

Technical Appendix D

Locational Data for TRI Reporting Facilities and Off-site Facilities

Table of Contents

- 1 Introduction..... 1**
- 2 TRI Reporting Facilities..... 1**
- 3 Off-site Facilities 2**
 - 3.1 Creating the Master Database 3
 - 3.1.1 Collapsing Reported Off-site Facilities 4
 - 3.1.2 Geocoding the Off-site Facilities 4
 - 3.1.3 Collapsing Off-site Facilities After Geocoding 4
 - 3.1.4 Cleaning and Conditioning the Data..... 5
 - 3.1.5 Identifying a Set of Best Names and Addresses 6
 - 3.1.6 Matching Records Within the Set of Best Names and Addresses 9
 - 3.1.7 Finding Indirect Matches 9
 - 3.1.8 Identifying and Assigning the Best State Match..... 10
 - 3.1.9 Results..... 11
 - 3.2 Matching Current Year Data to the Master Database..... 12
 - 3.3 RY 2011 Results 12

1 Introduction

The RSEI model uses latitude and longitude (lat/long) 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 population exposed to air releases and the discharge stream reach (if any).

There are two types of facilities included in the model, TRI reporting facilities and off-site facilities. The quality of data varies significantly between the two types. TRI reporters submitted their own addresses and, prior to Reporting Year (RY) 2005, submitted estimates of their latitude and longitude on Form R. As facility reports of lat/longs were subject to common reporting errors such as transposition of digits, confusion of latitude with longitude, lack of precision, and nonreporting, TRI no longer requires them to be reported, and instead relies on EPA's centralized databases of locational information. RSEI adopts EPA coordinates for most reporting facilities as well.

For off-site facilities, the data quality 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. Little or no standardization is performed by the TRI program, so minor differences in an off-site facility record, such as a slight misspelling of the name, or "St." instead of "Street", can make it difficult to automatically match records.

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 Versions 2.1 through 2.1.3, 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 the current version, RSEI uses coordinates from EPA's Facility Registry System (FRS) for reporting facilities, and uses data from earlier versions only in those cases where the FRS data is not available. For off-site facilities, RSEI has created a master database of off-site facilities based on RSEI Version 2.1.3 that is updated with each new reporting year.

Section 2 describes the method and data sources used for reporting facilities. Section 3 describes how coordinates are determined for off-site facilities, including creating the master database, and matching the current year off-site facilities to that database.

2 TRI Reporting Facilities

The primary source for locational information for TRI reporting facilities is EPA's Facility Registry System (FRS), a centrally managed database developed by EPA's Office of Environmental Information (OEI). The FRS contains accurate and authoritative facility identification records which are subjected to rigorous verification and data management quality assurance procedures. FRS records are continuously reviewed and enhanced by a Regional Data

Steward network and active State partners. RSEI accesses FRS data through EPA’s Geospatial Data Access Project, which provides downloadable files of regulated facilities or sites in various formats. Within the file is key facility information, along with associated environmental interests for use in mapping and reporting applications.¹

The database of all TRI reporting facilities for 1988-2012 includes 56,974 facilities. Of the total number of facilities, 56,276 facilities were assigned coordinates from EPA’s FRS system. These coordinates were applied to the TRI reporting facilities in the RSEI model, and the facilities were assigned the codes “FRS” or “LRT” for LatLongSource². For facilities without valid coordinates in FRS, RSEI uses older data previously submitted to TRI and housed in earlier versions of RSEI, or zip code centroids. In some cases, manual quality assurance led to revision of existing coordinates; these records show a source code of “MANUAL.” Generally, these coordinates were generated through geocoding the address and visually checking placement on a satellite image.

The number of facilities with each type of coordinates (including FRS coordinates) assigned are shown in Table D-1 below.

**Table D-1
Sources for Coordinates for Reporting Facilities**

LatLongSource	Description	Number of Facilities	% of Facilities Without LRT Coordinates
FRS or LRT	Coordinates taken from EPA’s FRS system	56,276	99%
TRI	TRI preferred or submitted coordinates submitted to TRI prior to 2005.	5	<0.01%
ZIP	5-digit ZIP code centroid, either from an Arcview GIS program, or from a commercial web site.	195	0.03%
MAN	Coordinates determined manually, either because lat/long were transposed, or based on visual inspection of map.	498	0.09%

3 Off-site Facilities

Data quality issues with the set of TRI 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

¹ Data can be downloaded from this site: http://www.epa.gov/enviro/geo_data.html.

² Facilities with FRS as the assigned source used updated coordinates; facilities with LRT as a source maintained the same coordinates as the previous year.

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.

In RSEI Version 1.x, all off-site facilities had been located on the model grid using the centroid of the off-site facility's ZIP code. For Versions 2.1 through 2.1.3, the off-site locations were substantially improved using fuzzy matching to collapse the database of off-site facilities, and commercial geocoding to assign a location to each off-site facility. 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. This method was last used in RSEI Version 2.1.3 (RY 2003).

Starting with RSEI Version 2.1.6 (RY 2006), a new method was used that preserves the matches made in previous years (rather than starting from scratch each year). The master database created for Version 2.1.3 was used as a starting point. This master database contains a key that identifies each unique off-site record, and contains every permutation of that record that has been submitted to TRI over the course of TRI reporting. For instance, if a reporting facility listed an off-site transfer as going to "ABC Waste Management" in 2000 and another facility listed one as going to "ABC Waste Mgmt." at the same (or similar) address in 2005, each name would have a separate record in the master database, but be keyed to the same unique record. The table "offsite.db" that is distributed with the RSEI model contains a single record for each unique key.

Each year, the current year TRI data (for all years) is matched back to the master database, and any unmatched facilities are added to the master list. Because there are no unique IDs or keys, there is no easy way to match the current year TRI data to the master. Exact matching is performed first, and then a simplified version of the fuzzy matching used to create the master database is used to match the remaining facilities. The sections below describe the process in detail: Section 3.1 describes how the master database was created- the collapse, geocoding and fuzzy matching. Section 3.2 describes how the current year data is matched to the master database.

3.1 Creating the Master Database

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 3 million to almost 300,000 by removing the exact duplicates. Second, in order to expedite and improve the off-site facility locating process, TCS³ 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.

³ Thomas Computing Services, Lantana, FL.

3.1.1 Collapsing Reported Off-site Facilities

In Version 2.1.3 (RY 2003), there were approximately 3 million off-site facility records in TRI. However, many of these facility records actually represent the same facility; they were just reported in slightly different ways by the facilities transferring chemicals to them. In addition, approximately 1 million records were blank or not viable records. In order to make the geocoding process more efficient, it was necessary to first collapse the list of all reported off-site facilities into possible unique facilities. The first collapsing procedure removed all records that were not viable along with all of the records that are exact duplicates. This first stage collapsed the off-site facility records from approximately 3 million to approximately 300,000.

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 facility, where another form might be better for geocoding purposes. Therefore, to decrease potential error in geocoding unique facilities, the almost 300,000 facility address records were sent to the geocoding service.

3.1.2 Geocoding the Off-site Facilities

TCS evaluated the 300,000 off-site facility address records. Their geocoding efforts resulted in a 50% street address match; 0.18% ZIP+4 centroid match; 0.16% ZIP+2 centroid match; nearly 47% ZIP code centroid match; and nearly 3% unmatched records. At this point in the process, this numbers may be misleading, since many of the 300,000 facilities were 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.

3.1.3 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.

The major parts of FIND_UNIQUE.SAS are:

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.

3.1.4 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-2 lists states that TRI reporters frequently misreport.

**Table D-2
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		

State codes 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.”

3.1.5 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:



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 BAND AIDS,” 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.

3.1.6 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.

As part of this step of the program, the data are exported to an Excel spreadsheet, where some manual matches and corrections supplement the SAS matching. The data are then imported into SAS again, where processing continues.

3.1.7 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.

3.1.8 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 3 millions records), and (2) a file which represents the legend of unique off-site records based on the ID_MATCH identification number. 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 facility as it does for name, street, city, and zip.

To develop a state value for each of the unique off-site addresses, the 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 3 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 unique addresses that resulted from running the collapse program, less than 200 were not resolved with a BEST_STATE based on this methodology. These remaining 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) Some combination of city/zip/state was confirmed on www.usps.com;
- 2) The state was in the facility name;
- 3) Searched on some combination of the name/street address/city on www.google.com for an exact match.

Among these 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. Finally, only three records resulted with no state value at all; and those lat/longs were also deleted because two were in the UK and one was located in Canada. These records were then re-appended to the larger off-site address file resulting in the complete set of unique off-site addresses.

3.1.9 Results

The geocoding procedure and the SAS algorithms collapsed the number of off-site records from the initial 3 million to a final set of 46,768 records. As shown in Table D-3, 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. The master database contains all of these final records (each has a unique identifying number) plus all of the variations of the name and address for each record that were determined to be matches in the process described above. The unique final record and all of its matches share the same identifying number. This allows the final database (which is distributed each year with the RSEI model) to be created from the master database simply by collapsing based on the unique identifier.

Table D-3
Status of All Off-site Facilities in 2003 Final Database

Coordinate Matches	Number of Records	Match Percentage
Street address (including hand matches)	16,610	32%
ZIP+4 centroids	109	<1%
ZIP+2 centroids	48	<1%
5-digit ZIP code centroids	28,738	55%
Unable to geocode or match to geocoded	6,885	13%
Total Unique Facilities	52,390	

3.2 Matching Current Year Data to the Master Database

Once the master database was developed, the next reporting year's data could be matched to it. This involves a process of collapsing, exact matching, and fuzzy matching, as described below. Exact and fuzzy matching are both necessary because there are no unique IDs assigned in TRI to the off-site facilities, and with every reporting year update, all of the unique keys that link the data tables in the TRI data change. So there is no way to link an off-site facility from the current year TRI data, even for an older year, into the master database except by matching the text fields. And a link must be created between the current year records and the master database, because the link to the reporting facility goes through the current year off-site record. If this matching process were not performed, there would be no way to tell what reporting facilities were transferring chemicals to these off-site facilities, or what quantities they were transferring.

Creating the Collapsed Set

The entire set of off-site facilities for all years of TRI reporting include more than 5 million records. The invalid or blank records are first removed, then exact duplicates are removed, so that the collapsed set contains around 400,000 unique records. These records are unique in the sense that there are no exact duplicates; however, many of the records in reality refer to the same off-site facility.

Exact Matching

The set of current year off-site facilities with the blank/invalid/duplicate records removed is then matched to the master database based on exact field matches. Approximately 90% of the records from the collapsed set can be matched exactly to records in the master database. This is because the master database was created from the previous year's TRI data, and most of the data is duplicative (for reporting years prior to the latest).

Fuzzy Matching

The approximately 10% of the remaining off-site records from the collapsed current year set are then matched to the master database using a fuzzy matching program similar to the one described above. This process matches approximately 25% of the records from the collapsed set to records in the master database. Of the unmatched remainder, the POTWs and incinerators (the only off-sites that are modeled) are examined manually, resulting in several hundred additional matches. The remaining unmatched off-site records are added to the master database, which then becomes the starting point for next year's off-site processing.

3.3 RY 2012 Results

For RY 2012, the TRI all-years data contained over 6 million records. Once the invalid or blank and exact duplicate records were removed, the collapsed set contained approximately 400,000 unique records. These records are unique only in the sense that there are no exact duplicates; many of the records refer to the same facilities. Most records in the collapsed set matched exactly to the master database. The unmatched remainder (4,350 records) was fuzzy-matched to the master, and approximately half were matched. Of the remainder, approximately 563 were matched by hand, and the remaining records were added to the master list as unique records.

For RY2007 through RY2009, the facility finder function in the EPA application OTIS⁴ was used to improve the locational data for off-site facilities. Using name and address, the off-sites (incinerators and POTWs only) were matched into EPA's FRS system, from which the FRS ID was obtained. The FRS ID was used to pull the locational data from EPA's Geospatial Data download service.⁵ Using OTIS, 6,657 off-site records were linked to 5,399 unique FRS IDs, 2,029 of which had valid locational data available. OTIS is no longer updated by EPA.

Beginning in RY2010, EPA's Discharge Monitoring Report (DMR)⁶ Pollutant Loading Tool was used to supplement the locational information for POTWs. DMR data have been vetted by EPA's Office of Water and includes data on all facilities regulated under the NPDES permit program, as reported to both the Permit Compliance System (PCS) and ICIS-NPDES. The tool includes DMR data for all "major dischargers" (POTWs that discharge 1 million gallons per day or more) and "minor dischargers" when available. For some POTWs, DMR data includes the lat/long of one or more outfall pipes used to discharge effluent. Where off-site facilities in the master database could be linked to a record in the DMR data and did not have address-matched locational information, the DMR lat/long was adopted.

For the rest of the off-site facilities, results from the 2003 geocoding were used, if available. For off-site facilities that could not be matched using OTIS, DMR, or matched to geocoded facilities, zip code centroids were used. Note that this table (and all of the results discussed in this appendix) describes data in the all-years table. A subset of this table (containing TRI reporting years 1996-2011) is contained in the RSEI public release. Over half of the off-site facilities in the final distributed table have coordinates based on zip code centroids. However, many of the off-site facilities are not POTW or incineration facilities, so they are not modeled. Of the POTWs and incineration facilities, over half of the records have coordinates based on address matching or GPS techniques. This may not represent the actual number of facilities; processing the off-site records eliminates some duplicates, but duplicates remain even in the distributed file. The process errs on the side of not matching, since a false match could result in erroneous modeling results.

⁴ http://www.epa-otis.gov/otis/permit_finder.html

⁵ http://www.epa.gov/enviro/geo_data.html.

⁶ <http://cfpub.epa.gov/dmr/>.