

Technical Appendix D

Locational Data for TRI Reporting Facilities and Off-site Facilities

Table of Contents

1.	Introduction	D-1
2.	Results of Geocoding	D-2
3.	TRI Reporting Facilities	D-3
3.1	Inputs	D-3
3.2	Overview of Process	D-4
3.3	Details of Process	D-5
3.3.1	Extract the high-quality preferred lat/longs from pref94D and the LRT, and preserve them for the final dataset	D-5
3.3.2	Combine the old and new geocoded, add the LRT data, and select best gdt coordinates	D-6
3.3.3	Select best TRI coordinates	D-8
3.3.4	Calculate the distance between the best TRI and the best geocoded coordinates	D-8
3.3.5	In a series of steps, replace the low-quality preferred lat/longs with the geocoded lat/longs	D-9
3.3.6	Replace the reported county name with the geocoded name for the county in any instances where the county field is blank	D-10
3.3.7	Facilities with missing coordinates and quality assurance	D-10
3.3.8	Creating the Final Table	D-11
3.4	Method, Accuracy, and Description (MAD) Codes	D-14
4.	Off-site Facilities	D-24
4.1	Overview of Process	D-24
4.2	Collapsing Reported Off-site Facilities	D-24
4.3	Geocoding the Off-site Facilities	D-25
4.3.1	Collapsing off-site facilities after geocoding	D-25
4.3.2	Cleaning and conditioning the data	D-26
4.3.3	Identifying a set of best names and addresses	D-28
4.3.4	Matching records within the set of best names and addresses	D-30
4.3.5	Finding indirect matches	D-30
4.3.6	Identifying and assigning the best state match	D-31
4.3.7	Results	D-32

1. Introduction

The RSEI model uses latitude and longitude coordinates for each TRI reporting and off-site facility to locate each facility on the grid that underpins the model. The facility's location determines many of the modeling inputs, including the exposed population. With changes made to the model for Version 2.1, including more detailed air modeling close to the facility, and full-model results for surface water media, accurate locational data takes on additional importance.

There are two types of facilities included in the model, TRI reporting facilities and off-site facilities. The quality of locational data varies significantly between the two types. TRI reporters submit their own addresses and estimates of their latitude and longitude (lat/long) on Form R every year when they submit their release reports. These reports are subject to common reporting errors: transposition of digits, confusion of latitude with longitude, lack of precision, and nonreporting. The quality of reported data for off-site facilities is much worse, as the name and address of these off-site facilities are reported by the TRI reporters transferring the waste, not the receiving facility itself. The name and address tend to be reported in slightly different ways by different reporters, and often misspelled or misreported. Latitude and longitude are not reported at all. Little standardization is performed by the TRI program, therefore minor differences in an off-site facility record, such as a slight misspelling of the name, or "St." instead of "Street", can make two records look like two different facilities, when they are really the same.

In RSEI Version 1.x, reporting facilities were located on the grid using their reported latitudes and longitudes, and off-site facilities were located using the coordinates of the centroid of their 5-digit ZIP code. For Version 2.1, the lat/longs for both reporters and off-site facilities were improved using a commercial geocoding service. Geocoding is a process where a computer program uses street address, city, state, and ZIP code to match addresses to geographic points in Census TIGER files, and then determines the latitude and longitude of the address.

For reporting facilities, existing EPA data were also used, including the results of a one-time quality-assurance process performed in 1996 on the coordinates reported by facilities, as well as EPA's current centralized storehouse of locational data, the Locational Reference Table (LRT). These EPA data, the geocoded data, and the submitted and preferred coordinates as reported in TRI are all compared and the best set of coordinates is chosen through a series of tests, as described in the following sections.

For off-site facilities, the chief hurdle was to identify all of the different ways each true unique off-site facility could be reported, and then use the most accurate geocoded results for that unique facility.

Section 2 describes the geocoding process. Section 3 presents the geocoding results for TRI reporting facilities, and describes how the highest-quality final coordinates were determined. Section 4 presents the geocoding results for off-site facilities, and describes how the set of all reported off-site facilities were collapsed into a set of unique facilities with best coordinates.

2. Results of Geocoding

A commercial firm, Thomas Computing Services (TCS), was selected to perform the geocoding process. TCS geocodes data using the GDT software package Matchmaker 2000. This geocoding process involved matching records in the address databases to a reference street map. The reference street map with positioning information is based on the U.S. government TIGER census files.¹ Matchmaker links records in the two databases by matching street names and addresses. When the database records are successfully matched to a reference street map database, the record is considered a match and tagged with the correct latitude and longitude coordinates from the reference street map. After geocoding, some nonmatched records are matched manually, using Internet resources, other databases, and direct contact with the facility.

The matches are broken down into the following different types:

1. *Street segment exact match*- address is matched to a specific segment of a street, including matches that were made manually;
2. *ZIP+4 centroid match*- street address cannot be matched, and facility is assigned to the coordinates of the centroid of its reported nine-digit ZIP code;
3. *ZIP+2 centroid match*- street address cannot be matched, and facility is assigned to the coordinates of the centroid of its reported seven-digit ZIP code;
4. *ZIP code centroid match*- street address cannot be matched, and facility is assigned to the coordinates of the centroid of its reported five-digit ZIP code; and
5. *No match*- none of the above matches is detected.

TCS separately geocoded the TRI reporters and the off-site facilities. The results of each process are discussed in Sections 3 and 4, respectively.

¹ The data are based on public record and cannot be copyrighted, therefore the files do not have licensing constraints.

3. TRI Reporting Facilities

The database of all TRI reporting facilities for 1988-2002 includes 49,369 facilities. As described above, these facilities were geocoded by TCS. Some additional manual changes that were performed the previous year by TCS (map look-ups, etc) were carried over to the current year. The results are shown below in Table D-1.

Table D-1
Match Results for TRI Reporting Facilities

Coordinate Matches	Number of Records	Match Percentage
Street address (includes manual matches)	34,926	70.74%
ZIP+4 centroids	99	0.20%
ZIP+2 centroids	369	0.75%
ZIP code centroids	13,579	27.51%
<i>Total matches</i>	<i>48,973</i>	<i>99.20%</i>
<i>Unmatched</i>	<i>396</i>	<i>0.80%</i>

These results were combined with other data to determine the most accurate set of coordinates for each facility. The process and results are described below.

3.1 Inputs

Five tables were used in this analysis to determine the best set of coordinates:

- **‘Facility’**. This table contains the data on TRI reporting facilities from the March 2004 Public Data Release TRI data freeze, including the 49,369 facilities currently or historically reporting to TRI, up to and including TRI Reporting Year 2002. Fields relevant for this analysis are TRI ID (field name ‘FacilityID’), county, state, submitted latitude and longitude (as submitted by the reporting facility), and preferred latitude and longitude (the facility-submitted coordinates after some annual quality assurance performed by EPA).
- **‘Pref94D’**. This table was provided by Loren Hall of U.S. EPA in dbf format. It contains the results of a QA process done in 1996 on TRI reporting data up to and

including RY1994. Data includes, for 36,652 facilities, submitted lat/longs, preferred lat/longs, codes describing the level of accuracy of the preferred lat/longs (in field PREFER_AC), and the QA checks themselves that were done in order to determine the quality of the lat/longs that were considered “preferred” (PREFER_QA). There is a preferred lat/long for each facility, and some of them failed very basic tests, so “preferred” should simply be taken as a sign that a set of coordinates went through tests, not that they are necessarily of a high quality. The quality of the preferred lat/longs can only be determined by looking at the fields PREFER_QA and PREFER_AC.

- **‘Old_gdt’.** This table is based on a Lotus 1-2-3 worksheet, ‘trigdt’, also provided by Loren Hall. This file was generated by OIRM’s (Office of Information Resources Management) System Development Center in 1998, and was a geocoding effort of all regulated entities known to EPA (some 25 million addresses). In this exercise, duplicate facility records were not eliminated, since the Agency did not want to miss any changed addresses. This geocoded and QC’d data became the source of most of the data in the Envirofacts database. The file contained 17,286 records (after records with zero lat/longs were deleted), with geocoded results (lat/long, confidence level information) for each one. In this Appendix, this dataset will be referred to as the old gdt dataset.
- **‘New_gdt’.** This table is based on data provided by TCS, the private company contracted to geocode reporting and off-site facilities. The file contains all of the current and historically reporting facilities that were geocoded to differing levels of precision, from street segment address matches to 5-digit ZIP code matches. Because the data set that TCS geocoded was simply a later version (with some extra new reporters) of the data that was geocoded in 1998, the results were compared and found to be very similar. This is because TCS uses the same gdt software used by OIRM. In this Appendix, this dataset will be referred to as the new gdt dataset.
- **‘LRT.’** EPA’s Agency-wide repository of locational information is called the Locational Reference Tables (LRT). The LRT is a result of an EPA effort to obtain latitude/longitude coordinate information of documented origin for all of EPA's regulated facilities and sites. The LRT act as a storehouse for the actual locational data as well as the business rules that are applied to them in order to provide the most accurate information available for depicting the locations of federally regulated entities. LRT data was provided by Timothy Antisdell of U.S. EPA in a custom data extraction for all of the TRI reporting facilities. The data fields provided included latitude and longitude, method, accuracy and description codes, and codes describing the program origin and the conveyor of the data. For more information on the LRT, see the EPA documentation on the internet at <http://www.epa.gov/enviro/html/locational/>.

3.2 Overview of Process

There are several steps involved in determining a best set of coordinates:

1. Extract the high-quality preferred lat/longs from PREF94D and the LRT, and preserve them for the final data set;
2. Select a best set of gdt coordinates for each facility from the old gdt set, the new gdt set, and the address-match quality LRT data, deferring to the new_gdt;
3. Select a best set of TRI coordinates for each facility from the reported and preferred sets from 'Facility' and 'PREF94D';
4. Compare the best TRI coordinates to the best gdt coordinates, calculating the distance between the two sets of coordinates for each facility; and
5. In a series of steps, replace the TRI lat/longs with the gdt lat/longs if a) the distance between the two is greater than a determined minimum, and b) the confidence level of the gdt lat/longs is above a certain minimum.

3.3 Details of Process

This analysis was conducted in Access. Each step of the process is described in the sections below.

3.3.1 Extract the high-quality preferred lat/longs from pref94D and the LRT, and preserve them for the final dataset

In order to preserve the high-quality lat/longs from the round of quality checks performed in 1996, those considered 'high quality' were extracted from the table 'pref94D' and set aside. The following criteria for 'high quality' were developed in consultation with Loren Hall:

1. Facilities where the third position of PREFER_QA field is "1" (indicating that the submitted coordinates were found to be within 2 km of a reasonably good alternate coordinate value); or
2. Facilities where the fourth position of PREFER_QA field is "V", "A", or "D" (manually verified to have preferred coordinates); or
3. Facilities where the PREFER_AC field value is <150 meters (the coordinates are considered accurate to within 150 meters, based on the kind of check performed).

A set of preferred lat/longs meeting any one of these conditions was considered high quality and set aside. There were 18,036 facilities originally in this group. Sixty-three records with zero values in the PREFER_AC (which describes in meters the level of accuracy of the preferred lat/longs) were deleted from this set, and therefore went through the rest of the process like other

non-high quality lat/longs. Several other sets of coordinates have failed plotting in previous years, leaving 17,916 facilities designated as ‘High Quality (HQ Prefer).’

The next step is to extract the similarly high quality data from the LRT Table. A query was done selecting the records where the degree of accuracy was less than 150m OR one of the verification codes was equal to 015 (EPA Verification). There were 7,467 records matching one of these criteria. These records were then matched against the HQ Prefer table- 3,053 records were present in both tables. LRT coordinates were chosen over the HQ Prefer coordinates if one of the following conditions were met: 1) it was manually verified by EPA, 2) the accuracy was the same or more accurate, or 3) the collection method for HQ Prefer was UN (Unknown). In all cases one of these conditions was true, so LRT data were adopted for all 3,053 facilities in common. The table of LRT HQ data was renamed “All HQ” (since the LRT data were adopted in all cases of duplicates) and the HQ Prefer records that were not duplicative were added to that table. In all, there were 22,332 facilities with High-Quality coordinates- 14,865 sourced to HQ Prefer, and 7,467 to LRT.

The remaining 31,452 facilities without high-quality TRI coordinates then were put through the comparison with the geocoded facilities described in the next section.

3.3.2 Combine the old and new geocoded, add the LRT data, and select best gdt coordinates

There is a large degree of overlap between the new and old gdt data sets, with approximately 17,000 facilities in common. The coordinates of the duplicates facilities were compared. If there was no difference, then the new gdt fields were adopted as the more current source. If there was a difference, then the coordinates were compared based on the confidence level of the geocoding. The values in the following fields were compared:

Table D-2
Match Level Codes
(In descending order of quality)

Field 'XIN'	Field 'STAT'	Type of Match
0, S, V or I	B1, R1, B2, R2, B3, R3, R4, B5, R5,	matched to a street segment
4		matched to a 4-digit ZIP code centroid
0	B6, R6, B7, R7	matched to a placeholder
2		matched to a 2-digit ZIP code centroid
X		matched to a 5-digit ZIP code centroid

An XIN = 0, STAT = B1, R1 match is the most accurate, and a XIN = 5 is the least accurate. The most accurate match was chosen for each facility. In the cases where the confidence level was the same, the new gdt coordinates were chosen as the most recent source.

The duplicate facilities described above were combined with the unique facilities from each data set into one table containing one record for each geocoded facility with the best set of geocoded coordinates.

Some of the coordinates contained in the LRT set are based on address geocoding, and could potentially be of higher quality than the gdt records. A query was performed to pull out this data, using all of the following conditions: 1) collection method = 1 (Address match); 2) Accuracy >=150 and <=200; and 3) Conveyor is not TRIS-PREFERRED or TRIS-REPORTED. The last condition is because in the next step we compare the best geocoded coordinate to the best TRI coordinate- we want to avoid comparing TRI to TRI. There were 28,462 facilities that met all of these conditions.

The geocoded LRT data were then matched against the 'All gdt' records. The existing coordinates were replaced if the LRT data were of greater accuracy, or if the source was the old gdt set. There were 4,085 facilities whose gdt coordinates were replaced with LRT data. NOTE: Old_gdt is still included as a source, as there are 234 facilities this year that still use old_gdt and there are no good-quality LRT coordinates for the same facilities.

The resulting data set will still be referred to at 'All gdt' despite the inclusion of LRT data.

3.3.3 Select best TRI coordinates

The facilities without high-quality TRI coordinates were matched against the facilities in the combined gdt dataset. For each facility in this set, it was necessary to select a lat/long to compare with the geocoded lat/long. There are two possibilities in each dataset (the Pref94D database or the current set of TRI reporters called 'Facility'): the Preferred lat/long, or the Submitted lat/long. Following further consultation with Loren Hall, the decision was made to compare the geocoded results with the Submitted rather than the Preferred coordinates. This comparison was considered more appropriate because a number of instances were identified where one of the QA tests used to derive Preferred values erroneously rejected valid submitted coordinates. Therefore, using the Submitted coordinates will avoid perpetuating a situation where EPA rejected a valid submitted coordinate. Loren Hall also advised that it was better to use the submitted coordinates from the 1996 data set (Pref94D), because in some instances erroneous Preferred coordinates had been entered into facilities' Form R's before they were sent out to them (to simplify reporting for the facilities). In many cases where these facilities' Preferred coordinates were wrong, the error was probably not discovered. This may have resulted in the erroneous Preferred coordinates becoming erroneous Submitted coordinates, thereby perpetuating the error.

Using the logic described above, each facility was assigned its 'best' set of TRI coordinates (TRI lat/long), using the following hierarchy:

- 1994 Submitted (from 'Pref94D');
- 2002 Submitted (from 'Facility');
- 1994 Preferred (from 'Pref94D');
- 2002 Preferred (from 'Facility').

3.3.4 Calculate the distance between the best TRI and the best geocoded coordinates

The distance (in kilometers) from each facility's best TRI coordinate to its best gdt (geocoded) coordinate was then calculated using the following formula:

$$\text{Distance} = 6377 * \text{acos}(\text{cos}(\text{rad}(90 - \text{TRI Lat})) * \text{cos}(\text{rad}(90 - \text{gdt Lat})) + \text{sin}(\text{rad}(90 - \text{TRI Lat})) * \text{sin}(\text{rad}(90 - \text{gdt Lat})) * \text{cos}(\text{rad}(\text{TRI Long} * -1) - \text{gdt Long}))$$

The resulting value was then used in determining whether to retain the TRI coordinates or substitute the gdt coordinates, as described below.

3.3.5 In a series of steps, replace the low-quality preferred lat/longs with the geocoded lat/longs

The basic premise of the following steps is that one can have greater confidence replacing low-quality lat/longs with geocoded lat/longs if the confidence level associated with the geocoded lat/longs is very high, and the distance between the two sets of coordinates is very great. In these cases one can feel confident that the submitted lat/long is simply erroneous. As the confidence level of the geocoded lat/longs decreases toward the level of a ZIP code centroid, one cannot be sure that differences of a few kilometers do not simply represent real distances from a plant to the centroid of its ZIP code. Therefore, the replacement of the low-quality preferred lat/longs was done in a series of steps that accounted for both the distance between the sets of coordinates and the confidence level of the geocoded results.

At each step, the geocoded coordinates that matched the criteria were substituted for the TRI coordinates and set aside. For the field 'Final Source,' the following codes were used:

- QA_GDT. This code refers to when geocoded lat/longs from the combined gdt database were substituted for the TRI lat/longs.
- QA_TRI. This code refers to when the pairs of lat/longs did not meet any of the criteria above, so the TRI Form R-reported lat/longs were retained.

Step 1. Low-quality lat/longs were replaced with geocoded coordinates if the distance between the two sets of coordinates was greater than or equal to 2 km, AND the geocoded result was matched at a street segment or intersection— i.e., the GSTAT field showed B1, B2, B3, B5, R1, R2, R3, or R4 and the GDTXIN field showed 0 or S, V or I. In this step, 8,027 facilities were assigned the code 'QA_GDT,' and 9,314 facilities were assigned the code 'QA_TRI'.

Step 2. Low-quality lat/longs were NOT replaced with geocoded coordinates if the distance between the two sets of coordinates was greater than or equal to 2 km and less than 5 km, AND the geocoded result was matched at a street segment or intersection— i.e., the GSTAT field showed B1, B2, B3, B5, R1, R2, R3, or R4, AND the facility reported MORE THAN 1,000,000 pounds of total releases, including direct releases and off-site transfers. TRI release data from 2001 was used for this test, and 843 facilities in total were found to have more than 1 million pounds of total releases. In these cases we are assuming that the large plants know their locations well and may have a good reason to report addresses up to 5 km different from their lat/longs; for instance, the geocoded result may represent the 'front door' of the facility, but the submitted lat/long represents either the point of release or the center of production. In Step 2, we are assuming that facilities releasing less than 1,000,000 pounds in 2001 are not large enough to have such an issue, and that the difference in the coordinates represents a lack of precision on their part. In this step, 22 facilities were assigned 'QA_TRI'.

Step 3. Low-quality lat/longs were replaced with geocoded coordinates if the distance between the two sets of coordinates was greater than or equal to 10 km, AND the geocoded result was

matched at the ZIP+4 centroid level. In this step, 32 facilities were assigned to 'QA_GDT' and 158 were assigned 'QA_TRI'.

Step 4. Low-quality lat/longs were replaced with geocoded coordinates if 1) the distance between the two sets of coordinates was greater than or equal to 15 km, AND 2) the geocoded result was matched at either the ZIP+2 centroid level OR the B6/R6 placeholder match. In this step, 112 facilities were assigned to 'QA_GDT' and 699 were assigned 'QA_TRI'.

Step 5. Low-quality lat/longs were replaced with geocoded coordinates if the distance between the two sets of coordinates was greater than or equal to 20 km, AND the geocoded result was matched at the 5-digit ZIP code centroid level. In this step, 1,203 facilities were assigned to 'QA_GDT' and 7,346 were assigned 'QA_TRI'.

3.3.6 Replace the reported county name with the geocoded name for the county in any instances where the county field is blank

After the final table was created (see below), county names and FIPS codes were pulled in from the 'Facility' table. The county field was checked for blanks, but no blanks were found in RY 2002.

3.3.7 Facilities with missing coordinates and quality assurance

There were 115 facilities which had geocoded results, but did not have valid (nonzero) TRI coordinates for comparison. These facilities were assigned the source code, 'GDT_NOTRI,' and the coordinates were taken from the new gdt table.

There were 37 facilities in the set of facilities without high-quality TRI coordinates that were not found in the combined set of gdt facilities, for the most part because the facilities were located in places other than the fifty U.S. States and the District of Columbia. These facilities were assigned the source code, 'TRI_NOGDT,' and the coordinates were taken from TRI in the hierarchy described earlier. Eight facilities did not have any coordinates in any of the current source tables, and so zip code centroids were used (taken from an internet zip code lookup table).

Once all of the facilities had been assigned coordinates and assembled in a draft table, the final coordinates were plotted in a GIS (Geographic Information System) program. The state that the coordinates were plotted to were matched against the reported state and visually inspected for cases that did not match. Those that fell on a coast or a river state boundary were considered allowable, and those on straight state boundaries were given a one-mile tolerance before being counted as incorrect. The final result was 24 facilities that plotted outside of their reported state. Using the TRI hierarchy and the gdt coordinates, if available, additional coordinates were checked for each of these failed facilities. If no valid coordinates could be found in the TRI or

gdt data, EPA’s LRT system was checked. If that also failed, the facility’s reported zip code centroid was adopted, using an internet-based zip code lookup.

Additionally, the coordinates for three facilities that had been previously checked using geocoding and maps were provided by Loren Hall of U.S. EPA. These facilities were also assigned the code ‘MANUAL,’ and changed by hand in the Draft final table. These three facilities are listed in Table D-3.

**Table D-3.
Facilities with Coordinates Corrected After Mapping**

Facility Id	Original Lat	Original Long	Original Source	Final_lat	Final_long	Final_source code	Data Source
46517LRMDW58288				41.645432	-85.991839	MANUAL	EPA (Provided by Loren Hall, map look-up.)
46517LRNC 28858				41.648419	-86.019967	MANUAL	EPA (Provided by Loren Hall, map look-up.)
46517TTFRM28816				41.648398	-86.019364	MANUAL	EPA (Provided by Loren Hall, map look-up.)

3.3.8 Creating the Final Table

The final coordinates from the comparison process described in section 3.3.5 were combined with the original High-Quality coordinates that were set aside in section 3.3.2.

Tables D-4 and D-5 show the data fields that will be added to the ‘Facility’ table used in the RSEI model, and how the contents of each field were derived. These tables do not include the Method, Accuracy, and Description (MAD) codes, which are described separately in Section 3.4, below.

**Table D-4
New Data Fields in Final Table**

Field	Derivation of Contents
Facility ID	Facility ID used in TRI reporting, unique for each facility.
SubmitLat SubmitLong	Coordinates originally submitted by each reporting facility as reported in the current year TRI data freeze. NOTE: This field may not match the submitted coordinates used to map the facility in the RSEI model, as the submitted coordinates are taken preferentially from the internal EPA dataset 'Pref94D.'
PreferredLat PreferredLong	QA'd coordinates derived by EPA for each reporting facility as listed in the current year TRI data freeze. NOTE: This field may not match the preferred coordinates used to map the facility in the RSEI model, as the submitted coordinates are taken preferentially from the internal EPA dataset 'Pref94D.'
Latitude Longitude	These are the final coordinates that will be used in the RSEI model. Their derivation depends on what is in the 'FINAL_SOURCE' field (see below).
LatLongSource	<p>QA_GDT. Ultimate source is geocoded data using GDT software, performed either by EPA in 1998 or by TCS in the current year for EPA, or the address-matched records from the LRT. Compared to low-quality TRI lat/longs and determined to be more accurate (see Section 4.4)</p> <p>QA_TRI. Ultimate source is TRI, as reported either in Pref94D or in the 'Facility' table in the current year TRI data freeze. Compared to geocoded coordinates and determined to be more accurate (see Section 4.4).</p> <p>HQ PREFER. Ultimate source is 'PREF_LAT' and 'PREF_LONG' in Pref94D. Originally selected as High Quality Preferred Lat/longs; no comparison to geocoded results was performed.</p> <p>HQ LRT. Ultimate source is the high-quality LRT coordinates. Originally selected as High Quality Lat/longs; no comparison to geocoded results was performed.</p> <p>TRI_NOGDT. Ultimate source is TRI, as reported either in Pref94D or in the 'Facility' table in the current year TRI data freeze. One of approximately 400 facilities with TRI coordinates that did not have geocoding results to compare against.</p> <p>GDT_NOTRI. Ultimate source is geocoded data using GDT software, either by EPA in 1998 or by TCS in the current year for EPA. Adopted without comparison because no TRI coordinates were available.</p> <p>MANUAL. Checked by hand and corrected, either using a map or by transposing digits or coordinates.</p> <p>ZIP. Zip code centroid used, found on internet look-up table.</p> <p>LRT. LRT coordinates adopted during final QA process- facility's initial coordinates failed plotting.</p>

**Table D-5
Summary of Final Facility Coordinates**

Code Used	Type of Match/Comparison	Source	Description	Number of Facilities
QA_GDT	Street Address Match	GDT	TRI coordinates replaced with Geocoded (including address-matched LRT data)	8003
QA_GDT	Zip+4	GDT	TRI coordinates replaced with Geocoded (including address-matched LRT data)	32
QA_GDT	Zip+2 (or B6/R6)	GDT	TRI coordinates replaced with Geocoded (including address-matched LRT data)	108
QA_GDT	5-digit Zip	GDT	TRI coordinates replaced with Geocoded (including address-matched LRT data)	1200
HQ PREFER	Within 2 km of alternate, manually verified, or accuracy <150m	Pref94D	TRI Preferred coordinates retained	14857
HQ LRT	Accuracy < 150m, or verified by EPA (code 015)	LRT	LRT best coordinates retained	7456
QA_TRI	Step 1 (<2 km from geocoded address match)	Submitted*	TRI Submitted coordinates retained, after check against geocoded coordinates	9313
QA_TRI	Step 2b (2-5 km from geocoded address match, >1,000,000 lbs)	Submitted*	TRI Submitted coordinates retained, after check against geocoded coordinates	22
QA_TRI	Step 3 (<10 km from zip+4 geocoded match)	Submitted*	TRI Submitted coordinates retained, after check against geocoded coordinates	158
QA_TRI	Step 4 (<15 km from zip+2 or B6/R6 geocoded match)	Submitted*	TRI Submitted coordinates retained, after check against geocoded coordinates	700
QA_TRI	Step 5 (<20 km from 5-digit zip code match)	Submitted*	TRI Submitted coordinates retained, after check against geocoded coordinates	7347
QA_TRI	Map plotting		TRI adopted after plot failure	4
TRI_NOGDT		Submitted*	TRI Submitted coordinates retained; no geocoded results available to compare	33
GDT_NOTRI		GDT	Geocoded coordinates used; no TRI coordinates available to compare.	107
ZIP	Coordinates revised after plotting coordinates	From zip code lookup	Zip code centroid from Internet zip code look-up used.	15
MANUAL	Coordinates revised after plotting coordinates		Coordinates either mapped or coordinates adjusted (e.g., decimal place moved).	7
LRT	Coordinates revised after plotting coordinates		Best Value from EPA's online Locational Reference Table (LRT) System used.	7
				49369

* Submitted coordinates were preferentially taken from 'Pref94D'; if not available there, they were taken from the 'Facility' table.

3.4 Method, Accuracy, and Description (MAD) Codes

Method, Accuracy, and Description (MAD) codes are standardized codes that describe how a set of lat/longs were generated, what quality assurance checks were performed on it, and how accurate it is considered to be. The codes allow for comparison of different sets of coordinates that were generated at different times and by different processes. These codes are used by EPA offices, contractors, and by EPA's centralized Locational Reference Tables (LRT).

The 'Facility' table in the RSEI database contains some information on MAD codes. In some cases, for instance when the coordinates from EPA's 1996 QA process were adopted, the MAD codes already assigned were simply translated and carried over (the format of the codes was different in 1996). The MAD codes for facilities with LRT coordinates were carried without translation. In other cases, where coordinates were adopted as a result of the comparisons described above, new MAD codes were assigned and added to the table. However, due to resource and time constraints, in some cases not all of the codes could be filled in. Table D-6 describes each of the MAD codes included in the 'Facility' table. This information is taken from the documentation for the LRT. Detailed information on the codes is available on the Internet. Table D-7 describes how these codes were assigned to the TRI facilities in the model. In addition to the values noted in the table, all geocoded facilities were assigned the following values:

Reference Point = 002

Scale = 100000

Hdatum = 002

Table D-6
Description of MAD Codes in 'Facility' Table

Code	Length	Type	Description	Values
Accuracy	8.2	N	Accuracy of the coordinates (in m)	
Method	3	C	Collection method code for the coordinates	001 Address matching-house number
				002 Address matching-block face
				003 Address matching-street centerline
				004 Address matching-nearest intersection
				005 Address matching-primary name
				007 Address matching-other
				009 Census block/group-1990-centroid
				010 Census block tract-1990-centroid
				012 GPS carrier phase static relative position
				014 GPS code measurements (pseudo range) differential (DGPS)
				015 GPS code measurements (pseudo range) precise positioning service
				016 GPS code (pseudo range) standard position (SA off)
				017 GPS code (pseudo range) standard position (SA on)
018 Interpolation-map				

Table D-6
Description of MAD Codes in 'Facility' Table

Code	Length	Type	Description	Values
				019 Interpolation-photo
				026 ZIPcode-centroid
				027 Unknown
				028 GPS-unspecified
				030 Interpolation-digital map source (TIGER)
				037 Zip+4 centroid
				038 Zip+2 centroid

Table D-6
Description of MAD Codes in 'Facility' Table

Code	Length	Type	Description	Values
PREFER_DC	3	C	Code for the category of the feature reference by the coordinates	001 Unknown
				002 Plant entrance (general)
				003 Other
				004 Plant entrance (personnel)
				005 Plant entrance (freight)
				009 Water release pipe
				015 Loading facility
				016 Loading area centroid
				017 Process unit
				019 Administrative building
				020 Facility centroid
				021 NE corner of land parcel
				022 SE corner of land parcel
				023 NW corner of land parcel
				024 SW corner of land parcel
025 Center of facility				

Table D-6
Description of MAD Codes in 'Facility' Table

Code	Length	Type	Description	Values
Hdatum	3	C	Horizontal datum of the coordinate	001 NAD27
				002 NAD83
				003 WGS84
Scale	9	N	Value representing the geographic scale of the map	
ManualVerify	23	C	Results of manual verification test done during current QA, as described above, and in 1996 by EPA.	23 alphanumeric (six 3 character segments, colon delimited). For values, see below.

Table D-6
Description of MAD Codes in 'Facility' Table

Code	Length	Type	Description	Values			
Zac			Proximity to centroid	where		and	
Cac			Proximity to centroid (county)	a = distance		c = confidence	
Oac			Proximity to centroid (other)	value codes:		value codes	
Aac			Proximity to alternative facility coordinate	0	zero	H	High
Isc			Point in polygon (zip code)	1	500m	M	Medium
Nsc			Point in polygon (county)	2	1km	L	Low
Pac			Point in bounding box (zip code)	3	1-5km	N	Not applic.
Tac			Point in bounding box (county)	4	2km		
Rsc			Verified relative to map features (roads)	5	2-5km	X	Method does not support coord location
Xsc			Verified Relative to map Features (Railroads)	6	5-10km		
Lsc			Verified Relative to Map features (land use)	7	>10km		
Fsc			Verified relative to map Features (photo or satellite image)	9	other scale		
				E	1:24K		
				J	scale 1:100K		

Table D-6
Description of MAD Codes in 'Facility' Table

Code	Length	Type	Description	Values
Verify1 through Verfiy5	3	C	Results of verification tests. Available for data originating in the LRT; also used to describe QA process descibed above (for QA_GDT and QA_TRI, this field is duplicative of ManualVerify).	001 Proximity to polygon centroid (county)
				002 Proximity to alternative facility coordinate
				003 Proximity to polygon centroid (other)
				004 Point in polygon (other)
				005 Point in polygon (county)
				006 Point in polygon (zip)
				007 Verified relative to map features (1:24K)
				008 Verified relative to map (1:100K or TIGER)
				009 Verified relative to map features
				010 Verified, unknown method
				011 Ground truth conducted
				012 Proximity to polygon centroid (zip code)

**Table D-7
Summary of Final Source Codes and MAD Codes Assigned**

Final_Source Code Used	Type of Match/Comparison	Source	Description	MAD CODES
QA_GDT	Street Address Match	GDT	TRI coordinates replaced with Geocoded	Accuracy = 150 Method = 002
QA_GDT	ZIP+4	GDT	TRI coordinates replaced with Geocoded	Accuracy = 4000 Method = 026
QA_GDT	ZIP+2 (or B6/R6)	GDT	TRI coordinates replaced with Geocoded	Accuracy = 8000* Method = 026
QA_GDT	5-digit ZIP	GDT	TRI coordinates replaced with Geocoded	Accuracy = 11000 Method = 026
QA_GDT	Street Match	TCS	TRI coordinates replaced with Geocoded	Accuracy = 150 Method = 002
QA_GDT	ZIP+4	TCS	TRI coordinates replaced with Geocoded	Accuracy = 4000 Method = 026
QA_GDT	ZIP+2 (or B6/R6)	TCS	TRI coordinates replaced with Geocoded	Accuracy = 8000* Method = 026
QA_GDT	5-digit ZIP	TCS	TRI coordinates replaced with Geocoded	Accuracy = 11000 Method = 026

**Table D-7
Summary of Final Source Codes and MAD Codes Assigned**

Final_Source Code Used	Type of Match/Comparison	Source	Description	MAD CODES
HQ_PREFER	Within 2 km of alternate, manually verified, or accuracy within 150 km	Pref94D	TRI Preferred coordinates retained	Maintain existing MAD codes in Pref94D; translated into new format
QA_TRI	Step 1 (<2 km from geocoded address match)	Submitted	TRI Submitted coordinates retained, after check against geocoded coordinates	Manual Verify = A4H Method = 002 Accuracy = 150 Verify 1 = 002
QA_TRI	Step 2b (2-5 km from geocoded address match, >1,000,000 lbs)	Submitted	TRI Submitted coordinates retained, after check against geocoded coordinates	Manual Verify = A5M Method = 002 Accuracy = 150 Verify 1 = 002
QA_TRI	Step 3 (<10 km from ZIP+4 geocoded match)	Submitted	TRI Submitted coordinates retained, after check against geocoded coordinates	Manual Verify = Z6H Method = 026 Accuracy = 4000 Verify 1 = 012
QA_TRI	Step 4 (<15 km from ZIP+2 or B6/R6 geocoded match)	Submitted	TRI Submitted coordinates retained, after check against geocoded coordinates	Manual Verify = Z7M Method = 026 Accuracy = 8000* Verify 1 = 012
QA_TRI	Step 5 (<20 km from 5-digit ZIP code match)	Submitted	TRI Submitted coordinates retained, after check against geocoded coordinates	Manual Verify = Z7L Method = 026 Accuracy = 11000 Verify 1 = 012

**Table D-7
Summary of Final Source Codes and MAD Codes Assigned**

Final_Source Code Used	Type of Match/Comparison	Source	Description	MAD CODES
GDT_NOTRI	None	Submitted	TRI Submitted coordinates retained; no geocoded results available for comparison	NO MAD CODES
GDT_NOTRI	None	GDT	Geocoded coordinates adopted; no TRI coordinates available for comparison	Same MAD Codes as QA_GDT

* No information available accuracy of ZIP+2. Value is the rounded average of 5-digit ZIP and ZIP+4.

4. Off-site Facilities

Previously, all off-site facilities had been located on the model grid using the centroid of the facility's ZIP code. The geocoding effort represents a significant improvement from that methodology. However, the problems with the set of off-site facilities are longstanding and serious: most notably that unique IDs are not used by TRI, and the addresses are not reported by the facilities themselves, but by those facilities that transfer waste to them. Given this, the accuracy of the reported addresses is questionable. In addition, because many different reporting facilities may be transferring their waste to the same facilities, there are many instances of the same facility being reported with many different permutations of name and address. The biggest challenge in this exercise was to collapse the entire set of off-site facilities into a set of unique facilities. Briefly, the entire set of off-site facilities was geocoded by TCS, and then the whole set was run through a series of matching programs in SAS, designed to match facilities to each other, on name first (based on the assumption that a third party is most likely to get a facility's name correct), providing leeway for non-exact matches, and then moving through the rest of the facility's address and determining if it is a plausible match. After this exercise, the set of unique facilities was pared down from approximately 4.4 million to 45,000. Approximately 34% of the 45,000 facilities had an address match more accurate than a ZIP+4 match.

4.1 Overview of Process

There are several data processing steps in determining unique facilities and their coordinates. First, in order to best determine unique facilities, the facility records were collapsed from approximately 4.4 million to almost 330,000 by removing the exact duplicates. Second, in order to expedite and improve the off-site facility locating process, TCS geocoded the data and reported match rates. Finally, the geocoded off-site facility data was further collapsed in order to remove non-exact duplicates and determine truly unique off-site facilities and their addresses.

4.2 Collapsing Reported Off-site Facilities

There are approximately 3 million off-site facility records in TRI. However, many of these facility records actually represent the same facility; they are just reported in slightly different ways by the facilities transferring chemicals to them. In addition, approximately 1 million records are blank or not viable records. In order to make the geocoding process more efficient, it is necessary to first collapse the list of all reported off-site facilities into possible unique facilities. The first collapsing procedure removes all records that are not viable along with all of the records that are exact duplicates. This first stage collapses the off-site facility records from approximately 4.4 million to approximately 330,000.

Further collapsing, using algorithms in SAS to match addresses where the content is the same but the form is different (i.e., St. instead of Street), can bring the count down to approximately fifty thousand. However, the risk with this second collapse is in matching records that are not exactly the same, and also in picking one address form to represent that off-site facility, where

another form might be better for geocoding purposes. Therefore, to decrease potential error in geocoding unique facilities, the 330,000 off-site facility address records were sent to the geocoding service.

4.3 Geocoding the Off-site Facilities

TCS evaluated the 330,000 off-site facility address records. Their geocoding efforts resulted in a 47% street address match; 0.29% ZIP+4 centroid match; 0.23% ZIP+2 centroid match; nearly 49% ZIP code centroid match; and approximately 4% unmatched records. At this point in the process, these numbers may be misleading, since many of the 330,000 facilities are duplicates. Presumably, some portion of the ZIP code matches and unmatched facilities have problematic street addresses that may be “corrected” by accepting the better data of some other record of the same facility.

4.3.1 Collapsing off-site facilities after geocoding

A “fuzzy” matching SAS program (FIND_UNIQUE.SAS) was used to identify additional duplicate records that belong to a single unique facility. The term “fuzzy” refers to logical systems that do not require exact equality of two values in order to classify the two values as equal. In the name matching application, FIND_UNIQUE.SAS assigns two records to the same unique facility even if some identifying fields do not match exactly. This approach accommodates misspelled words and inconsistencies in how a facility might report its identifying information over time. For example, “DuPont,” “Du Pont,” and “E.I. DuDont” might all refer to the same facility. FIND_UNIQUE.SAS identifies a possible match based on similarity rather than exact equality in the name field and then decides whether to match the various spellings by examining the address fields.

Fuzzy matching always introduces the possibility of error. Two records may be matched that do not in fact belong to the same unique facility. Therefore, some discretion was applied in varying the program parameters and performing manual checks to balance two competing outcomes: a greater number of good/high confidence matches versus a greater number of erroneous matches.

There are four major routines of FIND_UNIQUE.SAS:

1. Cleaning and conditioning the data;
2. Identifying a set of best names and addresses;
3. Matching records within the set of best names and addresses;
4. Finding indirect matches (where two records are matched not to each other but to a common third record).

The following sections describe in detail the SAS program and its application.

4.3.2 Cleaning and conditioning the data

The first part of FIND_UNIQUE.SAS corrects common spelling errors and inconsistencies and prepares the data for the matching algorithms. Data cleaning begins with the removal of extraneous characters, regularization of spaces and conversion of all letters to upper case. Then, words that occur frequently but do not aid in matching are deleted. These words include “COMPANY,” “LIMITED,” “POST OFFICE BOX,” “NOT AVAILABLE,” and numerous other words and their associated abbreviations and variations. If such words remain in the match fields, then a name such as “COB CORPORATION, P.O. BOX 2” would appear very similar to “AC CORPORATION, P.O. BOX 10.” The conditioning process converts the two names and addresses to “COB, 2” and “AC, 10,” respectively.

Where relevant words commonly appear in various forms, the conditioning process substitutes a single form. For example, “NAT'L” and “NATIONAL” are both converted to “NATL.” A frequency analysis and visual review of words in the database led to some regularizations of facility names, such as “ADM,” “A.D.M.” and “ARCHER DANIELS MIDLAND,” or “EMPAC,” “EMPACK” and “EMPAK.”

For computational purposes, FIND_UNIQUE.SAS adds a leading blank and a trailing blank to each name and street address.

One example of how the conditioning might change a name field follows. If the reported name of a company is (the misspelling of “environmental” is intentional):

E	N	V	I	R	O	M	E	N	T	A	L		B	A	N	T	E	R		C	O	R	P	.
---	---	---	---	---	---	---	---	---	---	---	---	--	---	---	---	---	---	---	--	---	---	---	---	---

then the cleaned and conditioned version of the name would be:

	E	N	V	I	R		B	A	N	T	E	R	
--	---	---	---	---	---	--	---	---	---	---	---	---	--

The conditioning process concludes by correcting the state field when possible, based on the ZIP code field. FIND_UNIQUE.SAS does not assume that the ZIP code is correct whenever the reported state and ZIP code conflict. However, it does identify certain values in the state field as particularly susceptible to error. These suspect values were identified by checking reported state codes against reported city names and ZIP codes, using the *1996 World Almanac*. The conditioning process uses the state that corresponds to the reported ZIP code when the reported state is particularly susceptible to error or is not a valid state abbreviation. Table D-8 lists states that TRI reporters frequently misreport.

**Table D-8
Suspect State Abbreviations**

Reported State	Possible Actual State	Reported State	Possible Actual State
AR	AZ	MA	ME
AK	AR	MI	MS
AS	AR	MI	MO
CA	GA	MI	MN
IA	IN	MS	MO
IA	ID	NE	NV
II	IL	ON	OH
KT	KY	OP	OH
KU	KY	RH	RI
LA	AL		

Table D-8 lists state codes that were discarded in favor of the state corresponding to the reported ZIP code if and only if:

1. The reported state is listed in the “Reported State” column, and
2. The state corresponding to the reported ZIP code is the state listed in the “Possible Actual State” column of the same row.

For example, if the reported ZIP code is “85607” and the reported state is “AR,” then the program corrects the state to “AZ.” However, if the reported ZIP code does not begin with “85,” then this section of the program makes no change to the state code.

Another section of the conditioning process corrects state codes in certain city name and state code combinations. For example, where the reported city name is “BALTIMORE” and the reported state is “MA,” the SAS program changes the state to “MD.” The program also changes Canadian province codes to “CN.”

4.3.3 Identifying a set of best names and addresses

The purpose of the second part of FIND_UNIQUE.SAS is to reduce the number of records to be matched as quickly as possible. Since the time required to match all records in a dataset to each other increases exponentially as the number of records increases, it is important to perform preliminary matching using a simpler method where possible. FIND_UNIQUE.SAS does this by sorting records by facility name and comparing adjacent records. Thus, this early round of matching compares each record only to the preceding record and finds a match only in cases where the similarity is quite strong.

Specifically, the program sorts the data by the first ten non-blank characters in the facility name. If a reported facility name begins with the same ten characters as in the preceding record, the program compares the street addresses and ZIP codes and assigns three scores that measure the closeness of match in these location fields. If the scores exceed specified thresholds, then the program matches the two records to a single facility. Similarly the program then sorts the data by the first ten non-blank characters in the facility street address and compares the names in adjacent records.

FIND_UNIQUE.SAS calculates three scores that measure how well two names or two street addresses match. The definitions below use two new terms: *source* and *target*. The *source* is the set of words – i.e., name or street address – for which a match is sought. The *target* is the set of words that is being compared to the source. In the current part of the program, which compares adjacent records only, it does not matter which comparison value is designated the target and which is designated the source.

1. **Match Score:** The match score is the weighted proportion of letter pairs in the source also found in the target. A score of 0 means that no letter pairs in the source occur anywhere in the target. A score of 1 means that 100 percent of the letter pairs in the source also occur at least once in the target.

Example:

Source =	<table border="1"><tr><td></td><td>B</td><td>A</td><td>N</td><td>T</td><td>E</td><td>R</td><td></td></tr></table>		B	A	N	T	E	R			
	B	A	N	T	E	R					
Target =	<table border="1"><tr><td></td><td>B</td><td>A</td><td>N</td><td>D</td><td>A</td><td>I</td><td>D</td><td>S</td><td></td></tr></table>		B	A	N	D	A	I	D	S	
	B	A	N	D	A	I	D	S			

The eight letter pairs in the source are: _B, BA, AN, NT, TE, ER and R_, where “_” represents a blank. Of these, _B, BA and AN also appear in the target. Therefore, the unweighted match score is 3/8 or 37.5 percent.

FIND_UNIQUE.SAS introduces variable weights to allow the user to apply expectations about where errors are most important. In the current application, weights for letter pair matches decline exponentially so that matches near the beginning of the target are more valuable than later matches. The use of this model was based on an informal examination of the data.

2. **Position Score:** The position score measures similarity in the sequencing of letter pairs. The reason this is important is that the match score gives credit for a letter pair match regardless of where the letter pair occurs. In the above example, if the target had been “AFTER BANDAIDS,” the match score would have increased to 7/8 or 87.5 percent because the letter pairs TE, ER and R_ occur in “AFTER.”

The position score depends on where a matched pair is with respect to the first matched pair. In the following example, the first pair matched is _B, which occurs in the target at position 7.

Source =		B	A	N	T	E	R	
Position =	1	2	3	4	5	6	7	8

Target =		A	F	T	E	R		B	A	N	D	A	I	D	S	
Position =	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16

The position score is similar to the root mean square (RMS) algorithm commonly used to measure error in diverse situations. A position score of 0 indicates that the matched letter pairs occur exactly in order and at the same relative positions in both the target and source. Higher scores indicate poorer matches.

3. **Leftover Score:** The leftover score measures the percent of the target that is not matched to any letter pairs in the source. The leftover score helps compensate for the tendency of the previous two scores to overmatch short sources. To illustrate, in the following example, the match score is 100 percent and the position score is 0 – both optimum values.

Source =		B	A	N						
Target =		B	A	N	D	A	I	D	S	

The leftover score measures the percentage of the target that is not matched to any letter pairs in the source. As in the match score, the leftover score uses a weighting system to give more weight to letter pairs that are most useful in discriminating between spelling variations and non-matching names. The best value is a leftover score of 0, and the worst value is 100 percent.

The comparison of adjacent records ends with one more iteration: by five digit ZIP code. The first two iterations examine records sorted by ten characters of the name and then by ten characters of the street address. The ZIP code iteration sorts all the records by five digit ZIP code and then compares adjacent records within each ZIP code for goodness of fit in both the name and street address fields.

FIND_UNIQUE.SAS allows the user to specify separate threshold values for each score and for each match field. In the iteration where names begin with the same ten non-blank characters, the thresholds for street address matches are relaxed slightly when five digit ZIP codes match exactly.

4.3.4 Matching records within the set of best names and addresses

The most powerful part of FIND_UNIQUE.SAS compares each record within a dataset to every other record, but it is also the slowest. For this reason, it is important to use Part II first to match closely-related records through comparisons of adjacent records.

Part III simultaneously scores and evaluates four match fields: name, street address, state and ZIP code. The program compares each record (source records) to all other records (target records). If the source record matches multiple target records, then the source is assigned to the target with the most frequently reported identifying data.

For example, assume that all of the following records match each other and they are all in the same state:

<u>Name</u>	<u>Street</u>	<u>ZIP</u>	<u>Frequency</u>
BANTER	10 MAIN ST.	12345	10
BANTER	P.O. BOX 40	12345	2
BATNER	10 MIAN ST.	72345	1

The “Frequency” column indicates how many times each version of the identifying data occurs in the database. Ten times, the facility reported its name as “BANTER,” its street address as “10 MAIN ST.” and its ZIP code as “12345.” Since this combination of identifying information occurs more frequently than the other two, FIND_UNIQUE.SAS assigns “BANTER,” “10 MAIN ST.” and “12345” to all thirteen records.

4.3.5 Finding indirect matches

In the final part of FIND_UNIQUE.SAS, the program consolidates all the information about matching records and finds a set of unique facilities. In particular, Part IV finds indirect matches, where record A matches record B and record B matches record C but a comparison of A to C fails the goodness of fit thresholds. In this case, A and C should be matched even though they fail in the direct comparison.

In the following hypothetical example, the first record might match the third record by a four letter-pair match in the name field (_B, BA, ER and R_) with an optimal position score of 0, combined with an exact match in the street address field and an exact five-digit ZIP code match.

<u>Name</u>	<u>Street</u>	<u>ZIP</u>
BANTER	10 MAIN ST.	123450040
TRENTON PLANT	P.O. BOX 40	12345
BATNER TRENTON PLT	10 MAIN ST.	12345

The second record might also match the third record based on good match and position scores in the name field and an exact match in the ZIP code field. Therefore, all three records pertain to a single unique facility, even if the first and second records might fail to match using a direct comparison.

4.3.6 Identifying and assigning the best state match

The fuzzy matching program results in two output files: (1) the original file of off-site facilities in which each observation is labeled with the identification number (“ID_MATCH”) of a unique off-site address for which it matches (approximately 4.4 millions records), and (2) a file which represents the legend of unique off-site records based on the ID_MATCH identification number (44,966 for Reporting Year 2002). The latter file contains the records used in the display of off-site facility information in the RSEI model, such as the best name and address or locational coordinates determined from earlier routines of the fuzzy matching program. However, this unique addresses file does not output the best state associated with each off-site facility as it does for name, street, city, and zip.

To develop a state value for each of the unique off-site addresses, the 44,966 facilities were plotted to retrieve the state in which they mapped to. Similarly, the state corresponding to the best zip value was also retrieved (i.e. BEST_ZIP as determined by the fuzzy matching program). A separate analytical routine was then performed in SAS to determine the BEST_STATE value. This analysis required the following preparatory procedures:

1. The original file of approximately 4.4 million reported off-site facility records was sorted based on the unique off-site identification number it was assigned;
2. The frequency of the reported state within each ID_MATCH group of records was calculated;
3. The state most frequently reported for each ID_MATCH group was retained.

As a result of these procedures, three different fields containing various state values could be compared for each unique off-site facility: the plotted state, the state corresponding to the BEST_ZIP, and the state reported with the highest frequency. The following rules were applied in their comparison and in the determination of the final BEST_STATE value:

1. If the plotted state = BEST_ZIP state = reported state, then the state was considered valid;
2. Alternatively, if any two of the three fields matched, then that state value was used;

3. Finally for instances in which the latitude or longitude = 0 or was blank, no plotted state could be determined, so the reported state, if available, was used.

Of the 44,966 unique addresses that resulted from running the collapse program, all but 34 resolved with a BEST_STATE based on this methodology. The remaining 34 off-site facilities were exported into Excel and manually evaluated since the three state fields were in disagreement. The three state fields were used as a guide and provided context during this manual verification of BEST_STATE. Some of the reasons for how the BEST_STATE was determined for these records, included:

1. They matched a facility record in the larger set that resolved with a BEST_STATE.
2. Some combination of city/zip/state was confirmed on www.usps.com;
3. The state was in the facility name;
4. A search on some combination of the name/street address/city on www.google.com found an exact match.

Among these 34 records were some facilities for which the lat/long coordinates were deleted. Reasons for deleting lat/long coordinates included: (1) they were erroneous (e.g., the facility was actually located in the UK or Canada, or the search on name and address revealed a different state that was NOT adjacent – if the state was adjacent, the lat/long was not deleted), or (2) no supporting data to make a determination could be found using any of the above mentioned methods. These 34 records were then re-appended to the larger off-site address file resulting in the complete set of 44,966 unique off-site addresses.

4.3.7 Results

The geocoding procedure and the SAS algorithms collapsed the number of off-site records from the initial 4.4 million to a final set of 44,966 records. As shown in Table D-9, approximately 34 percent of the unique facilities were matched to high-quality street addresses. Note that each unique address may represent multiple reports of off-site transfers from multiple Form R's.

Table D-9
TCS Off-site Geocoding Match Results after Collapse of Duplicate Records

Coordinate Matches	Number of Records	Match Percentage
Street address (including hand matches)	15,377	34.20%
ZIP+4 centroids	135	0.30%
ZIP+2 centroids	73	0.16%
5-digit ZIP code centroids	26,138	58.13%
Unable to geocode	3243	7.21%
Total Unique Off-site Facilities	44,966	100.00%