

# Exploiting Redundancy, Recurrence and Parallelism: How to Link Millions of Addresses with Ten Lines of Code in Ten Minutes

Yuhang Zhang

Tania Churchill

Kee Siong Ng

Australian Transaction Reports and Analysis Centre

Email: { yuhang.zhang, tania.churchill, keesiong.ng }@austrac.gov.au

## Abstract

Accurate and efficient record linkage is an open challenge of particular relevance to Australian Government Agencies, who recognise that so-called wicked social problems are best tackled by forming partnerships founded on large-scale data fusion. Names and addresses are the most common attributes on which data from different government agencies can be linked. In this paper, we focus on the problem of address linking. Linkage is particularly problematic when the data has significant quality issues. The most common approach for dealing with quality issues is to standardise raw data prior to linking. If a mistake is made in standardisation, however, it is usually impossible to recover from it to perform linkage correctly. This paper proposes a novel algorithm for address linking that is particularly practical for linking large disparate sets of addresses, being highly scalable, robust to data quality issues and simple to implement. It obviates the need for labour intensive and problematic address standardisation. We demonstrate the efficacy of the algorithm by matching two large address datasets from two government agencies with good accuracy and computational efficiency.

*Keywords:* record linkage, address linking

## 1 Introduction

Efficient record linkage is an important step in large-scale automated data fusion. Data fusion is a problem of increasing significance in the context of Australia's whole-of-government approach to tackling our most pressing social issues - including terrorism and welfare fraud - by combining and analysing datasets from multiple government agencies. Outside of personal identifiers like tax-file numbers and driver's licenses, names and addresses are the two most important attributes on which disparate datasets are matched. Whereas the problem of linking names is well-studied and there are specialised similarity measures like Jaro-Winkler for names (Cohen et al. 2003), not a great deal is known in the literature (Christen 2012) about best practices for matching addresses, especially address data with significant quality issues. Listed here are some address-specific challenges for efficient record linkage:

- **Incompleteness:** incomplete addresses that have no street type, no suburb name, no postcode, etc are common in address data.
- **Inconsistent Formats:** the structure of addresses can be different between countries, regions and languages. People can use variants to denote the same address, *e.g.*, unit 1 of 2 Elizabeth Street, 1/2 Elizabeth St, and U-1 2 Elisabeth Str all denote the same address. The presence of foreign characters in international addresses can also introduce issues.
- **Errors:** Wrong street types, invalid postcodes, non-matching suburb-postcode pairs, and various misspellings are widely seen in address data.
- **Unsegmented Addresses** Depending on the source, addresses can be captured as a single line of text with no explicit structural information.

These data quality issues may make two equivalent addresses look different and, by chance, make two different addresses look similar.

The most common way to tackle data quality issues is to standardise raw addresses before the linking operation (Christen 2012). Address standardisation usually includes two types of operations: parsing and transforming. With the parsing operation, addresses are parsed into semantic components, such as street, suburb, state, and country. For example, if an address contains three numbers they are in order unit number, street number, and postcode. In the transforming operation, variants of the same entity are transformed to a canonical format and typos are removed, *e.g.*, transforming Street, St, and Str all to Street.

The issue with standardisation is that it is in itself a challenging problem. For example, *Service Centre St George* might be interpreted as a business name *Service Centre of Saint George*; or a street name and a suburb name *Service Centre Street, George*; or a different business name *Service Centre of Street George*; or a suburb name and a state name *Service Centre, Saint George*. Three numbers in an address can also be street number, level number in a high-rise, and postcode. Interpreting an address is by nature ambiguous.

Address standardisation can be done using a rule-based system, or it can be done using machine learning approaches like Hidden Markov Models (Christen et al. 2004, Christen & Belacic 2005). Ongoing research is still being undertaken to improve standardisation accuracy (Guo et al. 2009). Perhaps the biggest drawback of address standardisation is that if a mistake is made during standardisation, it is usually hard to recover from it to perform linkage correctly. Rule-based standardisation also tends to be specific to the individual dataset, failing to generalise well.

## Using Redundancy to Avoid Standardisation

Instead of standardising raw addresses into canonical forms, we rely on the redundancy in addresses to resolve data quality issues.

We say an address contains redundancy if an incomplete representation is sufficient to uniquely identify this address. For example, if there is only one building in Elizabeth St that has Unit 123, then *U123 45 Elizabeth St* as an address contains redundancy, because specifying street number 45 is not really necessary. Redundancy exists widely in addresses. Not every suburb is covered by postcode 2600. Not every state has a street named Elizabeth. As an extreme example, three numbers like 18 19 5600, might be enough to identify a unique address globally, as long as no other addresses contain these three numbers simultaneously. Note that in this case, we do not even need to know whether 18 is a unit number or a street number.

Our working hypothesis is that address data, in general, contains enough redundancy such that:

1. each address is still unique even when meta-data distinguishing address components such as street, suburb, and state are missing.
2. equivalent addresses are still more similar to each other than to irrelevant addresses in the presence of errors or variants.

Our assumptions - which stem from earlier experiments using compressed sensing techniques (Cands et al. 2006) to represent and link addresses - are really stating that despite the data quality issues in addresses, two addresses, in their raw form, can still be separated/linked if they are different/equivalent. In particular, address segmentation - a problem that is arguably as difficult as the general address-linking problem - and address standardisation are not strictly necessary.

## Using Recurrence for Data-Driven Blocking

When linking two large databases, algorithm efficiency is as important as algorithm accuracy. An algorithm that takes days to finish is not only too expensive to deploy, but is also infeasible to repetitively evaluate during development.

Blocking is a widely used technique to improve linkage efficiency. Naïvely, linking two databases containing  $m$  and  $n$  addresses respectively requires  $O(mn)$  comparisons. Most of these comparisons lead to non-matches. To reject these non-matches with a lower cost, one may first partition the raw addresses according to criteria selected by a user. These criteria are called blocking keys, which may be postcode, suburb name, *etc.*. During linkage, comparison is only carried out between addresses that fall into the same partition, based on the assumption that addresses which do not share a blocking key are not a match.

Blocking key selection largely determines the efficiency and completeness of address linkage. If the keys are not meaningful, they will not help find matches and may even slow down the matching process. If too few keys are used, efficiencies will not be gained. If too many keys are used, one may fail to discover all possible links. If different blocking keys do not distribute evenly among the addresses, the largest few partitions will form the bottleneck of linkage efficiency. Moreover, the performance of blocking keys in previous work also depends on the accuracy of address standardisation.

In the spirit of (Halevy et al. 2009), we propose in this paper a data-driven approach to select blocking keys based on their recurrence. These data-driven blocking keys are by design adapted to the database at hand, statistically meaningful as address differentiators, evenly distributed, and provide comprehensive cover to all addresses. Since we implement no standardisation, our blocking keys do not depend on the success of standardisation either.

## Implementation on Parallel Platforms

Massively parallel processing databases like Teradata and Greenplum have long supported parallelised SQL that scales to large datasets. Recent advances in large-scale in-database analytics platforms (Hellerstein et al. 2012), (Zaharia et al. 2010) have shown us how sophisticated machine learning algorithms can be implemented on top of a declarative language like SQL or MapReduce to scale to petabyte-sized datasets on cluster computing. Building on the same general principle, we propose in this paper a modified inverted index data structure for address linking that can be implemented in less than ten SQL statements and which enjoys good scalability and code maintainability.

## Paper Contributions

The paper's contribution is a novel address-linkage algorithm that:

1. links addresses as free-text (including international addresses), obviating the need for labour-intensive and sometimes problematic address standardisation;
2. uses data-driven blocking keys to minimise unnecessary pairwise comparisons, in a way that obviates the need for address segmentation and avoids the usual worst-case scenarios encountered by using a fixed blocking key like suburb or postcode;
3. introduces an extension of the inverted index data structure that allows two large address datasets to be linked efficiently;
4. is practical because of its simplicity, allowing the whole algorithm to be written in less than 10 standard SQL statements; and
5. is scalable when the SQL statements are implemented on top of parallel platforms like the Greenplum Database (open-source parallel PostgreSQL) and Spark.

The algorithm is particularly suitable for integrating large sets of disparate address datasets with minimal manual human intervention. It is also possible to combine the algorithm with a rule-based system to produce a model-averaging system that is more robust than each system in isolation.

The remaining sections of this paper are organised as follows. We first explain how a single address can be linked to an address database utilising redundancy. We then show how the same algorithm can be carried out in batch taking advantage of recurring address components. Implementation techniques on state-of-the-art parallel computing platforms are then described. Finally, we demonstrate the performance of our algorithm with two large-scale address linkage applications, followed by our conclusion.

## 2 Address as Bag of Tokens

Without subfield structures, an address becomes a bag (or a multiset) of unordered tokens. For example,

No.	street	suburb	state	postcode
513	Elizabeth St	Melbourne	VIC	3000

becomes

$\{3000, 513, \text{Elizabeth}, \text{Melbourne}, \text{Street}, \text{VIC}\}$ ,

In this example, we implicitly define a token to be a word, or a maximal character sequence that contains only letters and numerics. We can also define a token to be a single character, where the address becomes

$\{0, 0, 0, 1, 3, 3, 5, \text{a}, \text{b}, \text{b}, \text{c}, \text{e}, \text{e}, \text{e}, \text{e},$   
 $\text{h}, \text{i}, \text{l}, \text{l}, \text{n}, \text{o}, \text{r}, \text{s}, \text{t}, \text{t}, \text{t}, \text{u}, \text{v}, \text{z}\}$ ,

or a two-word phrase, in which case we have

$\{513 \text{ Elizabeth}, \text{Elizabeth St},$   
 $\text{St Melbourne}, \text{Melbourne VIC}, \text{VIC } 3000.\}$

or generally anything we like. In the above example, two-word phrases preserve pairwise order information in the original address. We can also use two word tokens that do not contain pairwise order information.

Different types of tokens have different distinctiveness powers and different tolerances against data quality issues. To see the difference, note the word token, ‘Melbourne’, can match to any appearance of ‘Melbourne’ in other addresses, such as Melbourne Avenue, Mount Melbourne, Melbourne in Canada, *etc.*. By contrast, the phrase token, ‘Melbourne VIC’, can only match the co-occurrence of ‘Melbourne’ and ‘VIC’. The advantage of being distinctive is that we can reduce false matches. The disadvantage, however, is that we may miss a true match if the other address did not include the state information of ‘VIC’ or included it in a different form, *e.g.*, Victoria.

For the purposes of linkage, we do not need individual tokens to be distinctive. Instead, we want tokens to be tolerant to data quality issues. We lose nothing as long as a bag of tokens as a whole is distinctive enough to identify an address uniquely. However, for matching efficiency we prefer distinctive tokens. We will come back to this topic after we explain how to measure the similarity between two addresses as two bags of tokens.

## 3 Similarity between Bags of Tokens

We assess the similarity between two addresses using the Jaccard index between their bag-of-tokens representations. The Jaccard index of two multisets is defined as the ratio between the number of common elements and the number of total elements.

$$J(T_1, T_2) = \frac{|T_1 \cap T_2|}{|T_1 \cup T_2|} \quad (1)$$

Here is an example.

$$\begin{aligned} T_1 &= \{\text{this}, \text{is}, \text{an}, \text{example}\} \\ T_2 &= \{\text{this}, \text{is}, \text{another}, \text{example}\} \\ T_1 \cap T_2 &= \{\text{this}, \text{is}, \text{example}\} \\ T_1 \cup T_2 &= \{\text{this}, \text{is}, \text{an}, \text{example}, \text{this}, \text{is}, \\ &\quad \text{another}, \text{example}\} \\ J(T_1, T_2) &= \frac{|T_1 \cap T_2|}{|T_1 \cup T_2|} = \frac{3}{6} = 0.5 \end{aligned}$$

As one can see, the Jaccard index is always in the range between 0 and 1. Here 0 indicates that two multisets have nothing in common, and 1 that the two multisets are exactly the same. The more common elements two multisets share relative to the total number of tokens they have, the larger their Jaccard index is. We say two addresses are equivalent if their Jaccard index exceeds a threshold  $\tau$ .

We shall see in Section 9 that the algorithm admits other similarity functions too.

## 4 Inverted Index

Naïvely linking an address to a database requires comparing this particular address against each database address to obtain their similarity. An inverted index allows us to do the linking in sublinear time.

An inverted index keeps, for each distinct token, references to all the addresses which contain the token. When a query address arrives, an inverted index allows us to know which database addresses share common tokens with the query address without scanning through the database. More specifically, given a query address, we first break this query address into a bag of tokens  $Q$ . If a token is not included in the inverted index, we simply ignore the token. Each remaining token selects a segment from the inverted index. Database addresses appearing on these segments share at least one common token with the query address and therefore have a non-zero Jaccard index with the query address.

We can count the number of occurrences of each database address  $C_i$  on these segments, which gives us the value of  $|Q \cap C_i|$  for each  $i$ . We then derive the value of  $|Q \cup C_i|$  for each  $i$  using

$$|Q \cup C_i| = |Q| + |C_i| - |Q \cap C_i| \quad (2)$$

We can then calculate the Jaccard index between the query address  $Q$  and each candidate address  $C_i$  using Eq 1.

With an inverted index, we only compute the Jaccard index between a query address and those database addresses whose Jaccard indexes are non-zero. The efficiency of address linkage therefore depends on the number of addresses that share at least one token with the query address, not the size of the database.

## 5 Two-Round Linkage

Recall our earlier discussion that tokens of different types have different distinctiveness. The number of database addresses that contain a more distinctive token is by definition smaller than the number of database addresses that contain a less distinctive token. We therefore have better linking efficiency with more distinctive tokens. Yet in return, we may miss more matches due to data quality issues.

To maximise linking efficiency while minimising the number of missed matches, we propose a two-round linkage scheme. In the first round, we use distinctive tokens, *e.g.*, phrase tokens, and inverted indexes to short-list database addresses which have non-zero Jaccard indexes with the query address. In the second round, we compute the Jaccard index between the query and short-listed addresses using less distinctive tokens to account for data quality issues.

In this way, the distinctive tokens decide which database entries get shortlisted in round 1 of the linkage scheme. The less distinctive tokens decide the similarity between a query and a database entry. A

database entry gets shortlisted in the first round as long as it shares a distinctive token with the query. A database entry matches a query if they have enough less-distinctive tokens in common.

The two-round linkage strategy is similar to the one described in (Arasu et al. 2006).

## 6 A Batch Linkage Algorithm

Quite often, we need to find equivalent addresses between two large databases each containing tens to hundreds of millions of addresses. A naïve pairwise matching of every combination of addresses is obviously infeasible in that case. We describe in this section a simple extension of the inverted index data structure to allow efficient linking of two large address databases.

To do batch linking between two databases, we build separate inverted indexes for each database. From each inverted index, we eliminate all the tokens that recur more than  $k$  times. (More on that soon.) We then join the two inverted indexes by the common tokens they share. Joining a pair of common tokens essentially joins two sets of addresses from the two databases. Every pair of addresses from these two sets is a potential match because they share at least a common token. Between these pairs, we then compute the round 2 Jaccard index to identify the true matches.

We eliminate tokens that recur more than  $k$  times because if a token is too common, addresses linked by this token are not likely to be true matches. Moreover, examining addresses linked by a common token takes a lot of time but does not lead to proportionally more matches. Ignoring these common tokens will not miss many true matches because these matches are usually also linked by some more distinctive tokens. The price we pay is that addresses that contain *only* common tokens will not be indexed and therefore require special handling.

Linking one address at a time can be seen as a special case of batch linkage, *i.e.*, one of the databases contains only one address. The advantage of batch linkage over performing a single linkage many times is that in batch linkage we join the two inverted indexes only once, instead of many times.

Our batch linkage can be explained in the traditional framework of data linkage, where joining two inverted indexes implements (data-driven) blocking. There are also some notable differences. Instead of using fixed blocking keys like postcode and/or suburb, we use tokens as blocking keys. Importantly, multiple blocking keys are used and deciding which token is used as a blocking key is determined by the data, more specifically its recurring frequency. This allows the algorithm to adapt to characteristics of the specific databases to be matched, including foreign address databases that we may not understand well enough to construct blocking keys in the usual way.

### Computing Jaccard Index in Linear Time

The above extension of inverted indexes applies to the first of the two-round linkage scheme described earlier. The second round of pairwise Jaccard calculations of short-listed candidate address pairs is done using the algorithm described next.

We first sort the tokens in each multiset. This can be done efficiently since the number of distinct tokens is small. We then sort the tokens, and read from the two multisets at the same time following the rules below:

1. If the two tokens read in are the same, we increase the number of common tokens and the number of total tokens both by 2. We read one more token from each multiset.
2. If one token is larger than the other, we increase the number of total tokens by 1. We read one more token from the multiset whose current token is smaller.

We finish reading when either multiset is exhausted, and add the number of remaining tokens in the other multiset into the number of total tokens. The division between the number of common tokens and the number of total tokens then provides the Jaccard index.

For small tokens (like characters or 2-grams), the time complexity of the algorithm is  $O(l + r)$ , where  $l$  and  $r$  denote the number of tokens in the two multisets.

## Parallel SQL Implementation

The full algorithm in (almost ANSI) SQL is listed in Algorithm 1. The SQL code runs on Greenplum, an open-source parallel implementation of PostgreSQL that scales to hundreds of terabytes of data. The **DISTRIBUTED BY** keyword in the table creation statements specifies how the rows of a table are stored distributively across a cluster by hashing on the distribution key. The Greenplum database query optimiser will exploit the structure of the SQL query and the underlying data distribution to construct optimal execution plans.

With minor modifications, the SQL code can be modified to run on other parallel databases like Teradata and Netezza, and parallel platforms like Spark (using Spark SQL) and Hadoop (using HIVE, HAWQ (Chang et al. 2014) or Impala (Kornacker et al. 2015)). It is also straightforward to implement the algorithm in Scala/Python running natively on Spark.

## 7 Experiments

We demonstrate the performance of our proposed algorithm in two scenarios: linking an address dataset against a reference address dataset, and linking two arbitrary address datasets. In the first scenario, for each address in the first dataset, it can be assumed that there exists a match in the reference dataset. In the latter scenario, we have to provide for the case where there is no match for an address.

### 7.1 Linking with a Reference Dataset

This scenario usually occurs during address cleansing. We deal with two address databases. The first database contains raw addresses, whereas the second database contains reference addresses. For each raw address, we search for its equivalent reference address, which provides a cleansed or standardised representation of the raw address. In this experiment, we use the following two address databases:

- **AGA1** is a raw database collected by an Australian Government Agency. The database contains around 48 millions addresses most of which are Australian addresses. Addresses in this database are known to have significant data quality issues, with many incomplete and inaccurate addresses.
- **OpenAddress.Australia** contains more than 19 millions Australian addresses. All addresses

**Algorithm 1** SQL Code for Batch Linkage

```

1: %% Original address data
CREATE TABLE address_db
( address_id bigint,
  address text
);
DISTRIBUTED BY (address_id);
2: %% Compute 2-word phrase tokens
CREATE TABLE address_db_phrase
( address_id bigint,
  token_phrase text
);
DISTRIBUTED BY (token_phrase);
3: INSERT INTO address_db_phrase
SELECT address_id,(regexp_matches(
  regexp_replace(address,'[^A-Z0-9]+' , '' , 'g')
  ,'[A-Z0-9]+' + [A-Z0-9]+' , 'g'))[1]
FROM address_db
UNION
SELECT address_id,(regexp_matches(
  regexp_replace(
    regexp_replace(address,'[^A-Z0-9]+' , '' , 'g')
    ,'[A-Z0-9]+' , ''
    ,'[A-Z0-9]+' + [A-Z0-9]+' , 'g'))[1]
FROM address_db;
4: %% Compute inverted index
CREATE TABLE address_db_phrase_inverted
( token_phrase text,
  address_ids bigint[],
  frequency bigint
);
DISTRIBUTED BY (token_phrase);
5: INSERT INTO address_db_phrase_inverted
SELECT token_phrase,array_agg(address_id),count(1)
FROM address_db_phrase
GROUP BY token_phrase;
6: %% Compute matched address arrays
CREATE TABLE address_db_phrase_matched
( token_phrase text,
  address_ids_1 bigint[],
  address_ids_2 bigint[]
);
DISTRIBUTED BY (token_phrase);
7: %% Here, we assume address_db_phrase_inverted_2 is the
second address dataset, processed in the same way.
INSERT INTO address_db_phrase_matched
SELECT l.token_phrase,l.address_ids,r.address_ids
FROM address_db_phrase_inverted_1 AS l
INNER JOIN address_db_phrase_inverted_2 AS r
ON l.token_phrase=r.token_phrase
WHERE l.frequency ≤ 100 AND r.frequency ≤ 100;
8: %% Unnest the candidate address pairs
CREATE TABLE address_db_proposed_match
( address_id_1 bigint,
  address_id_2 bigint
);
DISTRIBUTED BY (address_id_1);
9: INSERT INTO address_db_proposed_match
SELECT DISTINCT address_id_1,unnest(address_ids_2)
FROM (
  SELECT unnest(address_ids_1) AS address_id_1,
  address_ids_2 FROM address_db_phrase_matched
) AS tmp;
10: %% Compute round 2 Jaccard index
CREATE TABLE address_db_match AS
SELECT address_id_1, address_id_2,
  jaccard(t2.address, t3.address)
FROM address_db_proposed_match t1,
  address_db_1 t2,
  address_db_2 t3
WHERE t1.address_id_1 = t2.address_id
AND t1.address_id_2 = t3.address_id

```

are in standard form. This reference address database is open-source and can be downloaded from <https://openaddresses.io>. Almost all Australian addresses in AGA1 have a reference entry in OpenAddress.Australia.

We use the batch linkage algorithm to link addresses in AGA1 with addresses in OpenAddress.Australia. We extract order-preserving 2-word phrase tokens from the addresses and construct in-

verted indexes for both databases. We then compute character-based Jaccard index between each pair of short-listed candidates. We accept a link if the Jaccard index exceeds a threshold  $\tau$ .

Since we do not have ground truth for the address cleansing result, we cannot quantitatively assess the rate of false negatives (i.e. there exists a cleansed entry for a raw address but the algorithm cannot find it) in our linkage result. It is fair to say that essentially all data operations involving large databases have the same problem. It is therefore difficult to select the proper threshold value  $\tau$ . We propose the following mechanism for threshold selection. We implement address linkage with increasing thresholds, *e.g.*,  $\{\tau_1 = 0.6, \tau_2 = 0.7, \tau_3 = 0.8\}$ . We then use the result of the lowest threshold to benchmark that of higher thresholds for false negatives.

Figure 1 shows the percentage of true positives, false positives, and false negatives for the proposed method. These results are obtained by manually assessing 100 randomly sampled linked addresses. As we can see, when  $\tau = 0.6$ , which roughly requires a cleansed address to share 60% or more characters with the raw address, nearly 40% of raw addresses will find false cleansed forms. When  $\tau$  increases to 0.7, the percentage of false positives drops to 12%. Conversely, 2% of raw addresses which used to find cleansed forms can no longer find them. This missing rate rises to 31% when  $\tau$  increases to 0.8. Among the three values,  $\tau = 0.7$  gives the best tradeoff in performance.




Figure 1: the percentage of true positives (true), false positives (false), and false negatives (missed) of proposed addressing linkage algorithm with different thresholds.

Table 1 lists 10 example links between AGA1 and OpenAddress.Australia found by our algorithm. Besides the two linked addresses, we also provide the clean addresses found by Google Maps for AGA1 addresses. To protect anonymity we have encrypted the street names and suburb names. Interestingly, there are three addresses that Google Maps failed to process, yet were successfully linked by our algorithm. The fourth example in the table shows the limitation of using characters as tokens in the Jaccard index calculation. A simple tie-breaker postprocessing scheme using, for example edit distance, can be used to resolve such issues.

## 7.2 Linking Two Arbitrary Datasets

This scenario occurs when two distinct databases need to be fused together. To test this scenario, we use two databases AGA1 and AGA2.

Table 1: Address Linkage between AGA1 and OpenAddress\_Australia

	Address	Jaccard
AGA1	11 16-18 EEHMNTV DIRTTO NSW 2770	
Open	UNIT 11 16-18 EEHMNTV STREET MOUNT DIRTTO NSW 2770	0.77
Google	11/16-18 EEHMNTV ST MOUNT DIRTTO NSW 2770	
AGA1	21 360 ADGNR EFORST AKLR QLD 4178	
Open	UNIT 21 360 ADGNR AVENUE EFORST LAKE QLD 4078	0.79
Google	21/360 ADGNR AVE EFORST AEKL QLD 4078	
AGA1	435 6 CEOPRW LLLMOOOOOOOOW NEW WALES 2011	
Open	UNIT 435 6 CEOPRW AFHRW ROADWAY LLLMOOOOOOOOW NSW 2011	0.80
Google	NOT FOUND	
AGA1	301 100 AGGHILNNU AGGHILNNU ACT 2912	
Open	UNIT 310 100 AGGHILNNU PLACE AGGHILNNU ACT 2912	0.87
Open	UNIT 301 100 AGGHILNNU PLACE AGGHILNNU ACT 2912	0.87
Open	UNIT 103 100 AGGHILNNU PLACE AGGHILNNU ACT 2912	0.87
Google	301/100 AGGHILNNU PL AGGHILNNU ACT 2912	
AGA1	1 24 FGGIIS DEKNOR QLD QLD 4031	
Open	UNIT 1 24 FGGIIS STREET DEKNOR QLD 4031	0.80
Google	1/24 FGGIIS ST DEKNOR QLD 4031	
AGA1	4 NO 3 TO 5 CEELMNT ADDEGNNNO VIC 3175	
Open	UNIT 4 3-5 CEELMNT STREET ADDEGNNNO VIC 3175	0.85
Google	NOT FOUND	
AGA1	101 9 EEEGKNNORW ABEEHILTZ BAY 2011	
Open	UNIT 101 9 EEEGKNNORW AVENUE ABEEHILTZ BAY NSW 2011	0.81
Google	101/9 EEEGKNNORW AVE ABEEHILTZ BAY NSW 2011	
AGA1	MARGETIC 7 302 ABBDFOORST BEELMNORU VIC 3051	
Open	FLAT 7 302 ABBDFOORST STREET HNORT BEELMNORU VIC 3051	0.82
Google	NOT FOUND	
AGA1	61 1162 AACDHINSV AAGRITTV QLD 4122	
Open	UNIT 61 1162 ACDEHINSV ROAD MOUNT AAGRITTV EAST QLD 4122	0.76
Google	61/1162 ACDEHINSV RD MOUNT AAGRITTV EAST QLD 4122	
AGA1	38 91 ADELMNOR BELM VIC VIC 3011	
Open	UNIT 38 91 ADELMNOR STREET ACFOORSTY VIC 3011	0.73
Google	38/91 ADELMNOR ST ACFOORSTY VIC 3011	

- **AGA2** contains around 18 millions addresses collected by a large Australian government department. Most addresses in AGA2 are Australian addresses. Addresses in this database may be incomplete and inaccurate. AGA1 and AGA2 are collected by different government agencies from different sources and for largely different original purposes.

We again use the batch linkage algorithm with 2-word phrase tokens for round 1 of Jaccard computations and character tokens for round 2. However, in this second address-linkage scenario, we can no longer use a simple threshold  $\tau$  to reject false matches. This is because when linking with a reference dataset, if a street is included in the reference database, all individual addresses in this street are included. Therefore, if a raw address has a high score best match in the reference database, this best match is usually consistent with the raw address in every detail. However, in the scenario where we are linking two arbitrary databases, it is quite common for two databases to contain only two different addresses in the same street. These two addresses may have the highest matching score but still remain a false match. To complicate matters, a true match can also be a low score match due to data quality issues with both addresses.

One way to overcome this challenge is to require two matching addresses to have consistent numeric tokens. We say two sets of numeric tokens are consistent, if one set is a subset of the other.

We manually assess 100 randomly sampled AGA2 addresses. For each AGA2 address, we consider, in order, its top 3 matches in the AGA1 database. If a match has consistent numeric tokens and is a true match, we label this AGA2 sample as true and no longer consider the remaining matches. If a match has consistent numeric tokens but is a false match, we la-

bel this AGA2 sample as false and no longer consider the remaining matches. If none of the top 3 matches has consistent numeric tokens with the query, this AGA2 sample is labelled as not found. Figure 2 shows the percentage of three labels in the 100 samples. It can be derived from Figure 2 that,  $59/(59+6) = 91\%$  of the samples are correctly linked, which compares well against an in-house rule-based address matching system.




Figure 2: Percentage of correctly linked (True), incorrectly linked (False), and not linked (Not Found) when joining AGA2 addresses to AGA1 addressees using our proposed algorithm.

Table 2 lists 10 example links between AGA1 and AGA2 found by our algorithm.

### 7.3 Computational Efficiency

We now report on computation times. When dealing with a large database, algorithm efficiency is as important as algorithm accuracy, because an algorithm that takes days to finish is too expensive to deploy, and even more expensive to test under multiple con-

Table 2: Address Linkage between AGA1 and AGA2

	Address	Jaccard
AGA1	4 EGNNOORU AEHLMT VIC 3095	
AGA2	4 EGNNOORU CRT AEHLMT NORTH VIC 3095	0.84
AGA1	528 LLTUY HLLU LLTUY HLLU QLD 4854	
AGA2	528 LLTUY ADEHHLLSU RD LLTUY QLD 4854	0.87
AGA1	45 EGHIMMNS ADDEGNNNO VIC 3175	
AGA2	RM 8 45 EGHIMMNS ST ADDEGNNNO VIC 3175	0.88
AGA1	6 EEILLS CEIMNRTY ABILRSSUY DNOSW SA 5108	
AGA2	6 EEILLS CEIMNRTY RD ABILRSSUY DNOSW SA 5108	0.96
AGA1	EL EILOSU 137 AAILNT EFNRY EOY QUEEN SLAND 4055	
AGA2	137 AAILNT RD EFNRY EGOV QLD 4055	0.74
AGA1	80 ABEGLNRRU DEHILLS VI 3037	
AGA2	80 ABEGLNRRU DR DEHILLS VIC 3037	0.94
AGA1	141 ACEHLRS EHPRT 6005	
AGA2	141 ACEHLRS ST ESTW EHPRT WA 6005	0.80
AGA1	51 BENOOR BELMNOT VIC 3216	
AGA2	2/51 BENOOR DR BELMNOT VIC 3216	0.91
AGA1	97 ELOXY ABDPRUY WA 6025 TRA LIA	
AGA2	97 ELOXY AVE ABDPRUY WA 6025	0.87
AGA1	9 DLORS DENSYY 2077	
AGA2	9 DLORS AVE AHQSTU NSW 2077	0.68

figurations. Experiments show that our algorithm is highly efficient and scalable to large databases.

Using the open-source Greenplum Database running on 8 servers (1 master + 7 slaves), each with 20 cores, 320 GB, and 4.5 TB usable RAID10 space, linking 48 million AGA1 addresses with 13 million OpenAddress addresses using our algorithm takes about 5 minutes. Linking 48 million AGA1 addresses with 18 million AGA2 addresses takes about 7.5 minutes. The algorithm also scales essentially linearly in the number of servers in the Greenplum cluster.

Note that the processing time of our algorithm depends more on the similarity between two databases than on the sizes of the two databases. The efficiency of the algorithm is due to the following factors:

1. The quantity of Jaccard index computation does not depend on the size of the database, but the number of addresses sharing common distinctive tokens.
2. Finding addresses sharing common distinctive tokens is done jointly for all addresses at the same time. This overhead does not depend on the number of addresses, but the number of distinctive tokens during the joining between two inverted indexes.

## 8 Parameters of the Algorithm

The use of Jaccard index to assess similarity between addresses in our algorithm is optional. Our implicit assumption is that there exists a function  $d(x, y)$  which assesses the similarity between two addresses  $x$  and  $y$ . Blocking can reduce the number of evaluations of  $d(x, y)$  without missing links, if  $d(x, y) > \tau$  indicating  $x$  and  $y$  share a common token. In our two-round linkage, our implicit function is

$$d(x, y) = \begin{cases} J_{char}(x, y) & \text{if } J_{phrase}(x, y) > 0 \\ 0 & \text{otherwise} \end{cases} \quad (3)$$

One may design any other implicit function instead, replacing the Jaccard index with any other measurement.

The Jaccard index in round 2 of the comparison can also be replaced by almost any other similarity function, for example the Monge-Elkan function (Monge & Elkan 1997), which is suitable for addresses.

## 9 Conclusion

We have presented in this paper a novel address-linkage algorithm that:

1. links addresses as free text;
2. uses data-driven blocking keys;
3. extends the inverted index data structure to facilitate large-scale address linking;
4. is robust against data-quality issues; and
5. is practical and scalable.

The simplicity of the solution – an important virtue in large-scale industrial applications – may belie the slightly tortuous journey leading to its discovery; a journey laden with the corpses of a wide-range of seemingly good ideas like compressive sensing and other matrix factorisation and dimensionality-reduction techniques, nearest-neighbour algorithms like KD-trees, ElasticSearch with custom rescoring functions (Gormley & Tong 2015), rules-based expert systems, and implementation languages that range from low-level C, to R, Python, SQL and more.

In retrospect, our algorithm can be interpreted as an application of a signature-based approach to efficiently compute set-similarity joins (Arasu et al. 2006), where the abstract concept of sets is replaced with carefully considered bag-of-tokens representations of addresses, with a modern twist in its implementation on state-of-the-art parallel databases to lift the algorithm’s scalability to potentially petabyte-sized datasets.

## Acknowledgements

We would like to thank Mark Schloss, Mitko Stoyanov, and Elias Saab for setting up and maintaining the Greenplum cluster used for the experiments.

## References

- Arasu, A., Ganti, V. & Kaushik, R. (2006), Efficient exact set-similarity joins, *in* ‘Proceedings of the 32nd International Conference on Very Large Data Bases’, VLDB Endowment, pp. 918–929.
- Cands, E. J., Romberg, J. K. & Tao, T. (2006), ‘Stable signal recovery from incomplete and inaccurate



measurements’, *Communications on Pure and Applied Mathematics* **59**(8), 1207–1223.

- Chang, L., Wang, Z., Ma, T., Jian, L., Ma, L., Goldshuv, A., Lonergan, L., Cohen, J., Welton, C., Sherry, G. & Bhandarkar, M. (2014), HAWQ: A massively parallel processing SQL engine in Hadoop, in ‘Proceedings of the 2014 ACM SIGMOD International Conference on Management of Data’, ACM, New York, NY, USA, pp. 1223–1234.
- Christen, P. (2012), *Data Matching: Concepts and Techniques for Record Linkage, Entity Resolution, and Duplicate Detection*, Springer.
- Christen, P. & Belacic, D. (2005), Automated probabilistic address standardisation and verification, in ‘Australasian Data Mining Conference (AusDM05)’.
- Christen, P., Churches, T. & Hegland, M. (2004), Febrl – a parallel open source data linkage system, in ‘Proceedings of PAKDD 2004’, Springer, pp. 638–647.
- Cohen, W. W., Ravikumar, P. & Fienberg, S. E. (2003), A comparison of string metrics for matching names and records, in ‘SIGKDD’.
- Gormley, C. & Tong, Z. (2015), *Elasticsearch: The Definitive Guide*, O’Reilly Media.
- Guo, H., Zhu, H., Guo, Z., Zhang, X. & Su, Z. (2009), Address standardization with latent semantic association, in ‘Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining’, KDD ’09, ACM, New York, NY, USA, pp. 1155–1164.
- Halevy, A., Norvig, P. & Pereira, F. (2009), ‘The unreasonable effectiveness of data’, *IEEE Intelligent Systems* **24**, 8–12.
- Hellerstein, J. M., Ré, C., Schoppmann, F., Wang, D. Z., Fratkin, E., Gorajek, A., Ng, K. S., Welton, C., Feng, X., Li, K. & Kumar, A. (2012), ‘The MADlib analytics library or MAD skills, the SQL’, *PVLDB* **5**(12), 1700–1711.
- Kornacker, M., Behm, A., Bittorf, V., Bobrovitsky, T., Ching, C., Choi, A., Erickson, J., Grund, M., Hecht, D., Jacobs, M., Joshi, I., Kuff, L., Kumar, D., Leblang, A., Li, N., Pandis, I., Robinson, H., Rorke, D., Rus, S., Russell, J., Tsirogiannis, D., Wanderman-Milne, S. & Yoder, M. (2015), Impala: A modern, open-source SQL engine for Hadoop, in ‘CIDR’.
- Monge, A. & Elkan, C. (1997), ‘An efficient domain-independent algorithm for detecting approximately duplicate database records’.
- Zaharia, M., Chowdhury, M., Franklin, M. J., Shenker, S. & Stoica, I. (2010), Spark: Cluster computing with working sets, in ‘Proceedings of the 2Nd USENIX Conference on Hot Topics in Cloud Computing’, HotCloud’10, USENIX Association, Berkeley, CA, USA, pp. 10–10.