatabases (one for each carter) that could be imported into ArcGIS so that each route could be mapped.

Two different Excel templates were sent to the carters by BIC with a suggested data storage template. The first format had four tabs while the second had two. The two templates asked for the same information, however carters could choose which template to use. Despite having two templates, the carters returned the data in a variety of formats that were not wholly consistent with the formatting, which complicated the organization and cleaning of the data. The templates collected the following information:

1. *Load Information*: transfer station and garage information for each route
2. *Truck Information*: volume capacity, model, and VIN
3. *Customer Information*: a list of each account, their address, the sequence in which they were served, and which route served them
4. *Material Information*: material and container type for each route

The customer information provides the critical routing information that was used to map each route in ArcGIS. Most carters submitted data with two tabs (material and customer information), with most of the required data combined within these two tabs in varying formats. Some carters included three or all four tabs, while some had only one tab that included all the required data separated by spaces or other formatting irregularities. Additionally, some carters had separate sheets or tabs for each route following this two-tab format.

Python

In order to work with this large and inconsistent set of data it became necessary to design a set of scripts to parse and import the various formats into a workable GIS database. In general, these scripts performed two primary tasks: first to combine all route worksheets into one workbook with two tabs, and second to join the material data (garage and transfer station for each route) to the customer data (customer addresses and sequence). The first task was accomplished by using the *openpyxl* Python library to parse and combine a series of workbooks into one digestible merged workbook. In some cases, route data (i.e. route name and date) were only included in the file name. In these cases, the script would tag each line in each sheet with the file name containing the necessary information. These scripts are included in the project data package and are named ‘merge_excel.py’ with some containing carter names when they have been specialized for a certain carter’s unique data format.

The second task was accomplished using *openpyxl* as well as the tools built into *arcpy* (the built-in Python library for ArcGIS). This script first determines the type of each column in each sheet of the input workbook and creates respective fields in a material and customer table in GIS. The two workbook sheets are then parsed and saved in dictionaries. The customer dictionary is written into the customer GIS table using the same format and structure as the import workbook. Then the fields from the material sheet are added to the customer table. Each line of the customer data is matched to its respective material (route) data using the route name as the matching key. The material data is written and a final table for a carter is saved.

In some cases, the first step of this process (merging excel workbooks) is not necessary as some carters submitted workbooks that already had this data combined. In a few cases the data was formatted in such an unmanageable way that cleaning had to be done manually.

Information on usage of openpyxl can be found here: https://openpyxl.readthedocs.org/en/2.3.3/

Information on usage of arcpy can be found here: http://help.arcgis.com/en/arcgisdesktop/10.0/help/index.html?v=3&q=arcpy.mapping#/What_is_ArcPy/000v000000v7000000/

Garage and Transfer Stations

Almost all of the top twenty carters stored their garage and transfer information separately from the customer sequence information. In order to include these stops in the existing conditions analysis it was necessary to manually add a start and end point for each route corresponding to the origin garage and ending transfer station, respectively.

Multiple Loads per Route

Several routes tipped more than once, however the point in the route that the first tip occurred was not included in the data provided. It was assumed that the first tip was in the middle of the route by number of accounts served. This issue effected data from eight of the twenty companies.

Geocoding

ArcGIS StreetMap Premium dataset was used as the main platform to perform the geocoding process. The StreetMap dataset includes the building address information, street network, and turning restriction data for the entire United States for the year 2015. Due to the numerous discrepancies in the route dataset there were a number of addresses that were not geocoded using the ArcGIS geocoding tool. These discrepancies included errors in ZIP codes, city names, street names, and building numbers. GIS tools and techniques were used to minimize the number of unmatched addresses including the use of Google API in the QGIS software in addition to the use of the StreetMap software.

The geocoding results are summarized in Table 5, which show that nearly 98% of addresses were matched. When geocoding, a score is calculated for each address by the software, with a higher score meaning a better match to an actual address. A score of 80 or higher (on a scale from 0 to 100) is given a ‘match.’ Tied addresses (1.0% of the data set) are those which matched more than one location and were assigned to the address with the higher score, as determined by ArcGIS. Unmatched addresses had scores less than 80, and represent 1.2% of the data set analyzed.

Table 5: Geocoding match rate for the top twenty carters

Carter	Total Accounts*	Matched		Tied		Unmatched	
		#	%	#	%	#	%
Carter 1	16,992	16,523	97.2%	150	0.9%	319	1.9%
Carter 2	12,610	12,435	98.6%	101	0.8%	74	0.6%
Carter 3	2,696	2,696	100.0%	0	0.0%	0	0.0%
Carter 4	4,362	4,263	97.7%	21	0.5%	78	1.8%
Carter 5	6,569	6,367	96.9%	33	0.5%	169	2.6%
Carter 6	444	433	97.5%	1	0.2%	10	2.3%
Carter 7	2,640	2,420	91.7%	200	7.6%	20	0.8%
Carter 8	3,639	3,575	98.2%	40	1.1%	24	0.7%
Carter 9	1,301	1,285	98.8%	13	1.0%	3	0.2%
Carter 10	1,218	1,214	99.7%	0	0.0%	4	0.3%
Carter 11	1,286	1,273	99.0%	7	0.5%	6	0.5%
Carter 12	1,228	1,225	99.8%	3	0.2%	0	0.0%
Carter 13	3,448	3,359	97.4%	55	1.6%	34	1.0%
Carter 14	2,441	2,421	99.2%	2	0.1%	18	0.7%
Carter 15	959	954	99.5%	5	0.5%	0	0.0%
Carter 16	1,017	983	96.7%	34	3.3%	0	0.0%
Carter 17	2,662	2,596	97.5%	21	0.8%	45	1.7%
Carter 18	1,307	1,303	99.7%	0	0.0%	4	0.3%
Carter 19	562	557	99.1%	1	0.2%	4	0.7%
Carter 20	436	418	95.9%	12	2.8%	6	1.4%
Total	67,817	66,300	97.8%	699	1.0%	818	1.2%

* According to private carter collection route data set

Unmatched addresses were removed from the analysis, resulting in about 67,000 accounts overall that were in included in the routing analysis. Due to the sheer volume of accounts it was not possible to check each unmatched address to see if it could be manually geocoded. Similarly, tied addresses were not manually checked to see if the correct address was chosen in the tie. It is also possible that matched addresses were assigned to the wrong location due to errors with the data but also due to the geocoding software.

Incorrectly geocoded addresses would disproportionately impact the existing analysis because the routes and sequence are provided by the carters. For example, if the second stop of a route is supposed to be on the same street as the first and third stop, but was incorrectly geocoded to another location, that could greatly increase that route's VMT. This same problem would have little impact on the Zoned Analysis because it would be treated as just another stop within that zone. This would nominally increase the VMT for the Random Analysis and would have even less impact on the Clustered Analysis due to the way routes were created.

Efforts to fix geocoding errors were focused on existing routes that had a VMT of over 250 and contained tied addresses (this affected 70 of the 340 existing routes). In these cases routes were examined and accounts that were in the wrong location were manually re-geocoded to the correct location. Table 5 reflects the corrections made through manual geocoding.

Routing Methodology

The organized geodatabases for each carter being studied were imported into ArcGIS so that a discrete route length and travel time could be determined for each route. The ArcGIS Network Analyst extension provides network-based spatial analysis tools for solving complex routing problems. The extension uses a configurable transportation network data model, allowing users to accurately represent their respective unique network requirements. The software enables users to plan routes for an entire fleet, calculate drive-times, locate facilities, and solve other network-related problems.

StreetMap Premium, a network dataset that works in conjunction with ArcGIS Network Analyst extension, was used to provide high-quality cartographic display for geocoding and routing. StreetMap Premium is based on commercial street reference data from global and local street data suppliers, NAVTEQ and TomTom, for the North America roadways network coverage. When used with ArcGIS Network Analyst, this dataset can account for historic traffic and trucking restrictions data to provide route model outputs.

Existing Conditions

Private carting routes were documented and created based on the sequence order and customer accounts information from the carter dataset and using the ArcGIS Network Analyst extension. To simulate and account for certain conditions, restrictions and setting changes were applied within the ArcGIS Network Analyst interface such that during the analysis traversing the identified restricted elements would be prohibited entirely, avoided, or preferred. Restriction attribute is defined using a Boolean data type such that each network element either has the restriction (the Boolean value evaluates to true) or doesn't have the restriction (the Boolean value evaluates to false). For the purposes of this analysis, restrictions were specified as “driving a truck”, which applies a parameter for routing outputs to be restricted to truck routes. However, since not all customer pickup locations were accessible via designated truck routes a further adjustment was made in the attribute parameters that changed the “driving a truck” restriction from “prohibited” to “avoid: low”, which relaxed the prohibit restriction such that Network Analyst could route each truck to each account. This accommodates customer pickup locations located on non-designated truck routes.

Using the above parameters and the self-reported route data, existing private carting routes were geocoded, routed, and visualized spatially. Individual route information such as VMT and time travelled were calculated by ArcGIS and summarized in Excel. For each route, 34 seconds were added per customer pickup location to account for required collection and service time, based on similar estimates employed by DSNY. This total “travel + service time” was divided by the miles traveled to calculate the average speed of the route.

The same process was used to create the Optimized Existing routes, except that Network Analyst was used to optimize the order in which the accounts were served.

Zoned System

The set of accounts used in the Zone Analysis are the same ones that were extracted from the Existing Conditions. This helped to ensure consistency and compatibility across data sources. Customer accounts were separated first by material type (putrescible waste and recycling) and assigned to their respective zones. Two approaches to the Collection Zone methodology (described below) were developed to illustrate a range of potential VMT results and impacts.

The process of creating truck routes under the Collection Zone Condition is similar to the process described under Existing Condition. The restriction settings and attribute parameters remain identical to the Existing Condition while the order of customer account pickup sequence is reordered to allow ArcGIS Network Analyst to find the optimal route based on minimizing travel time, measured in minutes. This routing methodology was applied to both the Clustered and the Random Analyses.

Assigning Garages and Transfer Stations for the Zoned Analysis

The assignment of garages and transfer stations for the Zoned Analysis was based on the existing established relationship between routes, their garages, and transfer stations. The assignment of a garage to each zone was as follows:

1. Garages from the Existing Conditions were filtered by borough and sorted by how many accounts were served from that garage.
    - No garages used by the top twenty carters were in Manhattan.
2. Within each borough (with the exception of Manhattan) garages were ranked by the number of accounts served.
3. The highest ranked garages in each borough were assigned to a zone within the same borough, with the exception of Manhattan.
4. The next most popular garages, including garages in New Jersey, were then selected to serve the four zones in Manhattan.

The assignment of transfer stations to routes was based on the established relationship between garages and transfer stations from the Existing Conditions. For example, if Garage A has twelve routes in the Existing Conditions, and one-third go to Transfer Station 1 and the other two-thirds go to Transfer Station 2, then that same ratio would be used in the Zoned Analysis. The garage and transfer station assignments are summarized in Table 6.

Table 6: Garage and transfer station assignment by zone

Zone	Total Routes	Garage	Transfer Station	Number of Zoned Routes
1	10	22 Van St, Staten Island, NY	1499 US 1, Rahway, NJ, 07065	7
			183 Raymond Boulevard, Newark, NJ	3
2	38	73-10 Edsall Ave, Glendale, NY	130 Varick Ave, Brooklyn, NY	28
			141 6th Street, Brooklyn, NY	2
			172-08 Douglas Ave, Jamaica, NY	2
			492 Scholes St, Brooklyn, NY	6
3	43	126-46 34th Ave, Flushing, NY	130 Varick Ave, Brooklyn, NY	6
			172-08 Douglas Ave, Jamaica, NY	3
			3 Railroad Pl, Maspeth, NY	9
			600 Merchants Concourse, Westbury, NY	25
4	28	451 Frelinghuysen Ave, Newark, NJ	183 Raymond Boulevard, Newark, NJ	23
			920 East 132nd Street, Bronx, NY	5
5	16	421 Manida St, Bronx, NY	1221 East Bay Ave, Bronx, NY	3
			287 Halleck St, Bronx, NY	13
6	20	2630 Park Ave, Bronx, NY	2 Hope St, Jersey City, NJ	7
			227 Rider Ave, Bronx, NY	13
7	30	315 Casanova St, Bronx, NY	183 Raymond Boulevard, Newark, NJ	17
			55-06 43rd St, Flushing, NY	1
			920 East 132nd Street, Bronx, NY	12
8	29	58-35 47th Street, Maspeth, NY	115 Thames St, Brooklyn, NY	2
			183 Raymond Boulevard, Newark, NJ	6
			3 Railroad Pl, Maspeth, NY	2
			485 Scott Ave, Brooklyn, NY	14
			770 Barry Street, Bronx, NY	3
			920 East 132nd Street, Bronx, NY	2
9	17	168-10 Douglas Ave, Glendale, NY	172-08 Douglas Ave, Jamaica, NY	10

			187-40 Hollis Ave, Jamaica, NY	7
10	28	70 Hamilton Ave, Brooklyn, NY	110-120 50th St, Brooklyn, NY	2
			2 Hope St, Jersey City, NJ	2
			213 Broadway, Jersey City, NJ	1
			287 Halleck St, Bronx, NY	2
			31-33 Farrington St, Flushing, NY	1
			325 Faile St, Bronx, NY	1
			577 Court St, Brooklyn, NY	15
			920 East 132nd Street, Bronx, NY	2
			98 Lincoln Ave, Bronx, NY	2
11	34	505 Cozine Ave, Brooklyn, NY	183 Raymond Boulevard, Newark, NJ	2
			75 Thomas St, Brooklyn, NY	15
			854 Shepherd Ave, Brooklyn, NY	3
			920 East 132nd Street, Bronx, NY	2
			941 Stanley Ave, Brooklyn, NY	12

Existing Conditions Results

All existing accounts from the top twenty carters were mapped in ArcGIS, resulting in about 67,000 businesses spread across the five boroughs. As was expected, the largest carters have accounts in four or more boroughs, while some of the smaller carters that were analyzed have much more discrete ranges that focus on one borough or even one part of a borough. In order to understand how the ranges of the carters were distributed, and how they overlapped, a directional distribution analysis was conducted in ArcGIS. This analysis creates a standard deviational ellipse of the accounts served by each carter, which summarizes the following spatial characteristics: central tendency, dispersion, and directional trends. An example of a standard deviational ellipse is shown in Figure 4.

Figure 4: Example illustration of a standard deviational ellipse (source: Esri)

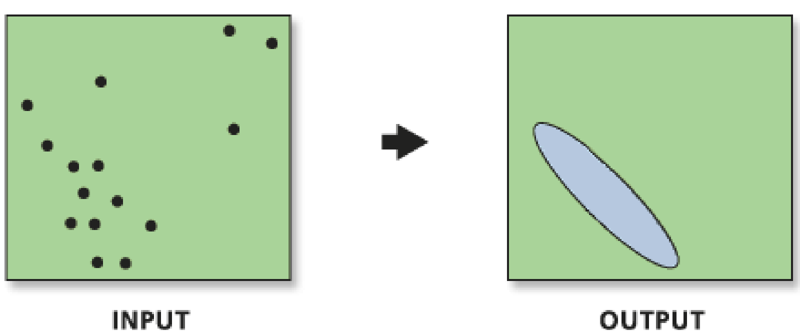

Each ellipse does not cover the full range of accounts served by a single carter, however it does show where they are concentrated. The larger the ellipse the more dispersed the accounts are. Figure 5 shows the standard deviational ellipses of the top twenty carters, illustrating how much overlap exists between these carters. The overlap is particularly acute in much of Manhattan, central and western Queens, and North Brooklyn. This overlap leads to increased VMT, a disproportionate impact of truck traffic on certain neighborhoods, and an overall less efficient carting system as compared with a system with less overlap.

Figure 5: The central tendency and dispersion of the businesses served by the top twenty carters, showing significant overlap between carters

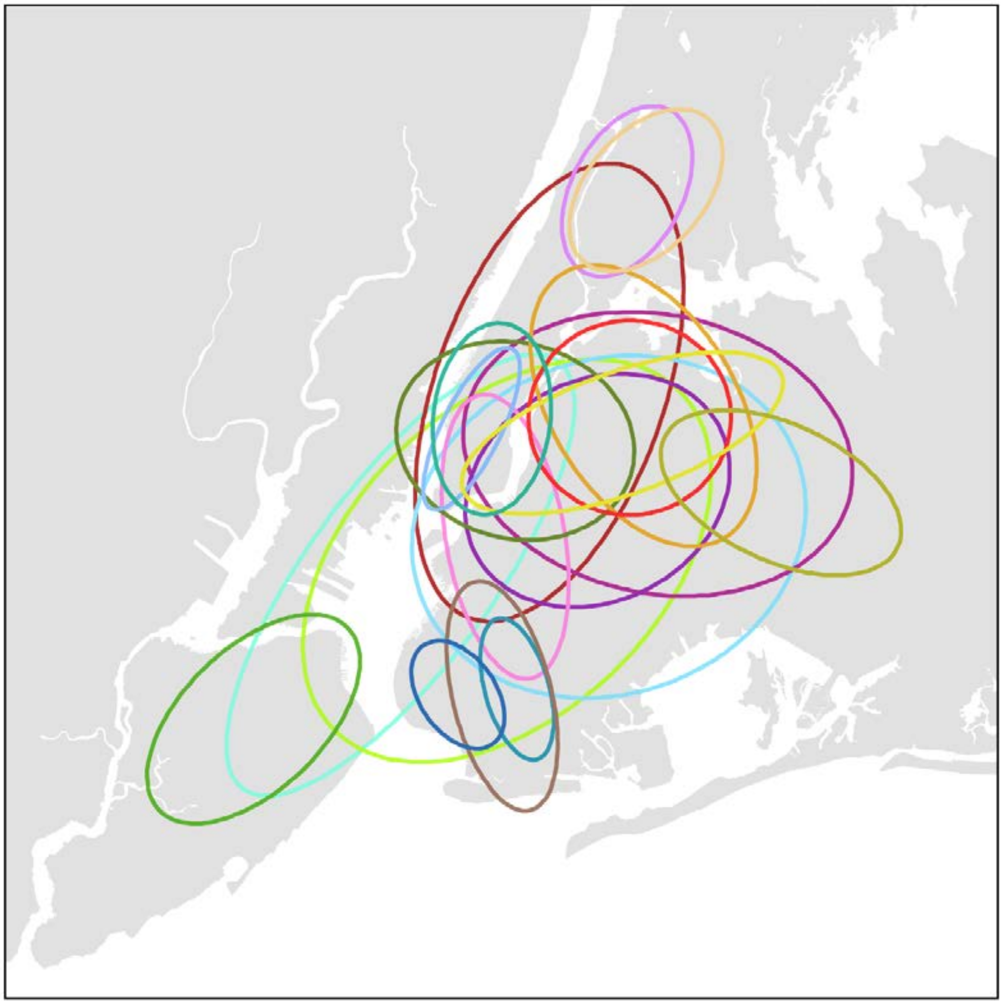

Existing Routes

All existing routes from the top twenty carters were routed and mapped in ArcGIS, resulting in 340 existing routes. Many of the routes that were mapped were unrealistically long in terms of the number of accounts served or the VMT. The number of accounts per route were reported by the carters and it was outside the scope of this study to understand why these routes were so long or if it was a reporting error. The histogram in Figure 6 shows the distribution of accounts per route. The majority of routes (72%) were under 250 accounts per route.

Figure 6: Histogram of Existing Accounts/Route

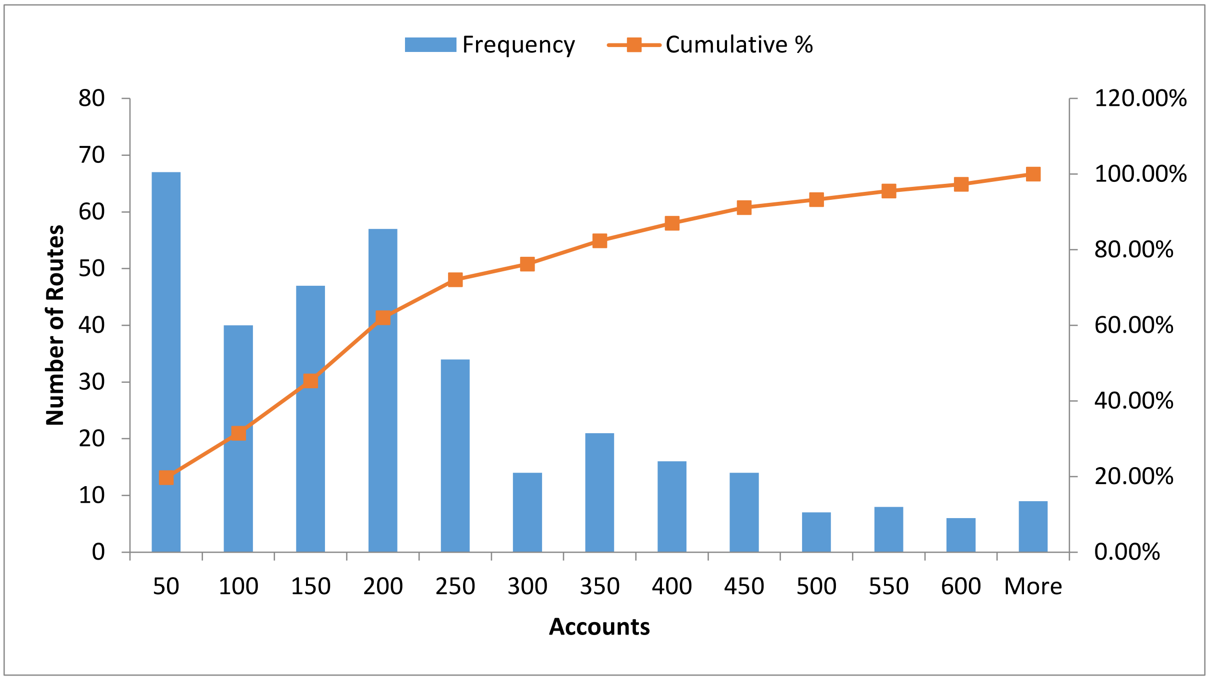

Figure 7 shows a histogram of the VMT of existing routes. The majority of routes (82%) were under 250 VMT. The routes with more than 250 VMT generally corresponded to the routes that had a high number of accounts per route.

Figure 7: Histogram of Existing VMT/Route

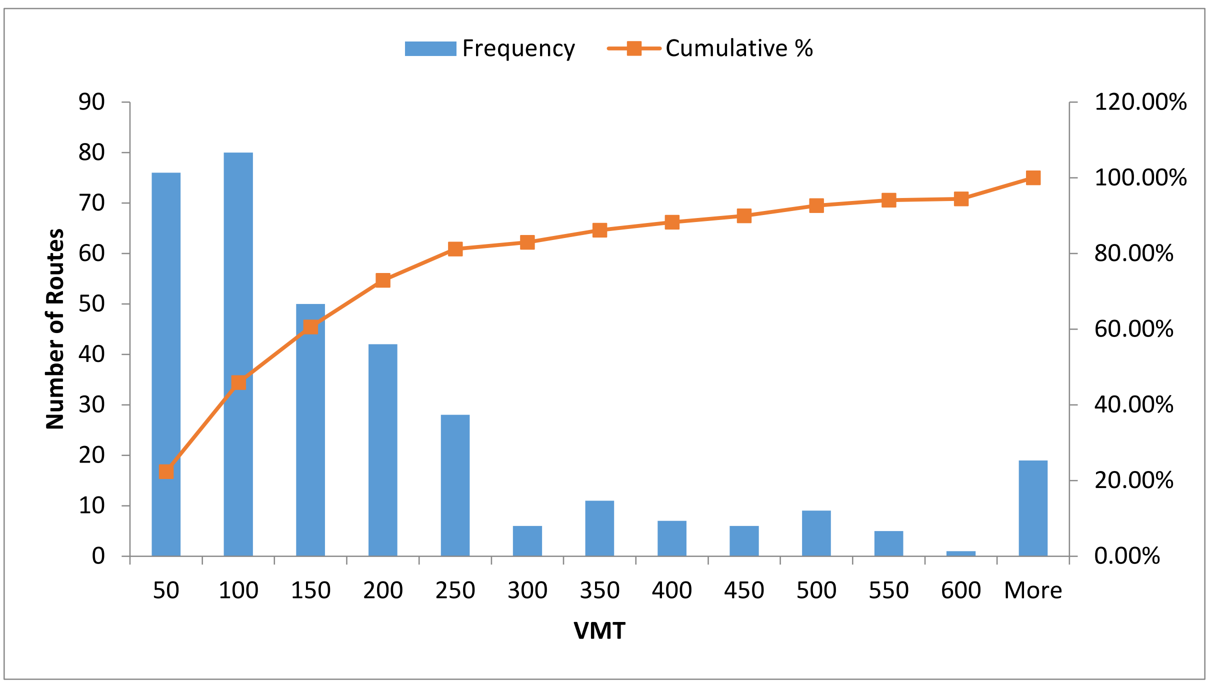

The total VMT of the 340 routes that were mapped for Wednesday, September 10, 2014 was about 61,900 miles. Using the established weights, this total was converted to a weekly and then annual amount, for a total of 23.1 million annual vehicle miles traveled. The results are summarized in Table 7.

Table 7: Summary of Existing Conditions routing results

	VMT (miles)	# of accounts	Travel + Service Time (hours)	Speed (mph)	Accounts / VMT
Total	61,928	67,772	N/A	N/A	N/A
Median	115	165	7.4	16.1	1.0
Annual	**23,100,000**				

The number of accounts served per vehicle mile traveled is an important indicator of route efficiency. The higher the number the more efficient a route is, picking up waste from more accounts in less distance.

Figure 8 shows the density of VMT of the routes from the Existing Conditions Analysis. There are heavy concentrations of VMT (and therefore trash trucks) in all of Manhattan and the Bronx, along the Gowanus, Brooklyn-Queens, and Long Island expressways, and in parts of South Brooklyn and central Queens.

Figure 8: VMT density of Existing Routes

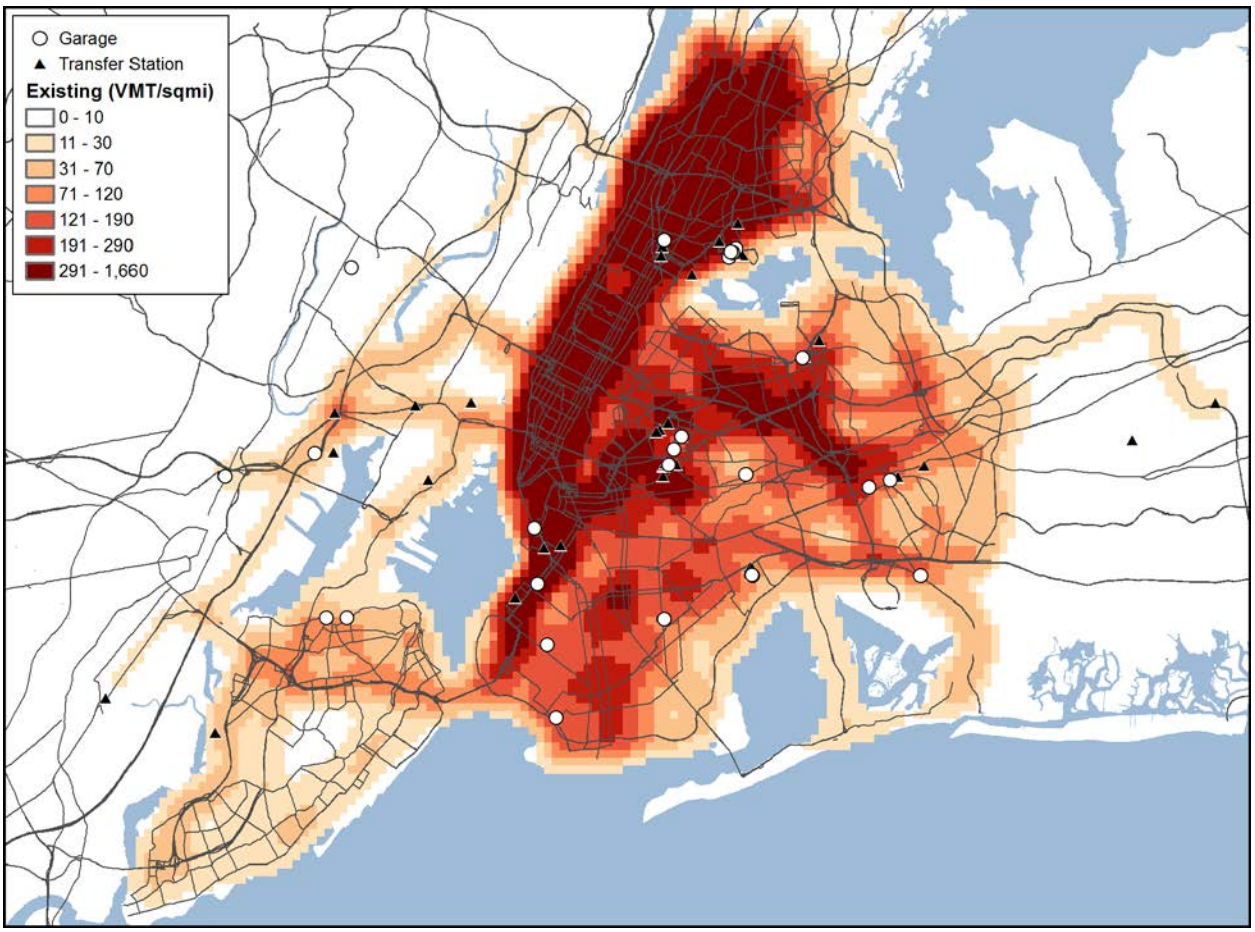

Optimized Existing Routes

The Optimized Existing Analysis shows a significant decrease in VMT from the Existing Conditions Analysis. The total VMT for the optimized routes was about 16,300 miles per Wednesday. Using the established weights, this total was converted to a weekly and then annual amount, for a total of 6.1 million annual vehicle miles traveled. The results are summarized in Table 8.

Table 8: Summary of Optimized Existing Conditions routing results

	VMT (miles)	# of accounts	Travel + Service Time (hours)	Speed (mph)	Accounts / VMT
Total	16,261	67,772	N/A	N/A	N/A
Median	46	165	4.3	11.9	3.2
Annual	**6,100,000**				

While the median route VMT and travel time are unrealistically low, this is expected since the optimization process creates an idealized system that is a useful comparison to the Clustered Analysis. It is important to note that this is not an indication that the existing private carting system could be made so efficient by optimized routing that annual VMT could be reduced from 23.1 to 6.1 million miles. The Existing and Optimized Existing scenarios cannot be compared because they represent two distinct analyses that are comparable to the Random and Clustered Analyses, respectively.

Figure 9 shows the density of VMT of the routes from the Optimized Existing Conditions Analysis. There are heavy concentrations of VMT in the Manhattan central business district (CBD) – where there is also a high concentration of businesses served by private carters – and along the major highways in the Bronx, Brooklyn, and Queens. VMT is also concentrated around the clusters of garages and transfer stations in Sunset Park, North Brooklyn/Maspeth, and the South Bronx (which is also where there are major highways).

Figure 9: VMT density of Optimized Existing Routes

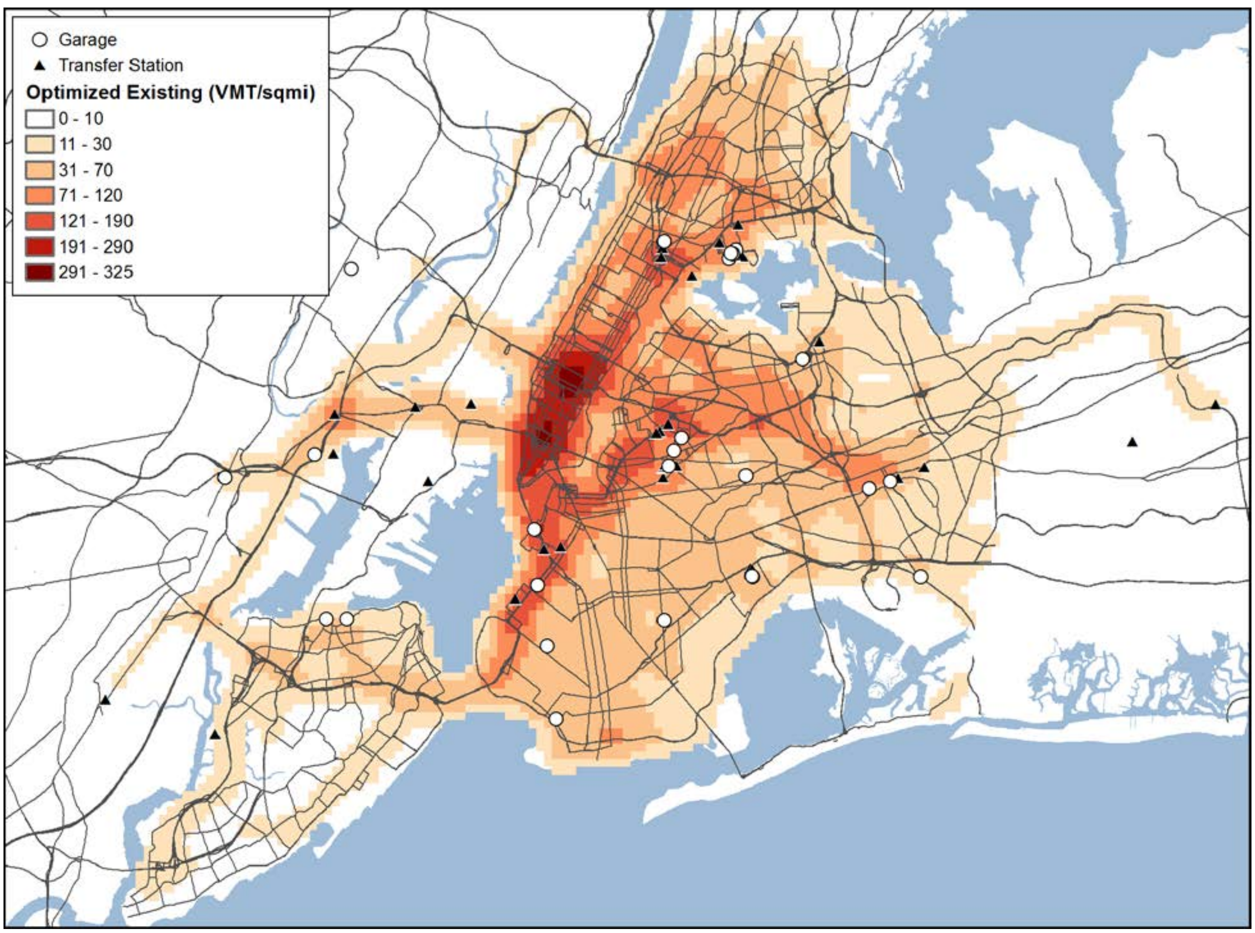

Zoned System Results

The Existing Conditions analysis used the routes as they were reported by the carters, and the Optimized Existing Analysis used those same routes but optimized them. As such, there were the same number of routes in each Existing Conditions Analysis (340) and each route had the same number of accounts between analyses. For example, Carter 1, Route 1 served 217 accounts in both the Existing and Optimized Existing Analyses. However, the number of accounts in a route varied widely, ranging from single roll-off routes to routes with hundreds of accounts.

In the Zoned System Analyses it was assumed that routes would be roughly the same size, between 200 and 250 accounts. This created more uniform routes across the board and eliminated very short and very long routes, resulting in a total of 293 routes, which is 47 fewer routes than in the Existing Conditions. The assumption of roughly equally sized routes has an inherent efficiency built into it, reducing the number of truck routes and therefore VMT. This is reflected in the following results.

Random Routes

The Random Analysis created routes by randomly assigning accounts to a route within a single zone. This analysis was conducted as a comparison to the Existing Conditions. The Random Analysis shows a significant decrease in VMT from the Existing Conditions. The total VMT for the random routes was about 7.4 million miles per year, which is a 68% reduction from the Existing Conditions. The results are summarized in Table 9.

Table 9: Summary of Random Analysis routing results

	VMT (miles)	# of accounts	Travel + Service Time (hours)	Speed (mph)	Accounts / VMT
Total	19,817	66,934	N/A	N/A	N/A
Median	66	229	6.3	10.9	3.4
Annual	**7,400,000**				

The accounts served per VMT for the Random Analysis is approximately 3.4 times greater than the Existing Conditions. This indicates that even though the routes are random, that by constraining them geographically and routing them efficiently, significant VMT savings and efficiency gains can be made.

Figure 10 shows the density of VMT of the routes from the Random Analysis. There are heavy concentrations of VMT in the Manhattan CBD and along the major highways, however to a much lesser degree when compared to the Existing Conditions.

Figure 10: VMT density of Random Routes

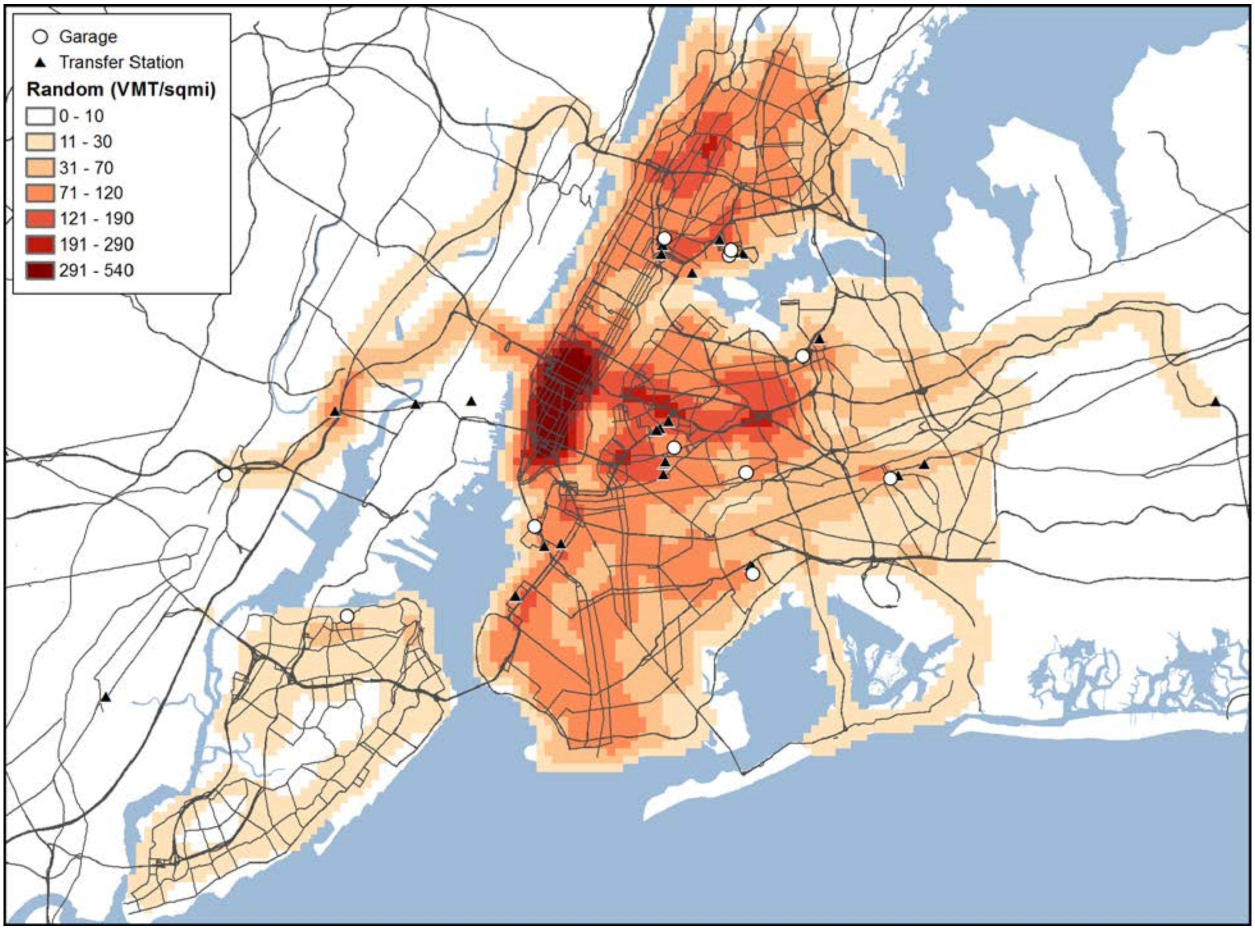

Clustered Routes

The Clustered Analysis created routes by grouping geographically proximate accounts into a route, creating an idealized zoned system. This analysis provides a comparison for the Optimized Existing Conditions. The Clustered Analysis shows a significant decrease in VMT from the Optimized Existing Conditions. The total VMT for the clustered routes was about 3.1 million miles per year, which is a 49% reduction from the Optimized Existing Conditions. The results are summarized in Table 10.

Table 10: Summary of Clustered Analysis routing results

	VMT (miles)	# of accounts	Travel + Service Time (hours)	Speed (mph)	Accounts / VMT
Total	8,224	66,934	N/A	N/A	N/A
Median	26	231	3.3	7.6	8.7
Annual	**3,100,000**				

The accounts served per VMT for the Clustered Analysis is approximately 2.7 times greater than the Optimized Existing Conditions. This indicates that even though the existing routes were optimized, that by geographically constraining them additional efficiencies can be gained.

Figure 11 shows the density of VMT of the routes from the Clustered Analysis. VMT is most concentrated along the Long Island Expressway (LIE) and Brooklyn-Queens Expressway (BQE), as well as in the South Bronx and parts of the Manhattan CBD.

Figure 11: VMT density of Clustered Routes

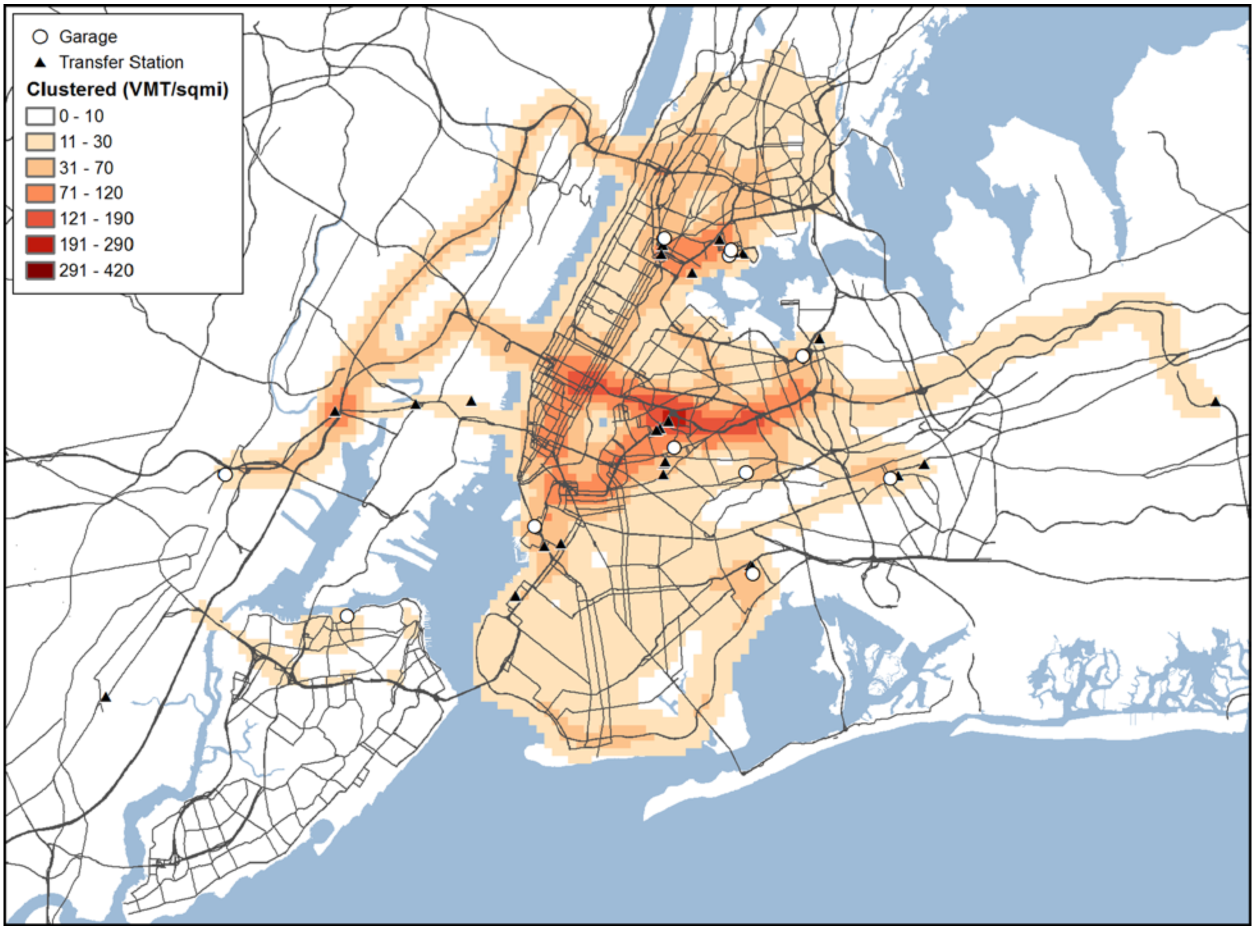

Analysis Comparison

Existing Conditions and Random Analysis

When the Random Analysis is compared with the Existing Conditions, there is a 68% annual VMT reduction. While the median number of accounts per route increased from 165 to 229, route times and VMT decreased owing to the increased geographic efficiency of the zoned system.

Table 11: Summary comparison of Existing and Random Analyses

	Annual VMT (million miles)	Median VMT/Route	Median Accounts/Route	Median Route Time (hours)	Median Route Speed (mph)	Median Accounts/VMT
Existing	23.1	115	165	7.4	16.1	1.0
Random	7.4	66	229	6.3	10.9	3.4

The heat map for the Random Analysis was subtracted from that of the Existing Conditions, resulting in a new heat map that shows how the density of VMT changed across the city from one condition to the other. Overall, VMT density was reduced nearly everywhere in the city. In Figure 12, the dark green areas show where the greatest VMT savings occurred: Manhattan, the Bronx, and along the BQE and LIE. The slight increases in VMT density (shown in pink) are mainly along the routes outside of the city that lead to the transfer stations in Nassau County and New Jersey. These routes saw an increase in VMT because the number of transfer stations used in the Existing Analyses was reduced in the Zoned Analyses. Thus, while overall VMT citywide decreased, more trucks did go to certain locations as compared with the Existing Conditions. These changes are a direct result of the way garages and transfer stations were assigned to the zones and do not represent a significant increase in VMT (maximum of 20 VMT/square mile).

Figure 12: Change in VMT density between the Existing and Random Routes

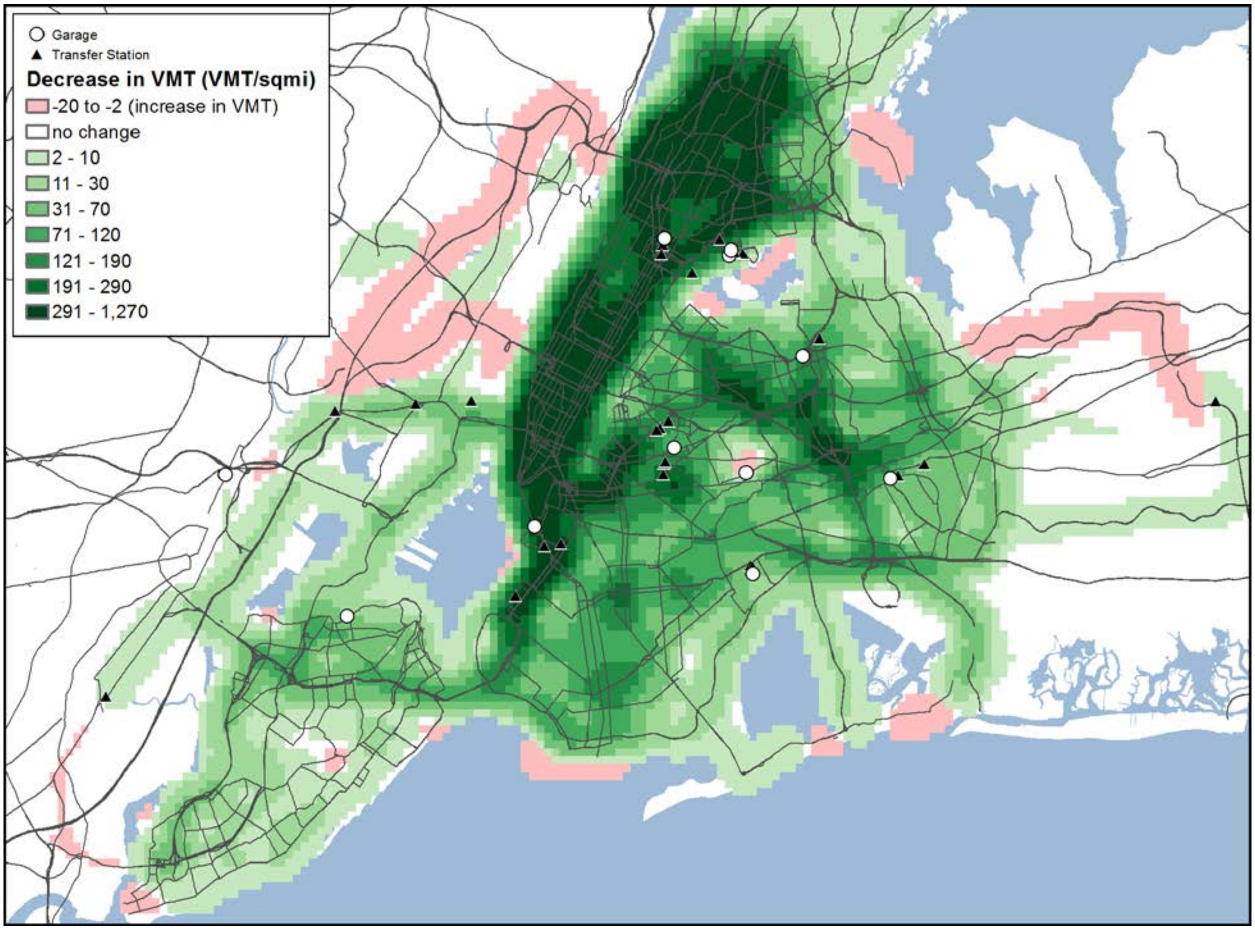

Optimized Existing Conditions and Clustered Analysis

When the Clustered Analysis is compared with the Optimized Existing Conditions, there is a 49% annual VMT savings. While the median number of accounts per route increased from 165 to 231, the median accounts per VMT increase by 2.7 times, owing to the increased geographic efficiency of a zoned system.

Table 12: Summary comparison of Optimized Existing and Clustered Analyses

	Annual VMT (million miles)	Median VMT/Route	Median Accounts/Route	Median Route Time (hours)	Median Route Speed (mph)	Median Accounts/VMT
Optimized Existing	6.1	46	165	4.3	11.9	3.2
Clustered	3.1	26	231	3.3	7.6	8.7

The heat map for the Clustered Analysis was subtracted from that of the Optimized Existing Conditions, resulting in a new heat map that shows how the density of VMT changed across the city from one condition to the other. Overall, VMT density was reduced nearly everywhere in the city. As shown in Figure 13, VMT savings were most significant in the Manhattan CBD and along the Gowanus Expressway. VMT density increased slightly along the LIE, the western portion of Jamaica Bay, and along the routes in New Jersey and Nassau County that led to transfer stations. These increases are a direct result of the way garages and transfer stations were assigned to the zones and do not represent a significant increase in VMT (maximum of 50 VMT/square mile).

Figure 13: Change in VMT density between the Optimized Existing and Clustered Routes

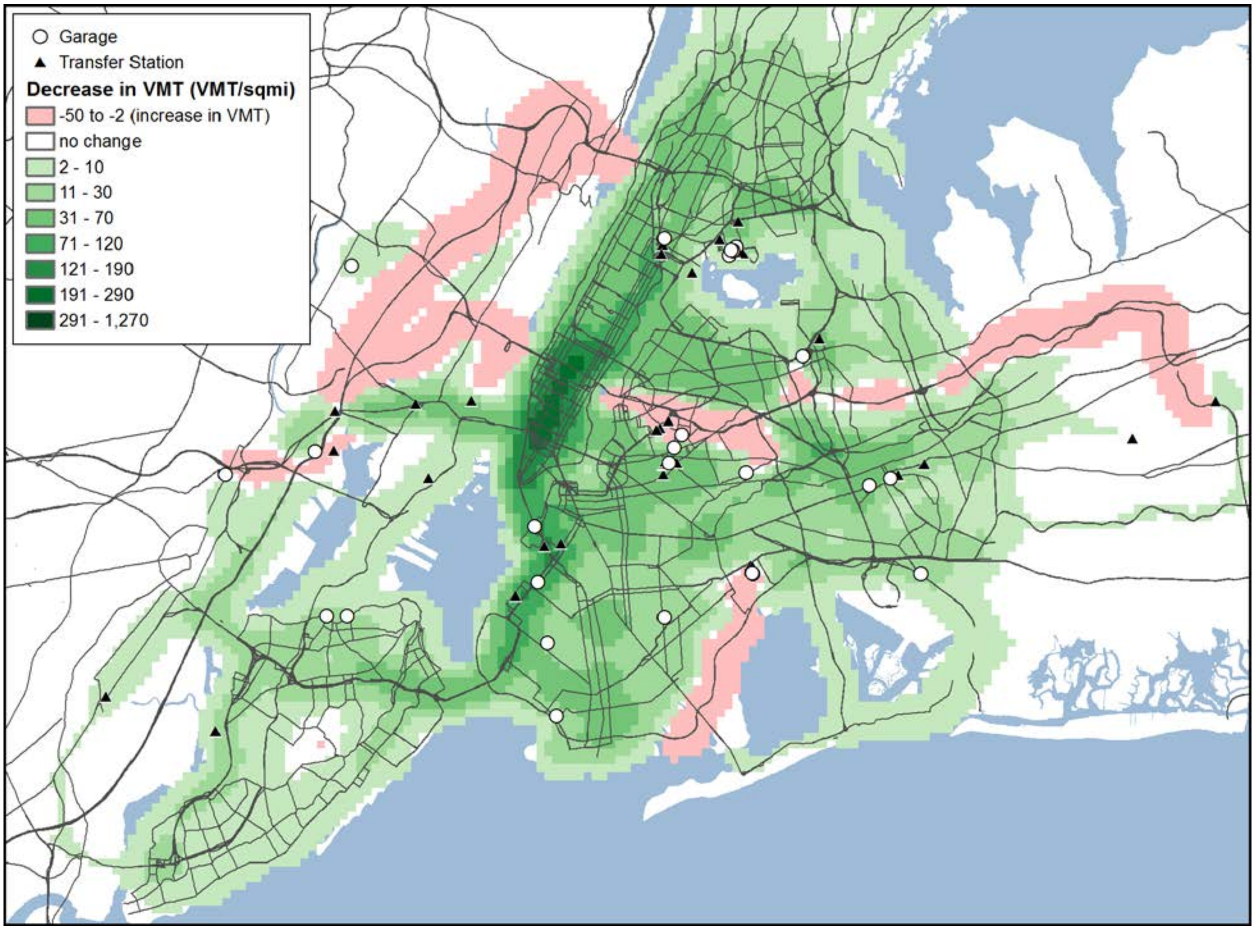

Key Findings

The private carting industry in New York City is highly inefficient from a vehicle miles traveled perspective. Using data reported by the carters through BIC, this analysis mapped the truck routes of the twenty largest carters in the city (as ranked by the total number of accounts served) to create a baseline to estimate the citywide VMT of the private carting industry. It was found that the private carting industry currently contributes approximately 23.1 million VMT annually to the streets of New York City and the adjacent counties, where many garages and transfer stations are located. It was found that many existing routes are geographically dispersed, often serving several neighborhoods across multiple boroughs; that routes from the same and different carters often overlap along key routes and neighborhood streets, creating duplicative services in nearly every corner of the city; and that for many routes, garages and transfer stations are far from the core service area of the route.

A hypothetical Zoned System was created to understand how VMT may change from the current system. Additional routing analyses were conducted in addition to the Existing Conditions, which were then used to estimate the VMT impacts of a Zoned System. The Existing Conditions Analysis was compared with the Random Analysis and the Optimized Existing Conditions Analysis was compared with the Clustered Analysis. In the Optimized/Clustered Analysis, where most of the factors in route creation were eliminated, the geographic constraints of the Zoned System reduced VMT by nearly 50%. This is a relatively conservative approach as compared with the Existing/Random Analysis, which showed a nearly 70% reduction in VMT. Simply put, a Zoned System's routes could be expected to be much more efficient than the existing system due to the geographic constraints created by exclusive zones. Overall, a Zoned System could reduce the VMT of the private carting industry in the range of 50%-70%.

Under a Zoned System, every neighborhood in the city would likely see a decrease in truck traffic, VMT, and the related emissions, however the areas with the highest concentrations of VMT will see the greatest benefits. These areas are in many cases the neighborhoods of greatest concern from an environmental justice standpoint (i.e. those with lower-income, disadvantage populations with worse health outcomes). The Manhattan CBD, which has the highest concentration of accounts in the city, would also see dramatic benefits.

There are many factors that could influence the VMT outcomes of an actual Zoned System being implemented in New York City, including the geography of the zones, how many zones there are, the policies and rules that shape the future private carting environment, who is awarded any future contracts, and how the balance of garages and transfer stations shifts, among many others. What is clear from this analysis is that the current system, from a VMT and geographic efficiency standpoint, is highly inefficient and duplicative in terms of the overlap of existing carter ranges. Furthermore, it is clear that implementing an exclusive Zoned System would dramatically decrease the VMT (and therefore the emissions) of the private carting industry. A Zoned System, by the nature of its geographic constraints, would create a more efficient private carting collection system in New York City. The benefits of a Zoned System would be felt citywide, and in particularly in the neighborhoods that are most disproportionately impacted by the current system.

Appendix

Individual Carter Data Challenges and Solutions

Carter 1 provided their data in the standard format close to the format of the data entry template provided by BIC. Carter 1 provided one workbook for each of the four weeks of data. Each workbook had two sheets, one for material information, i.e. material type, garage location, and transfer station location. A second sheet contained customer information with the first column displaying which truck serves the route. The data was filtered to only include the desired day, 9/10/14. The standard import script was then used to bring this data into an ArcGIS geodatabase for further analysis.

Carter 2 supplied data in nearly the same format as Carter 1. The same process, described above, was used to import Carter 2's dataset. The one difference was that Carter 2 combined all study period data into one sheet.

Carter 3 provided its data in the same 'Material + Customer Info' format described above. However, the data was separated into over 470 workbooks, one for each route for each day in the four weeks for which data was collected. After manually separating the sheets containing data for 9/10/14 the 'merge_excel' script was used to combine this data into one sheet prepared for import. The prepared data was then imported into a geodatabase using the standard import script. The Carter 3 data was also missing adequate transfer station location info for four routes, simply referring to the location as 'New Jersey'. BIC confirmed with the carter that the location was the Covanta Essex transfer station.

Carter 4 provided their data in nearly the same format as Carter 3. Carter 4 did not, however, have missing information related to garage or transfer stations. The same process was used to import Carter 4 Carting Corporation data as for Five Star Carting. The Carter 4 data also needed transfer stations to be added mid-route for routes with multiple tips as it was one of the companies with non-sequenced multiple load routes.

Carter 5 provided their data in a single workbook in the 'Material + Customer Info' format. The workbook included a tab for material and customer info for each of the four weeks included in the study. It was necessary to manually remove data for non-September weeks before using the data import script. Carter 5 used day names for its route names (i.e. WED), however, these days were off by one day when compared to the date column in each tab. For example, a Wednesday route was listed as being on 9/9/2014, which was a Tuesday. BIC confirmed with the carter that the route name was correct and the dates were adjusted accordingly.

Carter 6 provided data in a slightly different format with one workbook for each route. Each workbook had a separate tab for customer data for each day within the survey period with one material information tab. Non-9/10/14 tabs were removed manually. The merge_excel and GIS import script could then be used to import the data.

Carter 7 supplied one workbook for each week within the study period. These workbooks contained one tab for each day of the week. Each non-9/10/14 sheet was removed manually. Instead of having a separate tab for material information, the material information including garage and transfer station location was tagged to the first stop of each route. The ArcGIS excel to table function was used to bring this data into the project geodatabase.

Carter 8 provided customer sequence information for each route for each day in the four week study period. One material information lookup table was provided to connect the route data with the correct garage and transfer locations. First the routes in service on 9/10/14 were separated and combined using the merge_excel script. The material information sheet was then added to this merged workbook. That data was then ready to be added to the ArcGIS database using the standard import script. Some data clarifications from the carter were necessary as many of the route names found in the customer sequence data did not match the names used in the material information look up table.

Carter 9 provided data in the standard 'Material + Customer Info' format with one workbook containing all data. The carter also provided an alternative format that was not used. First the data was filtered to only the desired day, 9/10/14. It could then be imported into ArcGIS using the standard import script.

Carter 10 provided data in the standard 'Material + Customer Info' format with one workbook for each route for each day in the study period. This data was first combined using the merge_excel script. It was then imported to GIS using the standard import script.

Carter 11 provided data in the standard 'Material + Customer Info' format with one workbook for each week in the study period. The first step involved filtering by day and removing non-9/10/14 data from the workbook. This workbook was ready to be imported using the standard import script, however, while the route names did match between customer info sheet and the material information sheet they did not match exactly. This issue was fixed by manually editing the material route names to exactly match their equivalents in the customer sheet so that they could be matched accurately.

Carter 12 provided data in the same format as Carter 1 and Carter 2. The scripts were originally designed for this format so little to no data cleaning was necessary before importing into ArcGIS.

Carter 13 provided data in the standard 'Material + Customer Info' format with one workbook for each route for each day in the study period. After filtering to the desired day, this data was combined using the merge_excel script. It was then imported to GIS using the standard import script. When the Carter 13 data was transmitted it seemed to be missing data for some of the routes, however these routes (truck 41, 25, and 37) turned out to be extra 'routes' for overflow and can be used for any customers that require extra service.

Carter 14 provided data in the standard 'Material + Customer Info' format with one workbook containing all data. First the data was filtered to only the desired day, 9/10/14. This workbook was then ready to be imported using the standard import script, however, while the route names did match between customer info sheet and the material information sheet, a successful match could not be completed. This issue was fixed by manually editing the material route names to exactly match their equivalents in the customer sheet so that they could be matched accurately.

Carter 15 provided data in the standard 'Material + Customer Info' format with one workbook for each route for each day in the study period. After selecting the workbooks from 9/10/14, this data was first combined using the merge_excel script. It was then imported to GIS using the standard import script.

Carter 16 provided data in the standard 'Material + Customer Info' format with one workbook for each route for each day in the study period. After selecting the workbooks from 9/10/14, this data was first combined using the merge_excel script. The data originally provided by Carter 16 did not have matching route names between the customer and material sheets. After contacting the carter through BIC and getting updated route names it was possible to import the data through the standard GIS import script.

Carter 17 provided data in the standard 'Material + Customer Info' format with one workbook for each route for each day in the study period. After selecting the workbooks from 9/10/14, this data was first combined using the merge_excel script. It was then imported to GIS using the standard import script.

Carter 18 provided customer sequence information for each route for each day in the four week study period. One material information lookup table was provided to connect the route data with the correct garage and transfer locations. First the routes in service on 9/10/14 were separated and combined using the merge_excel script. The material information sheet was then added to this merged workbook. The data was then added to the ArcGIS database using the standard import script.

Carter 19 provided data in the standard 'Material + Customer Info' format with one workbook containing all data. First the data was filtered to only the desired day, 9/10/14. After filtering, significant manual data cleaning was necessary as there were arbitrary spaces and duplicate entries throughout the customer information sheet. Route names also did not match exactly between the customer and material information sheets so manual corrections were necessary. After sufficient cleaning the standard import script was used to bring the data into the ArcGIS geodatabase.

Carter 20 provided data in the standard ‘Material + Customer Info’ format with one workbook for each route for each day in the study period. After filtering to the desired day, this data was combined using the merge_excel script. It was then imported to GIS using the standard import script.