WiFi is Everywhere in California

Field Survey of Installed WiFi in Three California Cities

Ken Biba
Managing Director and CTO
Novarum, Inc.

Half of Internet traffic is conducted on WiFi, and for residential Internet likely more. This study attempts to measure the extent and quality of deployed residential WiFi by drive measurement and analysis of public WiFi beacons in three California cities in late 2018.

WiFi is Everywhere	There is at least one WiFi gateway per household in all three cities. In San Francisco, there is almost one per person. In San Francisco, there is one carrier WiFi hotspot every 35m of street.
Much of WiFi is very old, and likely rather bad	In El Centro and Merced, an estimated over 1/3 (in El Centro 59%!) of the installed WiFi gateways could be over a decade and three generations of WiFi old. These gateways are likely serving legacy WAN connections that are no longer considered broadband and have increasingly poor compatibility with modern WiFi user devices. These old gateways will also offer poor security choices.
Most of WiFi security is nonexistent or inadequate	An estimated 56% of WiFi gateways (Merced - 46% for San Francisco, 13% for El Centro) have either no or poor (WEP or WPA) security configured. The likely use of default and poorly chosen passwords compounds this WiFi security risk.

2 Ghz spectrum crowded, 5 GHz spectrum "empty"	In every city, all the spectrum in the 2 GHz band was occupied with WiFi networks. Only about 25% of the spectrum was occupied with WiFi networks in the 5 GHz band.

WiFi is important. Cisco[1] tells us that about half of projected 2020 Internet traffic begins or ends on WiFi - compared to 17% via cellular and 33% via a wire[2]. Most of mobile device traffic - 60% - is offloaded to WiFi rather than cellular. WiFi traffic from mobile devices has the highest growth rate of 57% CAGR substantially outpacing the overall Internet traffic growth rate of 22% CAGR and even the growth rate of mobile data - 51%. Wired traffic is growing at a comparatively modest 12% CAGR.

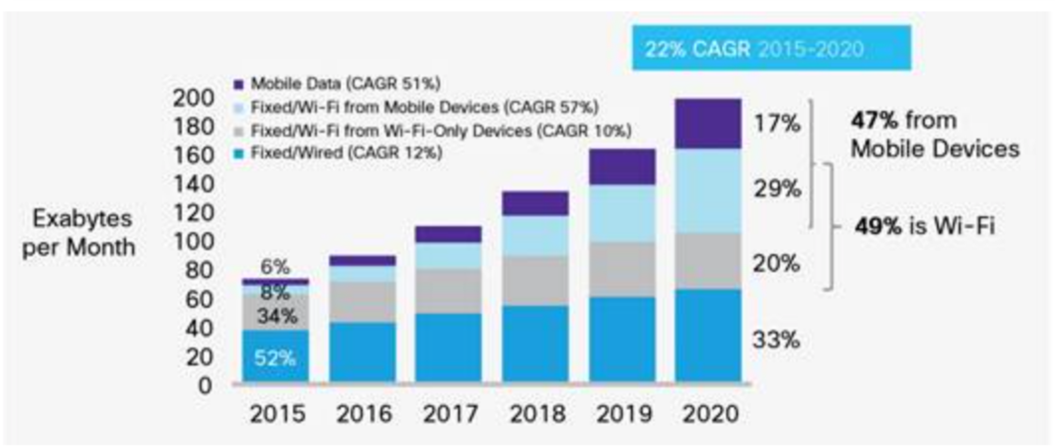

Our own lived experience tells us that all of our smartphones, tablets, smart TVs, and laptops connect to the Internet via - WiFi - most of the time. And at home , just about **<u>all</u>** of the time. After all - how many of us still use a desktop directly connected to Ethernet? Or … dial up via a phone jack? Or even own a laptop with an explicit, much less used, Ethernet connection?

Learning about the residential use of WiFi would tell us much about the residential use of the Internet. And all existing studies of WiFi usage, that I could find, were response surveys - rather than measured ground truth.

So I surveyed the residential WiFi (and incidentally the carrier public WiFi) in three California cities ranging from the high tech centrum of San Francisco to the more rural cities of Merced and El Centro. In each, I picked a representative residential neighborhood of single family residences and surveyed what WiFi access points could be found. Each WiFi access point, by design, publicly announces its presence via WiFi beacons. A survey by driving around with WiFi equipment to listen for these WiFi beacons gives on the ground evidence for the extent of WiFi deployment. And by inference, the wired equipment attached to those WiFi routers.

After collecting beacons, we analyzed them by beacon information to figure out how many there were and what purpose they were serving in those households. While not utterly precise, the

---

[1] "Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update, 2016–2021 White Paper", https://www.cisco.com/c/en/us/solutions/collateral/service-provider/visual-networking-index-vni/mobile-white-paper-c11-520862.html, February 2017

[2] Presumably that wire also includes servers and enterprise desktops.

resulting numbers are a good estimate of the population of WiFi access points in each of these neighborhoods. When extrapolated to the entire population of each city, we get a profile of what the WiFi access point installed base looks like in these three cities. Appendix A describes the methodology.

	San Francisco	El Centro	Merced
City Area (km^2)	121	29	60
City Population	884,363	42,598	78,959
City Households	356,797	13,108	24,899
Median Household Income	$77,734	$49,095	$47,739
Survey Area (km^2)	2.5	0.9	0.26
People Density (/km^2)	7281	2658	8977
Household Density (/km^2)	2938	818	2831
People/household	2.5	3.2	3.2

WiFi is Everywhere

WiFi is everywhere in our three cities - and by inference many other cities. In all three cities there is at least one WiFI gateway per household. It is only a modest exaggeration that in San Francisco and El Centro, there is almost one gateway per person.

	San Francisco	El Centro	Merced
Gateway APs/person	0.73	0.77	0.40
Gateway APs/household	1.80	2.50	1.28
Total APs/Person	0.84	0.78	0.47
Total APs/Household	2.08	2.55	1.48

There is more that one type of WiFi access point. Some are residential routers (with some provided by carriers and others provided by users from third party vendors), some are carrier public gateways, some serve for entertainment devices and some provide wireless access to printers. While Internet gateways are about 90% of the total population, these other applications are likely to grow. In the chart below, Internet gateways come from two sources - carriers and from third party consumers themselves[3]. The distribution between the two sources of gateways (consumers and carriers) is quite different in each city - with San Francisco dominated by consumer supplied gateways and Merced by carrier supplied gateways.

3 Attached to modem for the carrier of choice.

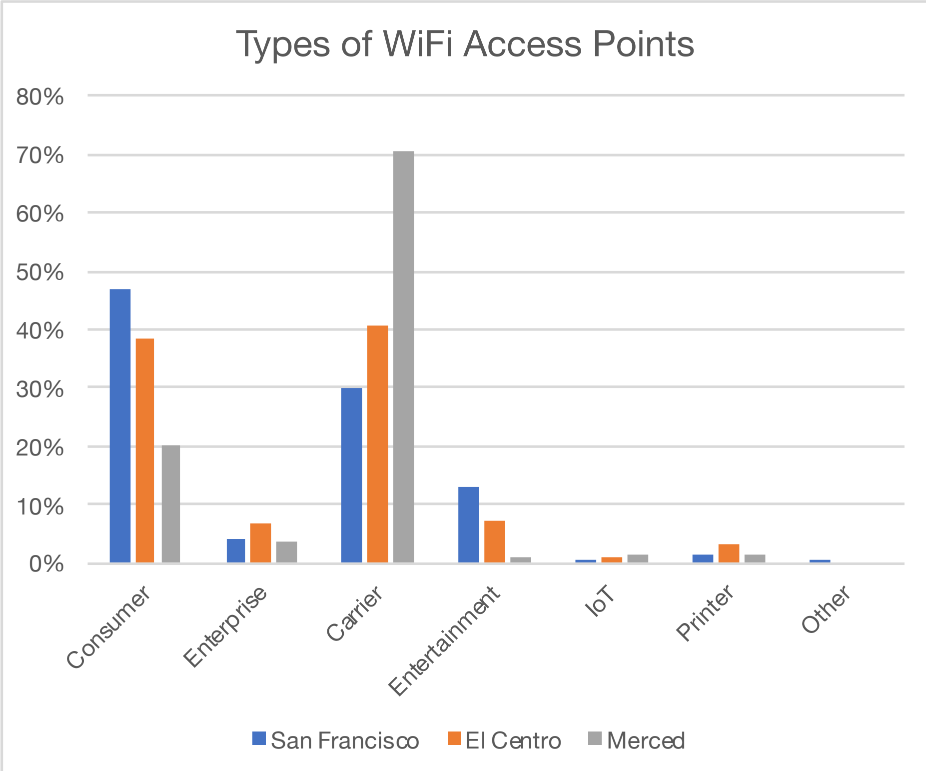

The above total carrier numbers include the public "hotspot" gateways of the carriers in each city offering public Internet access.   In each city this is is the cable Internet provider.

Scaling from each city sample to the entire city, is a lot of access points.

	San Francisco	El Centro	Merced
Speculative Total AP/city	740,372	33,420	36,943
City Population	884,363	42,598	78,959

A Mix of Old and New Technology

The installed base of access points, in the 2 GHz band strongly favors the 802.11n standard, in 20 MHz channels, on channels 1,6, and 11 in all three cities.

The installed base of access points, in the 5 GHz band strongly favors the 802.11ac standard in all three cities.   In San Francisco and Merced, there is under 10% use of 802.11n gateways at 5 GHz, and the predominant channel bandwidth is 80 MHz. El Centro differs, with a significant minority (30%) of gateways are using 802.1n on 40 MHz channels.   This would suggest that the dual band gateway population is older, since 802.11n dual band gateways are largely no longer sold.

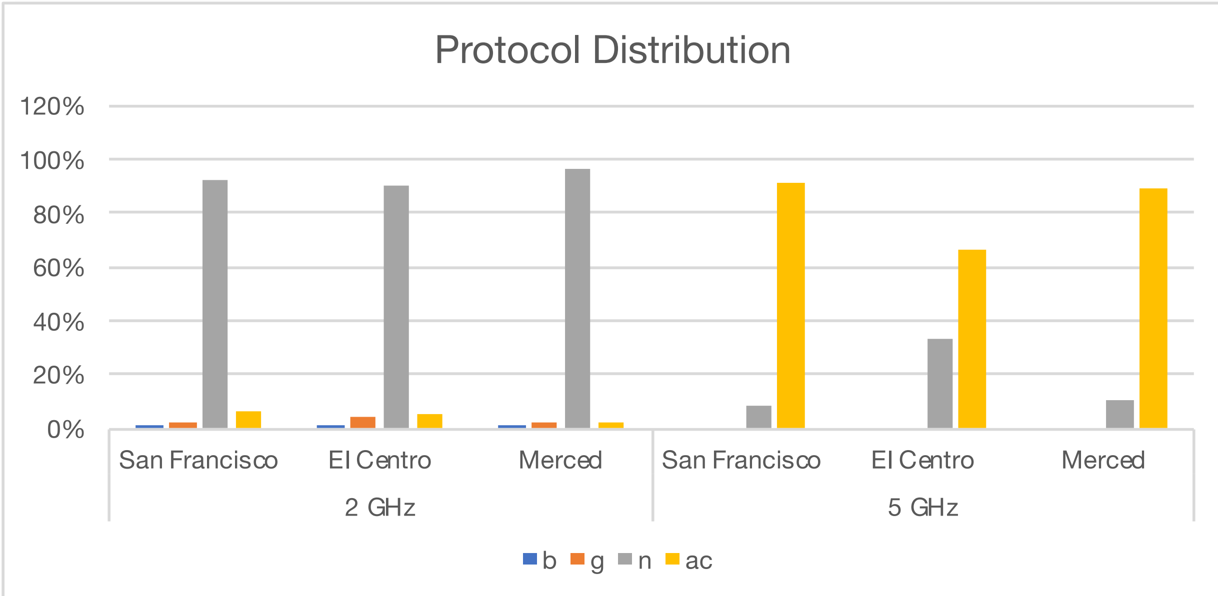

The distribution of channel bandwidths strongly favors 20 MHz channels in 2 GHz and a mix of channel bandwidths at 5 GHz.

20 MHz channels, at channels 1, 6 and 11, for the 2 GHz band is particularly good because minimizes co-channel interference. This minimizes interference from neighboring access points and increases the capacity of the band. It appears that only a small minority of users are configured to use channels 2,3,4,5,7,8,9,10 in the 2 GHz band - channels that overlap and increase interference and decrease WiFi capacity in the 2 GHz band.

San Francisco

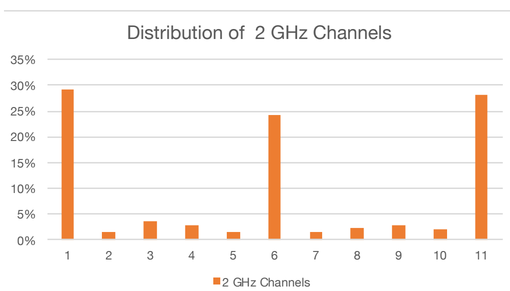

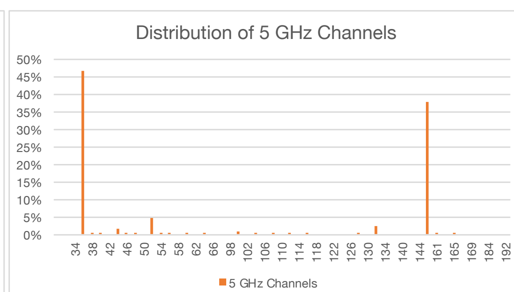

El Centro

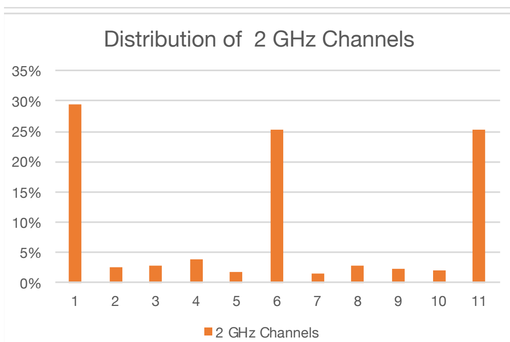

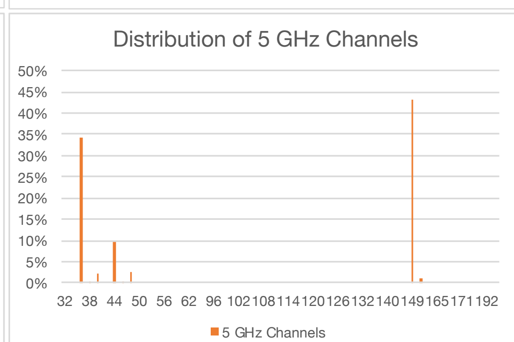

Merced

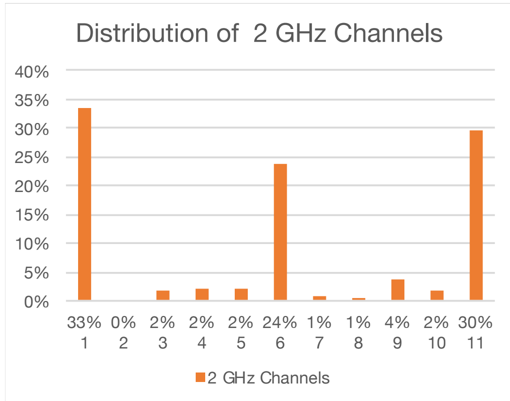

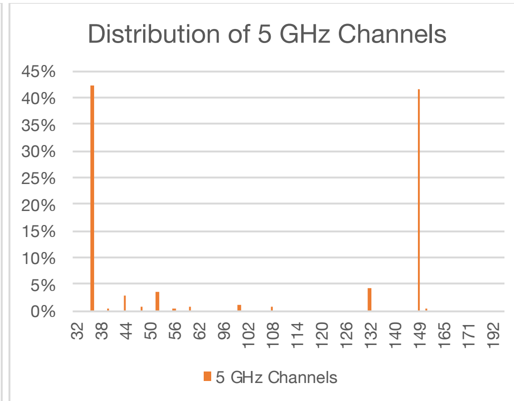

5 GHz channel usage is interesting. Older 5 GHz technology, like 802.11n, only support 20 and 40 MHz wide channels. Modern 802.11ac supports 20, 40, 80 and 160 MHz channels. In all three cities, 80 MHz channels are the majority, but El Centro shows a strong minority use of 3x3 MIMO 802.11n in 40 MHz channels (and in 20 MHz channels at 2 GHz).

The 5 GHz band in each city is extraordinarily lightly used. Over 80% of the channel usage in all cities lies on two channels - 36 and 149 - using only 160 MHz of the 685 MHz available in the US. The vast majority of the band - over 77% of the capacity is rarely used by residential WiFi. In these neighborhoods, there was small evidence of enterprise WiFi networks that might have more extensive use of the 5 GHz spectrum.

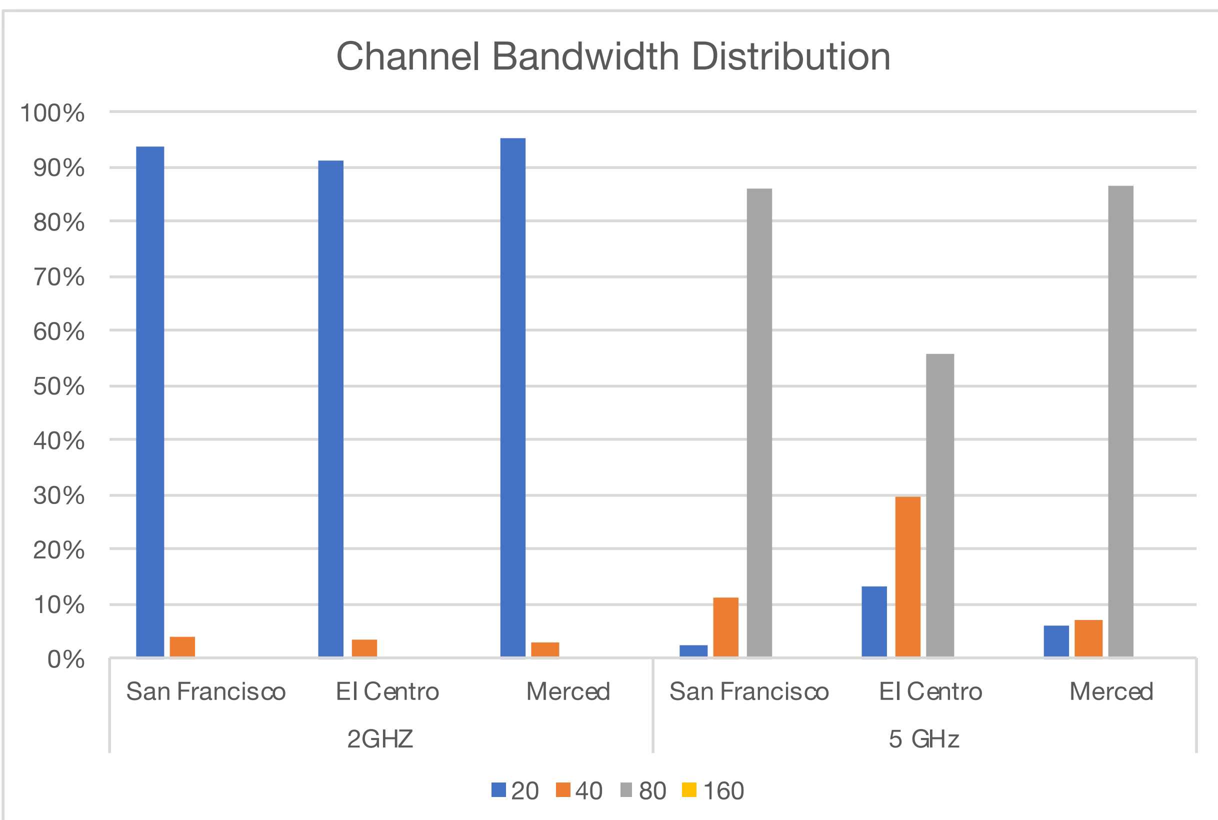

The dominant access point WiFi configuration at 5 GHz is 3x3 MIMO 802.11ac in 80 MHz channels in all three cities. San Francisco and Merced have a strong secondary installed base of 4x4 MIMO 802.11ac in 80 MHz channels. El Centro's secondary dual band installed base is 3x3 MIMO 802.11n in 40 MHz channels - an older technology.

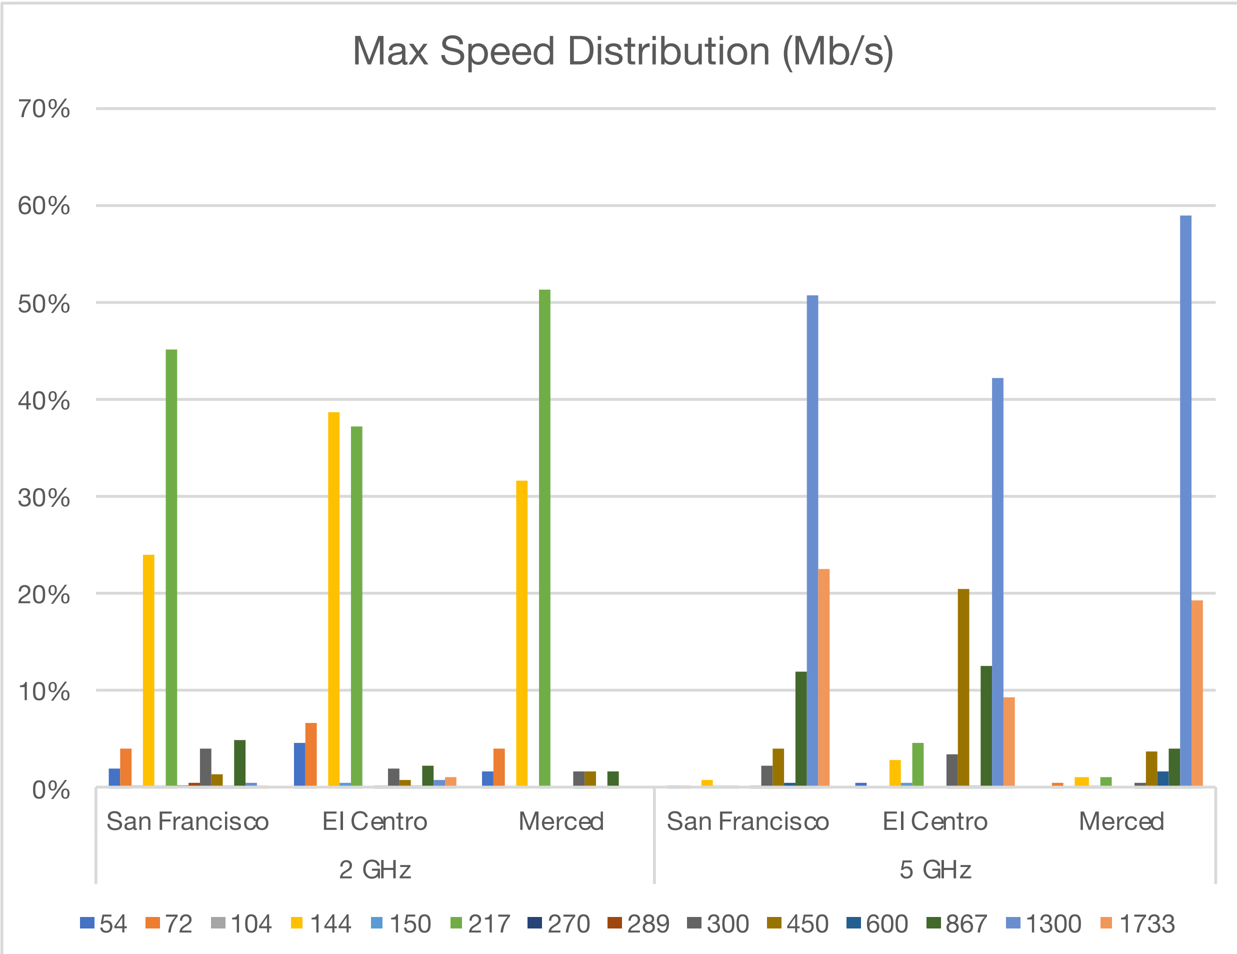

Looking at the type of modulation, we can make an estimate of the generation of WiFi included in a gateway. Very early (pre-standard, so-called "draft n" access points (likely deployed circa 2007) 802.11n gateways used single band 2 GHz operation (no 5 GHz service) with very limited use of MIMO (no more than 2x2 in 20 MHz channels). These were often deployed with carrier deployed DSL gateways where the limited WiFi throughput matched the limited throughput of the carrier connection.

First, let's calculate the population of single and dual band gateways. Then calculate the population of old (~2007 802.11) single band gateways, and the distribution of "oldish" (~2009) dual band 802.11n gateways. The combination of these gives an

	San Francisco	El Centro	Merced
% single band	24%	52%	80%
% dual band	76%	48%	20%
% single band "old"	45%	72%	43%
% dual band "oldish"	8%	45%	9%
% total gateway "obsolete"	17%	59%	36%
% dual band "newish"	65%	31%	16%

estimate of the percentage of gateways that can be considered as legacy. While San Francisco has a modest percentage, Merced has perhaps a full third of its single and dual band gateways as old and and El Centro has an impressive 59% of its installed base of gateways that are … elderly.

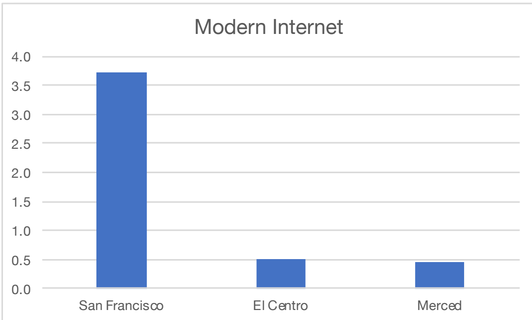

Gateways that support 802.11ac were all installed in the last 4 years and can be thought of as “newish”. They are all dual band and highly compatible with smartphones, laptops and tablets deployed in the same time period.   They are configured with 80 MHz 5 GHz channels and overwhelming support over 1 Gb/s link modulation rates from user devices. The ratio of “newish” (< 4 years old) decade old gateways to “obsolete” (~ 10 year old) gateways can be thought of a metric of the rate of turnover in gateways from old to new - a metric of how modern the Internet infrastructure is.

While all three cities have a high density of residential WiFi (and therefore residential WAN Internet connections) there is huge difference between the age of the deployed infrastructure.   Both El Centro and Merced are much, much older than San Francisco.   This study does not tell us, directly, about the age and capabilities of the WAN connections to each residence, but the age of the WiFi gateways is suggestive of older, slower WAN connections.

Why Old WiFi is Obsolete

WiFi is changing fast - as a technology that is just over 20 years old.   For example, here is a timeline[4] of technology in Apple’s now discontinued Airport Extreme line of WiFi Internet gateways - arguably, for its time, a high end WiFi access point.

Year of Introduction	Standard
1999	802.11b
2003	802.11 b/g
2007	802.11 a/b/g/n 3x3 MIMO dual band (non-simultaneous band operation)
2009	802.11 a/b/g/n 3x3 MIMO dual band (simultaneous band operation)
2013	802.11 a/b/g/n/ac 3x3 MIMO dual band (simultaneous band operation) (Phase I)
2015	802.11 a/b/g/n/ac 3x3 MU-MIMO dual band (simultaneous band operation) (Phase II)

4 https://en.wikipedia.org/wiki/AirPort_Extreme

In early 2019 the default standard technology is 802.11 a/b/g/n/ac 3x3 MU-MIMO (Phase II) dual band and pre-standard next generation 802.11ax “WiFI 6” gateways are already available.

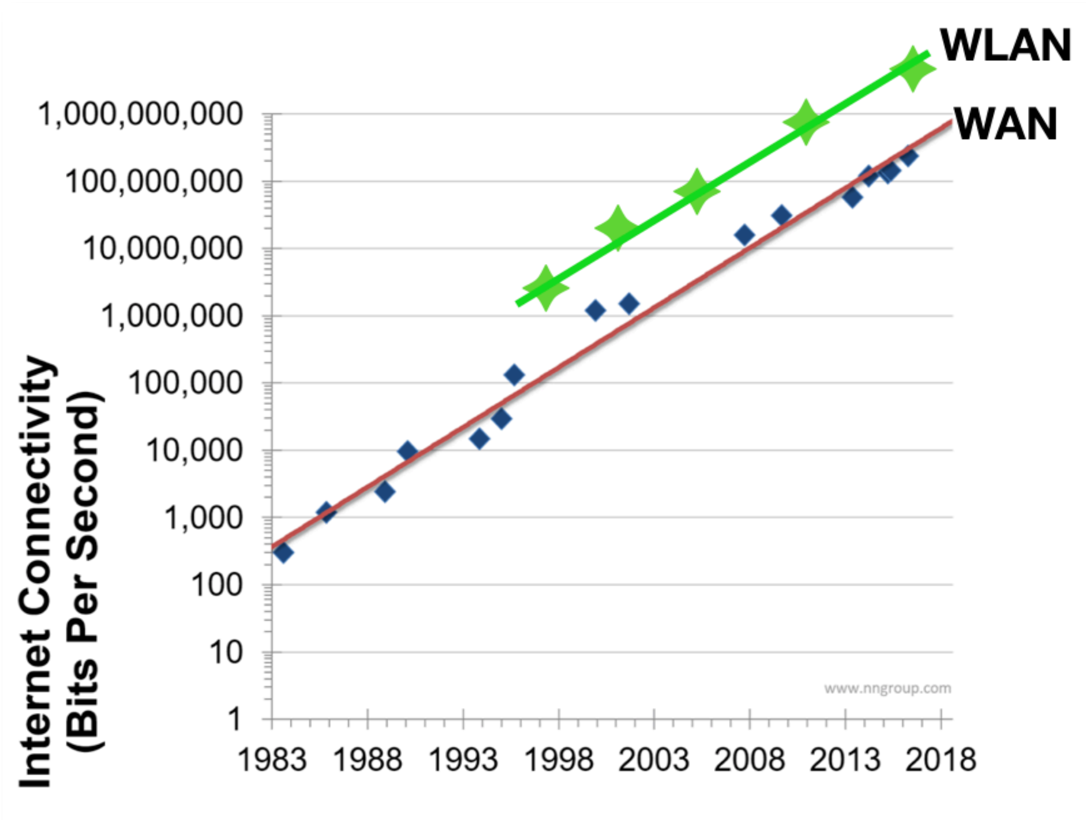

Tracking the growth in performance of both WAN connection speed (Nielson’s Law[5]) and WiFi connection speed (Biba’s Law) shows growth at 50% CAGR since … the Internet. Internet service and the WiFi gateway installed 10 years ago … are old - and comparatively slow. Are they obsolete in the face of 20x change in both carrier connection performance as well as WiFi performance?

So the questions are - how many of these old gateways are there? And do these gateways give their users a less than perfect Internet experience to their modern user devices? Modern user devices that are much younger (and faster) than the gateways and are highly WiFi intensive.

The data from our three cities suggest there is a substantial legacy of these old WiFi gateways - ranging from 17% (San Francisco) to 59% (El Centro) of the installed base. There is a speculative hint (that only more survey work could improve) that more rural, smaller cities like El Centro and Merced have a materially greater legacy of old gateways and WiFi and a slower turnover of modern gateways (and likely Internet connections).

---

[5] Best available WAN connection speeds grow at ~ 50% CAGR. https://www.nngroup.com/articles/law-of-bandwidth/

What about the quality of the Internet service? The data suggests two possible concerns.

1. WiFi gateways are generally purchased and installed when the WAN connection is installed or upgraded. Old WiFi is suggestive of an old WAN connection installed at the same time. 10 year old WiFi is 100-1000x slower than the best WiFi available today. Similarly, a 10 year old WAN connection is likely similarly slow - rather than a 1000 Mb/s cable or fiber connection - it is likely a legacy 6 Mb/s DSL connection or a 10 Mb/s cable connection.
2. While the WiFi Alliance tests for compatibility of new WiFi devices with legacy versions of WiFi, the criteria for acceptance are modest. There are no assurances of high performance of new devices with old technology access points. There is a risk of increasingly poor WiFi compatibility of old gateways and modern user devices.

The California Public Utility Commission has measured mobile broadband in California from 2012-2017[6] and is now engaged in measuring wired and WiFi broadband with the CalSPEED Home project. CalSPEED Home uses a residential measurement tool, a miniature Linux based computer, to measure both residential and WiFi residential broadband service[7]. While still early days, beta units[8] encountered an interesting problem in the field - difficulty with connecting to WiFi service in some residences - averaging connection failure rates of 30% up to 60% of all attempts - while Ethernet attempts at the same locations saw failure rates of 1-3% . These residences had a commonality - old single band 2 GHz WiFi routers from the beginning, pre-standard 802.11n days of over a decade ago. With seemingly poor backward compatibility with the modern dual-band 802.11ac 2x2 MIMO WiFi adapter in the CalSPEED Home measurement tool. These are early results from a small sample size in one California city.

But.

These early results are suggestive that the WiFi configured with modern user devices (usually 802.11ac but increasingly 802.11ax) will have difficulty with legacy single band 802.11n and perhaps dual band 802.11n gateways.

Security

WiFi security is provided by link level encryption regime with a shared encryption key. Highest security comes from both using the right encryption technology and a strong network shared encryption key. WiFi has had multiple generations of encryption technologies, as deficiencies are found and improved technology fixes them. The standard encryption regime for residential WiFi

---

6 Ken Biba, "CalSPEED Mobile: Final Report - Mobile Broadband Measurement in California 2012-2017", Novarum, January 2019.

7 CalSPEED Home's measurement instrument has both a 1000BaseT Ethernet connection and an 802.11ac 2x2 MIMO WiFi connection. In a controlled environment, cooked TCP throughputs of over 900 Mb/s on Ethernet and over 300 Mb/s on WiFi have been baselined. These are representative of modern laptop, smartphone and tablet network connections.

8 Ken Biba, "CalSPEED Home: Wired and WiFi Measurements - January 2019", Novarum, January 2019

today is WPA2. Older regimes, WEP and WPA, have known security vulnerabilities and should not be used.

The quality of the encryption key cannot be evaluated by this study, but sadly often users use keys are either manufacturer default or easily guessed. So even the best encryption technology can be defeated by poor choice of network encryption key. Poor choice of encryption key compounds the security risk of using a weak encryption regime.

Poorly Configured Security

	San Francisco	El Centro	Merced
Single Band	10%	7%	42%
Dual Band	36%	6%	14%
Total	46%	13%	56%

No Configured Security

	San Francisco	El Centro	Merced
Single Band	4%	3%	13%
Dual Band	15%	1%	3%
Total	19%	4%	16%

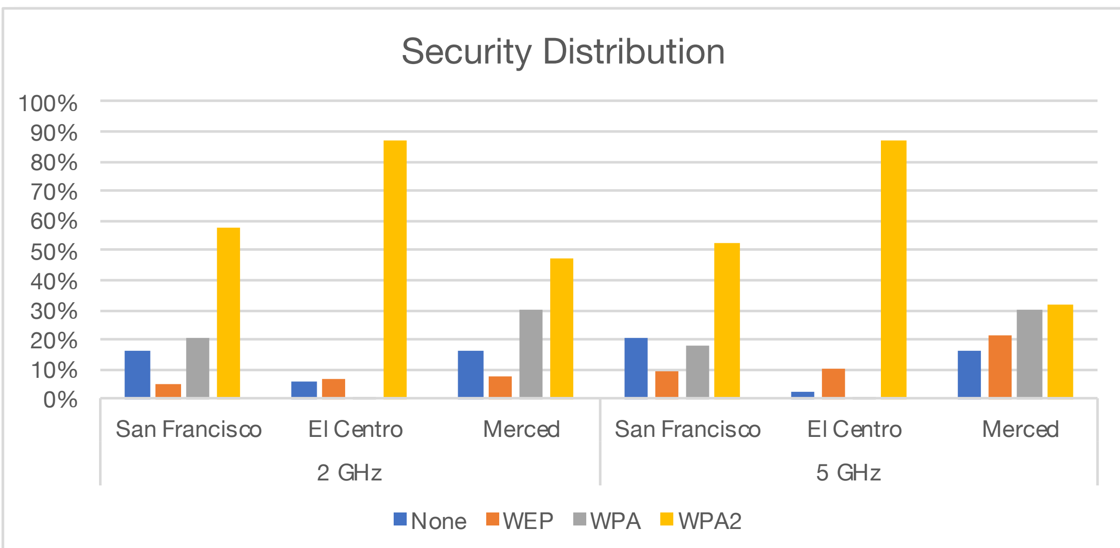

What we can evaluate is the quality of the encryption regime given a strong encryption key[9]. Sadly, many users have chosen to use no encryption or legacy encryption standards such as WEP and WPA that are weak with known vulnerabilities.

El Centro had the highest use of the strongest deployed security regime, WPA2, with 87% of access points making use of this strong encryption.

San Francisco was next strongest, with 54% of users making use of WPA2, but 46% of San Francisco users either chose no security or a weak encryption systems. In fact, 19% of San Francisco users chose to use no encryption regime and run open networks.

[9] https://www.netspotapp.com/wifi-encryption-and-security.html

Merced had the poorest Internet security, with only 44% of users choosing a strong security regime and 66% choosing no or weak security.   Merced also has the highest population of carrier provided gateways.

Obsolete gateways contribute to poor security since older gateways are either not capable of being upgraded or thru (user and/or vendor) neglect not upgraded to the WPA2 standard.

The next generation of WiFi gateways based on IEEE 802.11ax and security regime WPA3 will begin to be deployed in 2019.

Entertainment, IoT and Printers

In addition to WiFi access points used for both residential and public hotspot Internet gateways, a subset of access points provide specialized access for entertainment (particularly streaming video), management of Internet of Things infrastructure and wireless network printer access.

This analysis suggests there are only a small percentage of access points dedicated to IoT applications in all three cities - ~1%.   Similarly for wireless printer access - under 3%.

By far the largest application of these non-gateway access points is for entertainment.   I estimate that 13% of the observed access points in San Francisco are for video streaming, perhaps 7% in El Centro but only 1% in Merced.

Carrier WiFi

Some carrier provided Internet gateways provide public hotspot access.   A diagnostic for an estimate of the number of these gateways are gateways with both open security and the SSIDs typically used by these gateways.

	San Francisco	El Centro	Merced
Carrier AP/km^2	802	41	585
Distance Between AP (m)	35	156	41

With a number, an estimate of hotspot density can be simply made.
Both San Francisco and Merced have a very high density of carrier hotspot gateways - about every 40m in both cities.

City Profiles

Each city has a unique residential WiFi profile.

San Francisco	• Very high residential gateway density - ~ 1.8/household • Very high carrier public gateway density - ~ 1/35 meters • Modern WiFi gateway technology - IEEE 802.11ac, dual band, second wave as the dominant residential gateway • Emerging entertainment WiFi - ~.2/household
El Centro	• Very high gateway residential density - ~2.5/household • Low carrier public gateway density - ~1/156 meters • Older WiFi gateway technology installed base, much more legacy 802.11n
Merced	• Lowest residential gateway density - ~1.3/household • High carrier public gateway density - ~1/41 meters • The oldest installed base of WiFi gateways • Poorest security, highest population of carrier supplied APs

	San Francisco	El Centro	Merced
City Area (km^2)	121	29	60
City Population	884,363	42,598	78,959
City Households	356,797	13,108	24,899
Median Household Income	$77,734	$49,095	$47,739
Survey Area (km^2)	2.5	0.9	0.26
People Density (/km^2)	7281	2658	8977
Household Density (/km^2)	2938	818	2831
People/household	2.5	3.2	3.2

Appendix A: Methodology

Every WiFi access point broadcast information frames identifying the capabilities of the access point. This study captures those public frames by performing a car based WiFi survey driving (at slow < 15 mph) through selected residential neighborhoods in the selected cities with radio instrumentation to collect this public WiFi information. Visiwave with an internal Intel 2x2 MIMO 802.11ac WiFi adapter was used as the survey tool. San Francisco, El Centro and Merced were the selected California towns. Visiwave both listens to broadcast WiFi beacons and transmits probe requests frames to solicit probe response frames at both 2 GHz and 5 GHz emitted by surrounding WiFi access points by scanning all the WiFi frequency channels in both bands. It captures these beacons and probe response frames.

Both beacons and probe response frames contain key information about the surrounding WiFi networks:

MAC address	Uniquely identifies each radio (2 GHz and 5 GHz) in each access point, these MAC addresses also contain the manufacturer ID. The manufacturer ID is quite important since it is highly diagnostic in identifying the type of access point - some manufacturers only make IoT devices, some only printers, some only consumer retail-sold gateways, some enterprise class access points and some gateways only sold as OEM thru carriers. Almost all dual radio (2 GHz and 5 GHz) access points assign closely numbered MAC addresses to the two radios.

SSID	Symbolic name assigned either by a carrier or by a user identifying the WiFi network. Many carrier supplied access points acting as gateways have unique carrier supplied names, some in residences have been changed by the user to a different name. Third party access points have either default manufacture supplied SSIDs or user supplied. Dual radio access points will often have the same or closely related SSIDs. An access point can be configured to not advertise its presence by broadcasting a beacon, but it must respond to a probe request that broadcasts the same information as a probe response frame.
Frequency channel	The channel number/frequency of operation. Once configured, the access point will operate on that channel, broadcasting beacons advertising service until reconfigured.
Channel width	The frequency width of the WiFi channel: 20, 40, 80 or 160 MHz. The supported channel widths are an identifier for the generation of WiFi. IEEE 802.11b/g/a support only 20 MHz channels in both bands. 802.11n supports 20 and 40 MHz channels in both bands. 802.11ac supports 20, 40, 80 and 160 MHz channels in both bands.
Maximum data rate	The maximum link level data rate supported by the access point on this channel. From the data rate, the number of MIMO antennas, the specific version of 802.11 protocol and channel width can be inferred.

Speed	# antennas	Channel Width	Protocol
72	1	20	n
144	2	20	n
217	3	20	n
300	2	40	n
450	2	40	n
867	2	80	ac
1300	3	80	ac
1733	4	80	ac

WiFi protocol	The highest level of WiFi protocol supported by the access point. It is assumed that all earlier versions of WiFi are supported as well. 802.11b, 802.11g, 802.11n, 802.11ac, 802.11a
WiFi security	The link level security regime: none, WEP, WPA, WPA2.

Beacons are generally repeated every 100 msec and are transmitted at the slowest compatible speed for that frequency band[10]. Some beacons will be not be received due to collisions or lack of propagation for 5 GHz beacons.

Once captured, a partially automated and manually assisted process was used to convert beacon information into a database of observed access points, first recovering lost beacons and then identifying the type of access point. The process is not 100% perfect but does give quite high fidelity.

We “recover” or compensate for missing beacons to identify single band and dual band access points using these rules:

- Single band APs that are only 2.4 GHz 802.11n at 72 and 144 Mb/s maximum are either early, legacy first wave 802.11n gateways or a very recent IoT or entertainment access points. Each of these use minimal MIMO. The SSID usually disambiguates.
- Dual band APs that are second wave 802.11n or 802.11ac are highly likely to be a gateway router. Manufacturer and SSID identify the type of gateway: enterprise, carrier or consumer.
- There are no single band 5 GHz APs. An unmatched 5 GHz beacon implies a missed 2 GHz beacon. This is a dual band access point.
- There are no single band 2 GHz 802.11ac access points. If 802.11ac modulation is seen at 2 GHz, this is a dual band access point and the 5 GHz beacon was missed.

The type of access point can be highly reliably classified by a combination of the radio manufacturer, the access point SSID and the type of radio (speed, protocol, amount of MIMO).

Access Point Type	***Primary Method of Identification***
Carrier supplied residential gateway	SSID and radio manufacturer from MAC address. Primarily dual band except for legacy DSL gateways. Could be DSL, cable, satellite or fiber carrier.
Carrier supplied public hotspot	SSID and radio manufacturer from MAC address. Always dual band. Usually 3x3 MIMO. Always cable carrier.
Consumer supplied residential gateway	Radio manufacturer from MAC address and SSID. Primarily dual band with some legacy single band. Could be DSL, cable, satellite or fiber carrier.
Enterprise access point	Radio manufacturer from MAC address and SSID. All dual band. Almost always 3x3 or increasingly 4x4 MIMO.
Entertainment	Radio manufacturer and SSID. Usually single band with some dual band. Rarely MIMO.

[10] 2 GHz beacons are almost always transmitted for residential access points with 1 Mb/s 802.11b modulation. 5 GHz beacons are almost always transmitted for residential access points with 6 Mb/s 802.11a modulation. This slow modulation rate increases the distance these beacons can be received.

IoT	Internet of Things access points for environmental control, etc. Almost always single band 2 GHz, no MIMO.
Printers	SSID and radio manufacturer. Almost always single band 2 GHz, no MIMO.

Satellite photos were used to count the number of residences in each survey area. Census statistics were used to compute average number of people in each residence.

San Francisco

El Centro

Merced