

Using the Canadian Address Register in the Labour Force Survey

Implementation, results and lessons learned

Claude Turmelle, Jean-François Rodrigue and Gavin Thompson

C. Turmelle, Statistics Canada, claude.turmelle@statcan.ca
J-F. Rodrigue, Statistics Canada, jean-francois.rodrigue@statcan.ca
G. Thompson, Statistics Canada, gavin.thompson@statcan.ca

1. Introduction

The Canadian Labour Force Survey (LFS) is a monthly survey of around 50,000 households, selected using a stratified two stage design. The primary sampling units (PSU) are geographic areas called clusters and the secondary sampling units are dwellings. The clusters are regrouped into 6 rotation groups. Each month, all the selected dwellings of one rotation group rotate out and are replaced by another set of dwellings from the same clusters. This results in each cluster remaining in the survey for several years and dwellings remaining in the survey for 6 months.

Since the LFS was first introduced, the primary source of addresses used to form the second stage sampling frame were actual listing exercises performed directly in the field for each selected PSU, something that was obviously very expensive. But since the last redesign, which was introduced in 2004/2005, a new source of addresses is now being used: the Canadian Address Register (AR). Being a list of nearly 90% of all the dwellings in Canada, the AR had the potential to reduce the cost/time associated with the creation of the frame and maybe even increase its quality.

Even though the availability of such a list opened the door to a simpler more efficient design than the two stage design used in the past for the LFS, it was decided because of time and budget constraints that the two stage design would be kept for now. Therefore, a strategy to use the AR in the best possible way, in the context of a two stage design, was developed and implemented, and this is what we present here.

In the next section, a description of the AR along with its main characteristics will be given. In the following section, the proposed (and eventually implemented) methodology to use the AR will be described. To test and refine the proposed strategy, a field test was performed in the fall 2003 and we will briefly describe this test in section 4 as well as present some of the main results. Based on the positive results of the field test, the proposed strategy was accepted and we will describe next the details of it as well as some implications on the field procedures. Now that the new design has been introduced for a few months, we were able to examine the new frame and evaluate the success of the new methodology. Some of the main results of that study will be presented in section 6. In section 7, we will discuss some of the issues encountered during the implementation along with the lessons learned. We will finally conclude and say a few words about the future of the AR in the LFS.

2. The Address Register

The AR was initially created for the 1991 Canadian Census of Population, with the purpose of improving the Census' coverage. It was first created using several administrative files, such as telephone billing files, building permits, etc... After the Census, it was updated using the lists of addresses created during the Census' enumeration process. Since then, the process has always been the same; the AR gets updated through admin files before the Census and then reconciled with and updated from the Census' lists. For a more detailed description of the creation and maintenance of the AR see Swain and al. (1992).

The AR is now a list of over 12.5 million addresses. The majority of these addresses were valid during the last Census (in 2001), while the rest come from pre-2001 addresses deemed invalid in the 2001 Census, and from updates from admin files (in preparation for Census 2006).

The AR was originally designed to be maintained for communities with a population over 50,000. But, it was expanded following the 2001 Census to include less populated regions. Because of the history of maintenance for large communities the AR is more accurate in these communities. With time it should become more accurate in less populated regions as well.

Because the Census is only taken every 5 years (years ending in ‘1’ or ‘6’, most recently 2001) the accuracy of the AR can deteriorate between the Census updates. While most of the change will be from new construction, there are also other possibilities including businesses being converted to private dwellings (and vice versa), single family dwellings being converted into apartments, dwellings being demolished, etc...

In spite of these two issues, the overall coverage rate of the AR was estimated from post-Census studies to be approximately 96% in the covered areas. On the other hand, the coverage rate is also known to vary quite a bit from region to region, something that definitely had to be taken into account when developing the methodology for using the AR within the LFS.

One very important aspect about the AR is that, in order to be a useful listing tool in a two stage design context, we must be able to assign the addresses to one and hopefully only one specific LFS cluster. There are two ways to do this. The first is by linking the AR to the National Geographic Database (NGD), which contains the street network of the entire country, along with address ranges for most street sections, also called block-faces — like the AR though, the accuracy of the NGD can vary quite a bit from region to region. For all addresses linked to the NGD, through the address ranges, we know which block-face they fall on, making it possible to assign them to a specific cluster. Such addresses are called ‘structured’. For addresses that don’t link to the NGD but were valid during the last Census we can use another geographical link, called the Census block, which is the city block that they belonged to in 2001, according to the Census enumerator. Since the LFS clusters were built using Census blocks as a base, we can then also assign those addresses to a specific cluster. But since we don’t know exactly which block-face they fall on we call those records ‘unstructured’. At the moment, about 86% of all the addresses on the AR are ‘structured’.

Note that the 2001 city block is available for all 2001 valid records but, to be consistent with the map that is prepared from the NGD, it was decided to use the block-face link first, when available. It's also important to mention that any growth record, as well as any pre-2001 addresses found to be invalid in 2001, that can't be block-face geocoded can't be assigned to a specific cluster (we don't have a 2001 city block for them) and is consequently ‘lost’ for LFS¹.

When, for a specific cluster, a preliminary list is prepared from the AR and sent to the field for verification, the addresses need to be in a certain order that ideally would optimize the listing exercise. To do that, an algorithm was developed that orders the block-faces in the cluster in such a way that it covers the whole cluster while trying to minimize the total distance travelled when verifying the list. Since the sequence is defined from block-faces, it's impossible to use it to put the unstructured records at the right place in the list. Therefore, they are added at the end of the list, sorted by street name.

It's easy to see from the process described above, that the quality of an AR list for a cluster depends as much on the NGD as it depends on the AR itself.

3. The proposed strategy

As mentioned earlier, because of time and budget constraints, the possibility of using the AR as a list frame in a one stage design context was put aside early in the development stage. Therefore, ways to incorporate its use within the traditional LFS two stage design were looked at instead. A few possibilities were considered, but in the end the most promising one was the following:

1. Based on the estimated AR quality, within each LFS cluster, the clusters are separated in three groups
 - a. In AR group 1, where the AR is expected to be of excellent quality, no listing in the field is performed and the sample of dwellings is directly selected from the AR.
 - b. In AR group 2, where the AR is expected to be of good enough quality, a preliminary list from the AR is created and then verified/updated by the field interviewer during a listing exercise.

¹ A new step is currently being implemented to be able to assign to a specific cluster a growth record that is not block-face geocoded but that appears on a street that is entirely contained within the boundaries of the cluster.

- c. In AR group 3, where the AR is absent or of poor quality, no use of the AR is attempted and a traditional listing is performed. The route determined by the sequencing algorithm appears on the map and should be followed, as in AR group 2.
2. After that, since the AR update mechanism is not run on a regular basis at this point, the list of addresses in each cluster, no matter what AR group, is maintained in the field on average once a year.

It's easy to see why this strategy was promising. Essentially, it tries to make as much use of the AR as possible while at the same time taking into account the fact that its quality varies a lot from one area to another. Although it looks very simple, this strategy requires that a measure of the current AR quality be developed, which is not a simple task. The rules that were developed and used to assign the clusters to AR groups 1, 2 and 3 will be presented in section 5.

4. The field test

In order to test the strategy presented above, a field test was performed during the fall 2003. The main goals of that test were to:

1. Develop the infrastructure needed to use the AR within the LFS
2. Evaluate the quality of the AR in general and identify its weaknesses
3. Develop and refine a method to estimate the current quality of the AR at the cluster level and eventually use it to assign clusters to AR groups
4. Test the proposed strategy and make a final decision

The test was performed in 5 Canadian cities: Vancouver, Calgary, Toronto, Ottawa and Trois-Rivières. The first four are considered large and "active" cities, in the sense that significant growth happened since the last Census, whereas the latter is a much smaller and more stable city. In each city, a stratified sample of 40 clusters (20 in Trois-Rivières) was selected. Each cluster was listed twice, by two different interviewers, using two different methods of listing — once using the group 2 method, i.e. using a preliminary list from the AR, and once using the traditional method (no use of the AR). In order to control for time and interviewer effect, the order in which the listings were done was random, and each interviewer involved in the test was to do both methods of listing, in about the same number of clusters.

After that, for each cluster, a third interviewer reconciled both lists and went back in the field to verify any discrepancies. This way of doing things allowed us to produce measures of undercoverage and overcoverage for each method of listing as well as for the preliminary AR lists.

Table 1 shows the rates of undercoverage for the three different lists for the five selected cities. It's interesting to see that the AR, even though it was over a year old when the test was performed, seems to suffer from an overall undercoverage of the same order as what the traditional method gives. The preliminary list method definitely gave the best results, which makes sense since it starts from a list that is already pretty good and improves it. We can also see that the results seem to vary from city to city. Actually, it's not presented here but the results vary even more from cluster to cluster, even within the same city.

City	AR	Pre-List Method	Traditional Method
Trois-Rivières	4.1	2.5	7.6
Ottawa	4.6	0.7	3.9
Toronto	1.3	0.5	2.2
Calgary	6.0	0.3	3.2
Vancouver	5.4	3.1	6.0
Total	3.3	1.1	3.5

Table 2 below shows the rates of overcoverage. One thing easily noticed from that table is that the lists from the AR contained a lot more invalid addresses (out of boundary, listed twice, non-residential, etc...) than the lists created with the traditional method. Even after being verified (pre-list method) there are still more invalid addresses left on the lists. This is explained by the fact that it's easier to mistakenly leave an invalid record on a list than it is to add one. Here again, the results varied quite a bit from city to city and even more from cluster to cluster.

Table 2: Overcoverage rates (%), by City			
City	AR	Pre-List Method	Traditional Method
Trois-Rivières	9.7	6.5	2.5
Ottawa	4.4	0.6	0.8
Toronto	11.3	3.6	2.3
Calgary	8.1	0.01	0.9
Vancouver	5.6	2.6	1.8
Total	8.8	2.7	1.8

From the results of Table 1 alone, it wouldn't seem so unreasonable to use the AR everywhere as a frame, since the overall undercoverage of the frame would end up being very similar to what we had before. On the other hand, by looking at Table 2 it's easy to see that by doing this we would end up with a frame that suffers from significant overcoverage, an undesirable feature. Also, we have to realize that the field test was only performed in a small sample of clusters, covering only the urban part of only 5 Canadian Cities, not randomly chosen, which makes the results not representative of the country. Still, the results from the test seemed to indicate that the proposed strategy is promising, as long as we can assign the LFS selected clusters to the right method of listing, i.e. to the right AR group.

One other thing that we looked at during the test was the amount of time needed to verify a preliminary list compare to the amount of time needed to prepare a list from scratch (the traditional method). Overall they both took about the same time. The preliminary list method took on average 15 minutes less than the traditional method.

So, from the results presented here and also from the fact that the feedback received from the field was mostly positive, it was decided to go ahead and implement the proposed strategy. The details of that implementation will be discussed next. For a detailed description of the field test and its results see Thompson and Turmelle (2004).

5. Implementation

Let's now discuss the details of the implementation. The strategy that was finally agreed on is the three AR group strategy described in section 3. The main difficulty with this strategy was to develop a method to properly assign the LFS clusters to the three AR groups. There were two competing forces at work: the desire to maximize the number of clusters not requiring checking (AR group 1) and the desire to maximize the quality of the resulting frame. Within the desire to maximize the quality of the frame there are two issues; the quality of the AR for the cluster and the ease with which the interviewers can use the AR when verifying the cluster. One recommendation made by the interviewers who participated in the field test was to minimize the number of unstructured addresses on the lists sent for verification. Remember that those addresses are not sequenced properly (they just appear at the end of the list), so having a lot of them on the list can involve a lot of driving and/or walking around to verify their status, making the listing exercise quite tedious and prone to error. All of these concerns were taken into account when developing the strategy.

The initial cluster allocation strategy was basically developed using the results from the field test as well as the results from the 2006 Census dress rehearsal that occurred in 2003-2004. From these two tests, a set of important characteristics was identified and then used to develop a series of rules that assigns clusters to the AR groups. The main characteristics used were: The number of unstructured dwellings in the cluster and how scattered they are — a hundred single family houses on 15 different streets is a lot more problematic than one big apartment building —, the AR coverage after the 2001 Census, the number of records added since the last Census, and the number of multi-unit buildings of size 2 or 3 — this category of dwellings is known to contain a fair amount of overcoverage and therefore it is safer to verify it. The basic idea to decide in which group a cluster would be put was the following:

1. If there are too many unstructured records or if they are too scattered **or** if the AR coverage in 2001 was less than 90% then put the cluster into AR group 3 (It would be too risky or too much trouble to use the AR here).
2. Otherwise, if the AR coverage in 2001 was between 98% and 102% **and** if the number of unstructured records is very low or very concentrated **and** if the amount of size 2 or 3 multi-unit buildings is small enough **and** if the number of records added since the last Census is very low, then put the cluster into AR group 1 (the idea was to put in group 1 only the low risk stable clusters).

3. Otherwise, put the cluster into AR group 2.

These rules were applied to almost all the clusters. Only for a few problematic clusters was manual intervention needed. In the end, the initial allocation distributed about 39% of the sampled clusters to AR group 1, 24% to AR group 2 and 37% to AR group 3. This is the allocation that was used in production when the listing exercises started at the end of the summer 2004.

After a couple weeks of listing we started getting some feedback from the field, and unfortunately it was quite negative. One common complaint was related to the fact that sometimes the street network shown on the cluster maps was terribly out-of-date, especially in high growth areas. Another major complaint was related to the sequencing algorithm. They were complaining that sometimes the route just didn't make sense and/or was quite far from optimal. These two issues were obviously making the listing exercise quite confusing and burdensome, which consequently was increasing the time and money spent listing those problematic clusters, on top of increasing field staff frustration towards the new methodology.

Since nothing could be done at that point to improve the quality of the street network or the optimality of the sequencing, our only way to address the problem was to modify the way things were done in the field. Therefore, for the extreme problematic cases, it was decided to allow the interviewers to do things the old way, which meant:

- For AR group 3 clusters: ignore the predetermined route that appears on the map and use your experience and judgment to come up with the most optimal route possible.
- For AR group 2 clusters: ignore the AR list and the predetermined route and start from scratch as for the problematic AR group 3 clusters.

Since the new LFS sample was phased-in one rotation group at a time, over a period of six months, the listing exercise was also done one rotation group at a time. So, when we started getting the feedback from the field, the listing of only one rotation group had started. That gave us the opportunity to go back to the set of rules used to assign the clusters to AR groups and adjust it, based on the new information, for the other rotation groups. To do that, we meticulously examined several problematic clusters that had already been listed and identified some common characteristics that had not been used that far and that could help in identifying problematic clusters. The characteristics that seemed to be linked to problematic clusters were: large increase in the length of the street network since the last Census, very large area cluster and high number of census blocks within the cluster. So, we used those characteristics to adjust the rules. In the end we transferred over three hundred clusters from AR group 1 and 2 to AR group 3.

We can find the final allocation by province, at the end of the phase-in, in Table 3.

Province	AR group 1	AR group 2	AR group 3	Total
Newfoundland & Labrador	27 (12%)	38 (17%)	157 (71%)	222
Prince Edward Island	25 (20%)	21 (17%)	80 (63%)	126
Nova Scotia	88 (26%)	59 (17%)	197 (57%)	344
New Brunswick	58 (18%)	45 (14%)	217 (68%)	320
Quebec	429 (33%)	314 (24%)	553 (43%)	1296
Ontario	935 (46%)	459 (22%)	644 (32%)	2038
Manitoba	215 (51%)	49 (11%)	162 (38%)	426
Saskatchewan	151 (35%)	80 (18%)	207 (47%)	438
Alberta	213 (38%)	125 (23%)	220 (39%)	558
British Columbia	333 (41%)	242 (30%)	232 (29%)	807
Total	2474 (38%)	1432 (22%)	2669 (40%)	6575

It's easy to see that the distribution among AR groups varies quite a bit from province to province. The AR is definitely not that good and useful in the maritime provinces (The first 4 rows), but it's much better in the other provinces, specially in Ontario, Manitoba and British Columbia where we have over 40% of the sampled clusters in AR group 1.

6. Results

Now that the listing has been done for all AR group 2 and 3 clusters and one round of listing maintenance has been done for all clusters (including AR group 1), it's possible to match the current frame back to the AR used to prepare the preliminary lists and evaluate how well our strategy worked, especially the strategy used to assign the clusters to the AR groups. Remember that the idea behind the three AR groups was to separate the clusters based on their expected AR quality, in terms of coverage and in terms of usefulness in the field. To evaluate the coverage part, we assumed that the lists that came back from the field after initial listing and maintenance represented the truth — this is obviously a very strong assumption but there was no other option that didn't involve going back in the field — and we matched them back to the AR. By doing this we were able to measure the undercoverage and overcoverage of the AR for each selected cluster and also for the frame as a whole.

Note that for the AR group 3 clusters, for which we didn't send preliminary lists to the field, the matching exercise involved a lot of manual work and it was decided to do it only for one rotation group. Also, and for obvious reasons, the results only include the clusters that were covered by the AR since the 2001 Census. The clusters not covered by the AR fall automatically in AR group 3 and are of no interest here.

One more thing to mention; when the evaluation study was performed, only 5 rotation groups out of 6 were available for analysis. The results are presented in Table 4 below.

Table 4: AR coverage rates by AR groups

Province	Undercoverage (%)			Overcoverage (%)		
	AR group 1	AR group 2	AR group 3	AR group 1	AR group 2	AR group 3
Newfoundland & Labrador	1.1	1.9	62.2	1.7	4.1	29.6
Prince Edward Island	2.3	8.7	45.3	3.0	6.9	23.3
Nova Scotia	1.4	14.2	32.5	2.4	12.6	20.7
New Brunswick	3.4	9.4	23.7	3.2	8.6	18.8
Quebec	1.7	8.1	33.8	3.5	9.1	16.9
Ontario	1.1	4.9	18.5	1.0	6.4	15.7
Manitoba	0.9	4.3	37.9	0.8	5.1	17.9
Saskatchewan	0.9	7.4	31.2	0.7	6.5	18.9
Alberta	0.4	12.9	44.0	0.8	7.9	21.8
British Columbia	2.3	5.9	27.2	2.3	6.9	17.8
Canada (without Territories)	1.4	7.1	32.2	1.7	7.4	18.5

It's clear from the table above, that the methodology developed to distribute the selected clusters into the AR groups worked pretty well overall. Even though in each AR group there were a few clusters for which the results were much worse than elsewhere, the strategy was in general successfully able to predetermine what the coverage of the AR would be and consequently assigned them to the appropriate AR group.

In some regions though, the results look so good that they could be considered suspicious. For example, in Ontario, Manitoba, Saskatchewan and Alberta the overcoverage of AR group 1 clusters is less than or equal to 1%, which is surprisingly low. Even the undercoverage in AR group 1 in those four provinces seems to be quite a bit lower than in the other provinces. It's not impossible that the strategy to assign the clusters did an exceptional job in those provinces, but it's more likely that those amazing results come from a "less meticulous" listing maintenance than elsewhere.

Remember that in AR group 1, no initial listings were performed. The clusters were visited by the field staff for the first time at the time of interview, with the instructions to completely verify/update the list at the most convenient time, within the collection period. The reason for this was to put the interviews as priority number one. Since these instructions could be interpreted a bit differently, it's possible that certain regional offices and/or interviewers decided to put a lower priority on the listing maintenance and consequently put less effort into updating the lists in AR group 1 clusters, making it look better in our analysis. The only way to verify this though would be to go back to the field to assess how well the maintenance was done.

Results by rotation group were also looked at, but only a little variation was observed from one rotation group to the other. We also compared, for AR group 1 and 2 clusters, the results between rural and urban areas and between the old AR covered areas and the 2001 AR extension areas. The results are presented in Table 5.

For AR group 1 there don't seem to be any big differences, but for AR group 2 it's obvious that the AR coverage is much better in the urban areas and in the old AR areas. This doesn't come as a surprise since the NGD and the AR are fairly new in the rural and extension areas. On the other hand, even if the AR coverage was not as good in these areas, the strategy used to determine which clusters should go into AR group 1 seems again to have worked pretty well.

Table 5: AR coverage rates by AR groups

Areas	Undercoverage (%)			Overcoverage (%)		
	AR group 1	AR group 2	Combined 1&2	AR group 1	AR group 2	Combined 1&2
Urban	1.3	6.8	3.4	1.7	7.2	3.8
Rural	2.1	10.3	6.1	1.7	9.9	5.7
Old AR	1.3	6.3	3.1	1.7	6.8	3.5
2001 Extension	1.7	9.5	5.5	1.9	9.2	5.4

7. Lessons Learned

It was shown in the previous section that coverage-wise the strategy implemented worked pretty well. Now let's take a step aside and let's look at it from a different angle. One of the main motivations to use the AR within the LFS was to reduce the total cost associated with listing, through a reduced number of clusters to be listed. The lower number of clusters to be listed was achieved with the creation of the AR group 1, but unfortunately the cost reduction didn't quite meet our expectations and there are a few reasons for that.

First, as mentioned before, the AR and the NGD are most accurate in large urban areas, which means that most of the AR group 1 clusters will be concentrated in those areas where clusters are in general smaller, more accessible and inexpensive to list. As for the other clusters, i.e. the rural/remote ones and the ones where significant growth is happening, most of them ended up in AR group 2 or 3. The consequence of that is that almost all the expensive clusters to list still had to be listed, either with or without the help of the AR.

Second, even though the field test had given very encouraging results in terms of coverage and time spent listing, generalizing these results to the rest of the country was risky. The quality of the AR and/or the NGD in the newly AR covered zones (AR extension) was never really tested yet the same strategy was implemented. For AR group 1 we saw that the strategy worked well, but for AR group 2 it resulted in many clusters having to be listed with the help of a not so good AR and/or an out-of-date map. In these cases, the use of the AR actually created more problems than helped, thus increasing the time needed to perform the listing, and the cost.

Third, it's always risky to use tools/methods that haven't fully been tested. One good example of this is the sequencing algorithm that was developed for the Census of population and used by LFS to determine the route to follow while listing a cluster and also determine the order of the addresses on the list. At the time of implementation a few flaws of this algorithm were known but the full impact of these flaws was not known. In a significant number of clusters in AR group 2 and 3 the sequencing ended up being so bad (e.g. a route that keeps jumping from one side of a highway/river to the other or a route that makes you travel erratically from one city block to the other in a very large cluster) that it created a lot of confusion in the field and a significant increase in the time and cost associated with listing.

Finally, any new procedure being implemented always involves a certain learning curve. In LFS, we luckily have experienced field interviewers that have been doing listing for quite some time and that were very familiar with the old methodology. But because of that, they were not all open and ready to change the way they work, even if it was for the better. On top of that, due to time and budget constraints it was decided that most field procedures would be developed at Head Office and that the training would consist of self-training manuals that would be distributed to the field staff just before they start using the new methodology. Obviously, it took some time for the interviewers to become comfortable with the new

procedures and at first a certain degree of confusion and frustration caused the listing exercises to take longer and consequently to cost more.

To summarize all that, here's a list of lessons that we learned along the way:

- Before implementing a new strategy/methodology, make sure to test it in as many possible situations/environments as possible. Otherwise it might mean taking some big risks.
- Always involve as much as possible the people that will actually do the work. Most of the time they are the ones with the expertise.
- Don't try to save on training; you might pay for it later.

8. Conclusion and discussion

This paper has presented the new methodology implemented during the last LFS redesign regarding the use of the AR for the creation of the sampling frame. The details of the implementation were given as well as some preliminary results on the performance of the three AR group strategy. In general the new methodology proved to be a success.

A few issues encountered during the implementation stage were described as well as their negative impact on the total listing cost. These issues were mostly related to the unreliability of the AR and the NGD in some areas of the country and also related to the fact that it always take a certain amount of time for the staff to adjust to any new procedures. But fortunately, the interviewers eventually got more comfortable with the new methodology and, after making a few adjustments to the field procedures and to the cluster allocation strategy (as explained in section 5), everything got on track and started going well.

As for the future, it was mentioned earlier that the availability of a register that covers close to 90% of the dwelling population was opening the door to a simpler more efficient sampling design, and this is something that is definitely going to be explored. As part of the 2006 Census enumeration process, the AR will once again be updated, improved and potentially extended. This, and the fact that a certain amount of experience with the use of the AR will have been acquired by then, make the implementation of a stratified single stage design something to seriously consider. The gains in sample efficiency of such a simpler design could be significant.

REFERENCES

- Swain, L., Drew, J.D., Lafrance, B. and Lance, K. (1992). The creation of a residential address register for coverage improvement in the 1991 Canadian census. *Survey Methodology*, 18, 127-141.
- Thompson, G. and Turmelle, C. (2004) Classification of address register coverage rates – a field study. In *Proceedings of the section on Survey Research Methods*, American Statistical Association, 4477-4484.