

In this document

Introduction	1
Overview	2
Sources	2
Detailed description	3
Step one: Determine areas ineligible for placement of new HU, HH and POP	3
Step Two: Determine road segments eligible for new HU, HH and POP.....	4
Step Three: Place HU and POP along new roads.....	6
Identify blocks that contain new roads indicative of new development	6
Add HU and POP counts along new roads.....	10
Step Four: Add or remove HU and POP to individual blocks to align with county-level estimates.....	11
Adding HU or POP	11
About annual county-level HU and POP figures.....	12
Removing HU or POP.....	12
Step Five: Calculate the number of HH in each block based on the number of HU in the block	14

Introduction

This document describes the methodology used by Federal Communications Commission staff to create block-level estimates for the number of housing units, households and population for mid-year 2016 (or any year for which data are available). The first section provides an overview of the methodology; the second section provides a description of the data sources staff relied on; and the third section describes the methodology in more detail, providing illustrative examples to show the effect of the code.

The results of this approach are necessarily estimates. We lack data about which houses in each county were actually destroyed, or which homes sit vacant in 2016 that were occupied in 2010. We attempt to minimize estimation errors by basing the estimates on an authoritative data source: U.S. Census Bureau county-level population estimates. And, as described below, we try to avoid over-stating the footprint of development – areas with housing units, households and/or population (e.g., to avoid creating an area unlikely to ever get broadband deployed but which we count as unserved). As with any such estimate, the larger the area of interest, the less these estimation errors matter.

Overview

The end product of this process is a table (in .csv format) with the number of housing units (HU), households (HH) and population (POP) for each census block in the country for mid-year 2011 through 2016.

The process begins with block-level counts for HU, HH and POP from the 2010 U.S. Census. Staff then take county-level estimates from the U.S. Census Bureau for HU and POP for mid-year 2016 (or any later time period) and “distribute” any increases or decreases along eligible roads as described below. This provides block-level estimates that, when summed to the county level, will align with U.S. Census county-level estimates. Staff then calculate the HH count based on the HU count in each block.

Commission staff use 2010 census blocks as a standard geography in this calculation (and, e.g., in collecting fixed broadband data on Form 477). All calculations here are made using 2010 block polygons.

There are five basic steps to determine block-level estimates:

1. Determine areas ineligible for placement of new HU, HH and POP. These areas are meant to capture, for example, parkland where new development is much less likely
2. Determine road segments (TIGER edges) eligible for new HU, HH and POP (i.e., find the roads along which new homes will be built and people will live)
3. Identify census blocks that contain with new roads indicative of new development; and add HU and POP counts along these new roads
4. Add or remove HU and POP counts to individual blocks to align with county-level estimates
5. Calculate the number of HH in each block based on the number of HU in the block

The approach of placing housing units and population along roads is consistent with work done in creating the National Broadband Map. In addition, the use of U.S. Census Bureau county-level estimates and the methods for adding and removing counts of housing units is consistent with the approach taken in developing cost models for high-cost Universal Service.¹

Sources

- OpenStreetMap (OSM) data to identify areas less likely to have new development
 - [Downloads](#) for the 50 states and DC
 - [Downloads](#) for the five territories: American Samoa, Guam, the Northern Mariana Islands, Puerto Rico and the Virgin Islands
- U.S. Census Bureau [TIGER shapefiles](#)
 - Census block shapes: 2010 census blocks
 - Road segments: Census all lines (edges)

¹ See, for example, Section Two of the [cost model methodology documentation](#).

- Additional U.S. Census data
 - 2010 block HU, HH and POP counts: [Summary File 1](#) (SF1)
 - Annual population and housing unit estimates: American Fact Finder, PEP_201x_PEPANNRES and PEP_201x_PEPANNHU (for each year from 2011 through 2016)
- Additional data source
 - Annual population estimates for territories other than Puerto Rico: [CIA World Factbook](#)

Commission staff used PostGIS (GIS-enabled Postgres) for the analysis; see, e.g., information about [OpenGeo Suite](#). Commission staff read Census shapefiles into postgis using shp2pgsql; and read osm data, in .osm.pbf format, into PostGIS using osm2pgsql.

All data sources are re-projected onto SRID 900913 to allow for geospatial analysis in a single SRID with a unit of meters.

Detailed description

Step one: Determine areas ineligible for placement of new HU, HH and POP

While Commission staff are not aware of a data set that captures areas that are off-limits to new development, there are some areas where such development is less likely. National Parks are one clear example, but there are others (e.g., airports, wetlands, nature reserves, hospital campuses, cemeteries). Other areas, like university and college campuses or military installations, may see growth in population, but are less likely to see growth in housing units and households (with people living in group quarters like dormitories).

OpenStreetMap data include a rich set of descriptions to identify such areas. Staff used the following flags and values to identify areas that we assume will not have new development in this process:

- a. leisure IN ('park','nature_reserve')
- b. boundary = 'national park'
- c. landuse IN ('forest','cemetery','golf course', 'landfill', 'military')
- d. military IS NOT NULL
- e. water IS NOT NULL or wetland IS NOT NULL or waterway='riverbank'
- f. tourism = 'zoo'
- g. aeroway = 'aerodrome' (i.e., airports)
- h. amenity IN ('university','college','hospital')
- i. natural IN ('water','wetland')

Staff create a table for each of the 56 states, territories and the District of Columbia with such areas that are treated as ineligible for development. Since OSM boundaries do not necessarily align well with U.S. Census shapes, staff buffer these polygons derived from OSM data (to avoid, e.g., placing a new home in a park by selecting a road that forms a park boundary).

The state-by-state data from OpenStreetMap do not necessarily include all bodies of water, however. Therefore, staff also exclude the area of any census block (2010 vintage) with no land area (aland10=0). This should exclude, for example, roads that cross bodies of water from having new housing units added to them.²

Staff append data for polygons for blocks without land area to the buffered “park” data to capture all the areas ineligible for new HU and POP placement.

There are examples of the impact of this approach in the next section.

Step Two: Determine road segments eligible for new HU, HH and POP

Commission staff start with the Edges (all lines) file from the U.S. Census, and extract all roads (using roadflg='Y') for each county. Next, staff uses code to identify whether each segment of road is eligible for placement of HU and POP along two dimensions: road type and whether the road falls into an area determined to be ineligible above.

For road type, staff takes two types as eligible based on the U.S. Census classification: Secondary Road; and Local Neighborhood Road, Rural Road, City Street.³ This represents the majority of roads in most areas, and excludes, for example, Primary Roads, Ramps, Alleys and Private Roads.

Next, code identifies road segments from the edges file that fall largely outside the ineligible area defined above. Excluding only roads that fall entirely in the excluded areas would mean that a road with only a tiny bit of its length outside a park could be selected for placement of a new housing unit or population which would likely fall inside the park. On the other hand, excluding any road that touched any part of a park would mean excluding too many roads. The method in place excludes any road with more than 25% of its length inside the area deemed ineligible.

See Figures 1 and 2 below for example of the areas excluded using this approach.

² There are some census blocks in the middle of bodies of water that the US Census Bureau classifies as having land area, perhaps reflecting the area associated with roads; see, e.g., census blocks 121030254051000 and 121030268091033.

³ Technically MTFCC of S1200 and S1400



Figure 1 Area north of Bethany Beach, Delaware showing roads excluded from eligibility for placement of new housing units and population in red. These ineligible roads are primarily in Delaware Seashore State Park and a National Guard facility. Base map © OpenStreetMap




Figure 2 Washington, DC showing roads excluded from eligibility for placement of new housing units and population. Roads in red are excluded based on their location – for example, near the White House, Capitol and monuments and museums; in Rock Creek Park; the National Observatory; and Howard and George Washington Universities. Roads in yellow are excluded because of the road type (e.g., highways). Eligible roads are shown as blue lines. Base map © OpenStreetMap

Step Three: Place HU and POP along new roads

Identify blocks that contain new roads indicative of new development

There are several challenges with U.S. Census Edge data that make it difficult identify new roads. First, while each road segment has a unique TLID (“Permanent edge ID”), Census assigns a new TLID when a road segment is redefined. For example, in the 2010 Edge data the segment of K Street NW in Washington, DC with addresses from 52 (lfromadd) to 99 (rtoadd) has TLID=76235443; however, in the 2015 Edge data that road segment is replaced by two others (TLID = 639337214 runs from 62 (lfromadd) to 81 (rtoadd) while TLID=639337215 extends from 76 (lfromadd) to 99 (rtoadd)).

Second, there may be new road segments in a later year that correct an omission in a prior year. See, for example, Figure 3 below, which shows a portion of K Street NW in Washington, DC overlaid with a satellite image of the area. Lines in orange represent roads in the 2010 TIGER Edge data, while blue lines represent roads in the 2015 TIGER Edge data (roads that appear in both will appear as blue). Between 2010 and 2015, the TIGER data corrected the data along K Street NW to reflect it as a divided road; but there was not new road added.



Figure 3 Area along K Street NW in Washington DC showing road lengths in 2015 (in blue) compared to 2010 (in orange). Imagery © 2016, Commonwealth of Virginia, DigitalGlobe, District of Columbia (DC GIS), Sanborn, U.S. Geological Survey, USDA Farm Service Agency

Furthermore, Census can re-classify a road that had been an ineligible road type in an earlier year as an eligible road type in a later year (see Step Two above for discussion about road types). In short, it is difficult to identify new roads from the Edge data alone.

Instead, staff calculate the amount of roads in each census block and then determine the amount of growth in eligible road types (secondary roads and local roads, as described in Step Two above) in each block over time.

The first step is to find the intersection between 2010 block boundaries and 2015 road segments for each county. Staff buffered 2010 block boundaries by 0.1 m to avoid problems where the road segments happen to cross the boundary. The next step is to find the intersection between those same blocks and the 2010 road segments for each county. Note that there can be changes to the definition of counties and their boundaries,⁴ so it is important to align the areas over time.

Next, code calculates the total road length in each block in the end year and the start year. Finally, code tests blocks to determine where the change in road length is more likely due to changes in the data set (e.g., corrections) than new road. The first test rejects blocks where changes in road length are small – 10 meters or less. The second test rejects blocks where total road length increases by less than half of

⁴ See, e.g., <https://www.census.gov/geo/reference/county-changes.html>

the change in eligible road length – this excludes blocks where most of the increase in road length comes from reclassification of existing roads to secondary or neighborhood roads. The third test rejects blocks where the total length of new road is small compared to the circumference of the block. This last test excludes, for example, where the second lane of a divided highway is added to a block.

This approach identifies blocks with new roads (see example below), while rejecting others where there do not appear to actually be new roads (including Figure 3 above). See, for example, Figures 4 and 5 below showing census block 121199114001002 in Sumter County, Florida. Between 2010 and 2015, there were over 37.5 km (over 23 miles) of eligible roads added to this block alone.

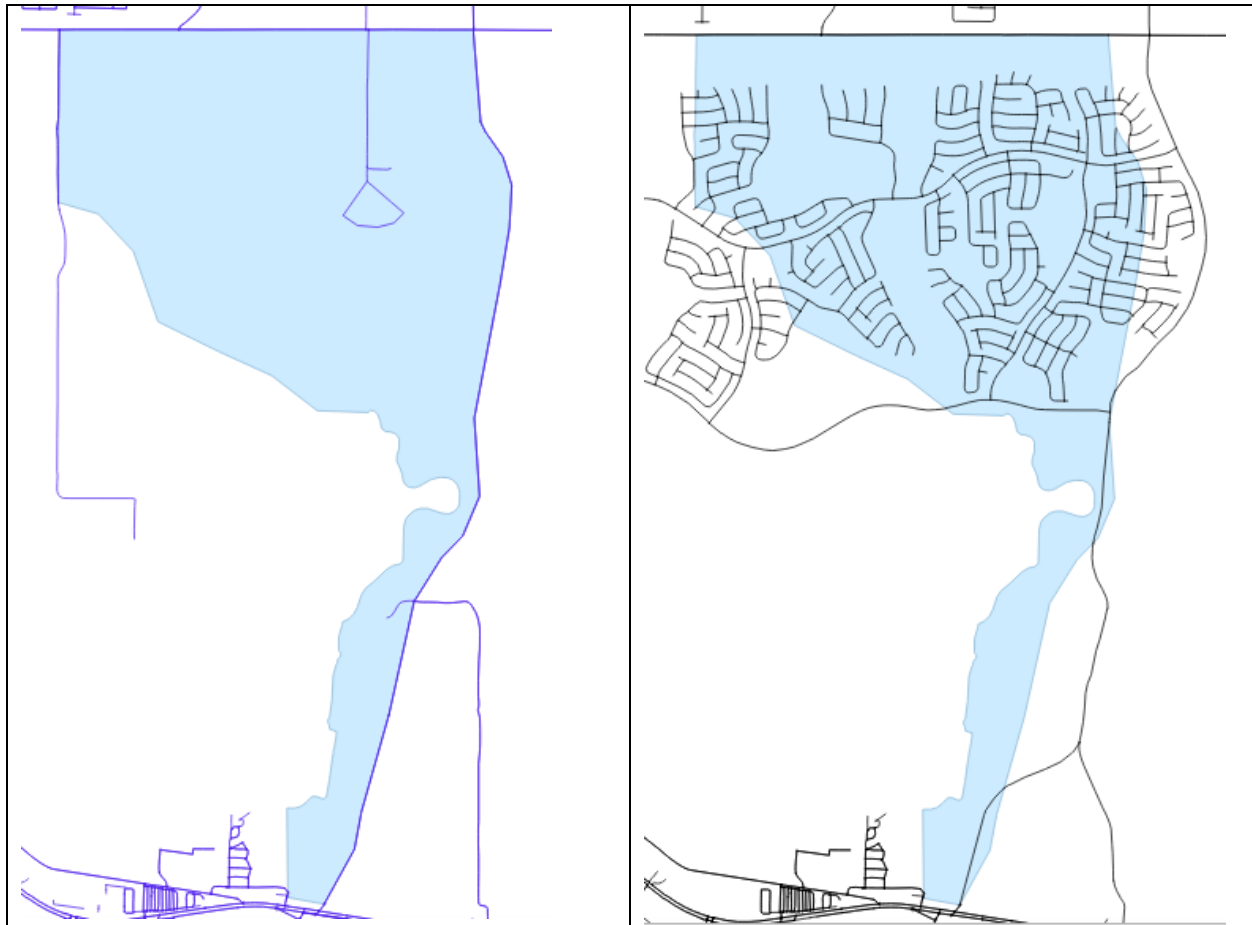


Figure 4 Census block 121199114001002 in Sumter County, Florida, shaded blue. On the left, roads from the 2010 TIGER Edge data; on the right roads from the 2015 TIGER Edge data.



Figure 5 Close up of block 121199114001002 overlaid on satellite image showing residential development. Imagery © 2016 Digital Globe, Lake County, U.S. Geological Survey

There are examples where the approach appears to miss new development. See for example, Figure 6.




Figure 6 Census block '121199114001000 in Sumter County, Florida shaded blue, showing 2015 TIGER roads (North faces left in this figure); this block is just east of the block shown in Figures 4 and 5. This block is not identified as having new roads because of the amount of new road relative to the size of the block.

However, the drawbacks of assuming the presence of housing units and population where none actually exist is significant. The Commission could, for example, treat such areas as lacking broadband; but if those areas lack development, they are unlikely to ever have broadband deployed. Thus, the approach attempts to be conservative in identifying areas with new development.

Unfortunately, even with this approach, there appear to be areas identified as having new development that reflect corrections to the underlying data.



Figure 7 Census block 121199107001123 in Sumter County, Florida shaded in blue with 2015 TIGER road data overlaid. The two roads running north-to-south across the block are new in 2015 compared to 2010 and lead this block to be identified as having new roads and development. Imagery © 2016, DigitalGlobe, US Geological Survey

Add HU and POP counts along new roads

The next step involves taking the blocks identified as having new roads, and the length of new roads in those blocks, to determine the number HU and POP to be added. Complicating this is what appears may be a lag between the TIGER road data and the decennial census – it may take longer for roads to appear in the TIGER data than it takes for housing units, households and population to appear in the decennial census. For example, block 121199112001349 as shown below in Figure 8, was reported as having 440 housing units (and a population of only 373) in the 2010 Census figures despite having very few roads according to 2010 TIGER road data (left-hand side). A more complete road network appears in 2015 TIGER road data (right-hand side).

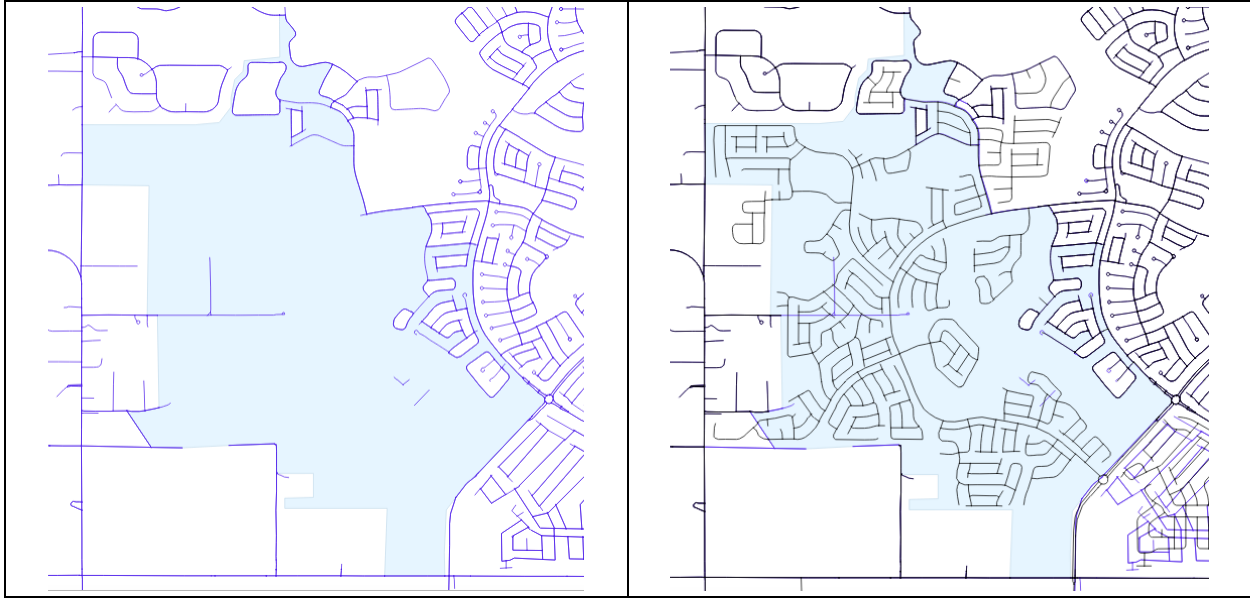


Figure 8 Census block 121199112001349 in Sumter County, Florida, shaded in blue. The left-hand side shows the roads according to 2010 TIGER data; the right hand side shows roads from 2015 TIGER data overlaid.

Therefore, the code treats blocks with a significant number of housing units or population differently than those with only a limited number of each. For blocks with fewer than five housing units or a population of ten in the 2010 census, code assumes one HU for every 30 meters (~100 feet) of new road length; and one POP for every 15 meters. For blocks with more housing units or population, the code assumes one HU for each new road segment (new TLID).

In some cases, there is very large growth in road lengths in the later-year TIGER road data (perhaps due to corrections of prior road omissions rather than new road construction) in counties with relatively few housing units and population counts overall. Allocating one HU for every 30 meters of new road and one POP for every 15 meters would mean putting close to – or on occasion more than – the entire HU and POP counts of some counties in just one or two blocks. To avoid this, the code will place no more than 10% of the 2010 HU and POP counts in any one census block due to road growth.

Step Four: Add or remove HU and POP to individual blocks to align with county-level estimates

Adding HU or POP

Where needed, the code adds HU or POP along road segments. The first step in the process is to randomly pick an eligible road segment as identified in Step Two above. Note the use here of TIGER Edge data, where a long road is split into many segment by intersecting roads, rather than road data. This increases the likelihood of placing a location in a more densely settled area because there are more road segments in areas with more intersecting roads. For each road segment identified, the code finds a point a random length along the road. Then, to ensure the selected point is in one block (not on the border between two blocks), the code shifts the point off the road by several meters; that offset is to a random side of the street in a direction perpendicular to the end points of the street segment. Finally, the code finds the census block associated with each point.

Note that it is possible for the offset point to be outside the county of interest (when the selected road segment makes up a border of the county and the offset moves the point outside rather than into the county). Due to such cases, there can be slightly fewer HU or POP added to counties than the total change in HU or POP for that county; however, such discrepancies are typically small (compared, e.g., to the difference between the 2010 census count for a county and the 2010 county-level estimate for that same county).

About annual county-level HU and POP figures

Each year, the U.S. Census Bureau publishes population estimates for each county in the 50 states, the District of Columbia and Puerto Rico; and housing unit estimates for each county in the 50 states and the District of Columbia. We estimate the change in housing units for each county in Puerto Rico as being in proportion to the change in population in that county.

For American Samoa, Guam, the Northern Mariana Island and the Virgin Islands, we take population for the territory from the CIA World Factbook; and allocate change in population in proportion to the 2010 population. We then allocate change in housing units in proportion to population. This avoid allocating housing unit or population to counties (or county equivalents) that have no population in 2010.

As noted above, county definitions can change. Each year's county estimates reflect the then-current definition of county boundaries. For example, three counties in Alaska had their boundaries change between 2010 and 2015; U.S. Census Bureau estimates in 2015 reflect the 2015 definition of those counties in reporting the values for 2010. Thus there's a significant difference between the 2010 census figures and the 2010 estimates.

In order to account for changes to county definitions, this approach uses the difference between the end-year estimates and the start-year estimates, rather than the difference between end-year estimates and 2010 census figures. As noted above, however, this approach relies on block-level figures from the 2010 census. *Therefore, because of difference between 2010 census figures and 2010 estimates (even for counties that have not had a change in their definition), the county-level totals will typically not agree perfectly with U.S. Census county-level estimates.*

Removing HU or POP

Removing HU or POP randomly could lead to removing the only housing unit or person from a given census block, effectively making that block no longer relevant for at least some policy-making purposes. Rather than risk dropping blocks in such a way, and recognizing that blocks with greater HU and POP

counts are more likely to have reductions, staff remove HU or POP from blocks with higher starting counts.

The approach removes HU or POP from the block with the largest initial count. That process continues until the reduction target is met or until the block that had the highest HU or POP count has the same number as the second-highest. The process then removes HU or POP from the highest two until the target is met or until those two blocks have the same count as the third highest, and so on.

An illustrative example will highlight the method taken. Imagine a county consisting of five census blocks, each represented by a vertical line on the left-hand side of Figure 9 below. Each blue dot represents a housing unit in that block. If the count of housing units needs to be reduced by three, the three red dots shown in the right-hand side of Figure 9 would be removed (i.e., the block with the most housing units would have the count in that block reduced by three). To remove four to seven HU, the green dots would be removed; and to remove eight to ten HU the orange dots would be removed




Figure 9 Example to illustrate how HU or POP are removed

The identification of new roads and placement of new HU and POP along those roads in steps two and three occurs independent of the net change in HU and POP in each county. The number of HU and POP added or removed in step four reflects any such additions

Step Five: Calculate the number of HH in each block based on the number of HU in the block

Since the U.S. Census Bureau does not release county-level estimates of households annually, Commission staff estimate the number of households in each block based on the number of housing units. In blocks where there are housing units in the 2010 Census, the number of households in the end year is derived by taking the number of housing units estimated in the block in the end year and multiplying by the ratio of households to housing units for that block. For blocks where there are no households in 2010, the code uses the ratio of households to housing units for the county in the 2010 census instead.

This methodology produces household counts that may not be whole numbers in each block. Over larger geographic areas, like counties, however, staff believe this approach will provide a better estimate for households than rounding the household count to whole numbers in each block. Typically staff will round to whole numbers of housing units at the tract or county level.