Mobile Data Trends

Executive Summary

Global industry trends show that the adoption of services such as streaming video and cloud storage continue to increase mobile data usage. The capabilities of the latest handsets and technologies such as LTE and LTE-A give operators solutions to manage this growth even as they fuel it. Drilling down and examining data usage at an individual-customer and specific-location level shows that the situation is complex. These growth trends are not consistent and uniform across the network for all customers. The disparity between average macro-level trends and detailed, specific behavior means that mobile network operators need detailed visibility in order to evolve their networks to cost-effectively satisfy the demands of their customers.

A very few customers use most all data

A recent Viavi Solutions® study of 1.6 million customers spread over a representative area of a Tier 1 mobile network found that 1% of customers consumed over 50% of data. The majority of data usage is confined to a tiny fraction of the overall customer base. Understanding the behavior of these customers is absolutely critical to efficiently deliver capacity and maintain acceptable quality-of-experience (QoE) for ALL customers.

Most data usage is in a tiny area

For the first time, the geographical dispersion of data consumption has been studied, and this produces an even more extreme view than the user distribution. Over 50% of data is consumed in less than 1% of the network area. 90% of data is consumed in less than 5% of the area.

Residential and industrial regions represent 3/4 of data usage

By analyzing the top 100 extreme data locations, the study reveals that data from residential areas makes up 48% of the total and industrial a further 25%. The dense urban areas that would be expected to contribute the most data only represent 13%.

Data consumption profiles in extreme locations are not typical of global trends

The most extreme residential locations have a downlink-to-uplink ratio of only 5.69:1 (5.69 times more data downloaded than uploaded) which is not consistent with consumption being dominated with streamed media, which we would expect to have a ratio that is an order of magnitude higher. The ratio for industrial locations is over double this figure while it would normally be expected to be lower than residential locations due to more mixed business use.

Taken collectively, these findings show that planning network evolution based on industry trends and assumptions runs the risk of failing to bring the best business outcomes. Having access to in-depth, detailed, and granular data about customer usage and experience, and using this to personalize the network for them, is essential to succeed in an ever-more competitive industry.

Introduction

Competitive market pressures continue to challenge operators to create new, innovative strategies to deliver ever-growing volumes of mobile data while driving profits and growth. Network congestion can result in a poor customer experience—making sure that adequate capacity is available to serve customers when and where they want to use services is critical to customer satisfaction.

Previous Viavi surveys have shown that each new device launched encourages customers to consume even more data than the ones before, and the advent of higher-speed LTE services has simply amplified this effect further. These studies have consistently revealed that a small number of 'extreme users' using the latest devices consume a large proportion of the data in a typical mobile network.

This year, the study again examines the distribution of data consumption among users, but has also been extended to examine the geographical distribution of this data use. We analyzed a larger sample of data covering a representative area of 17,461 km2 of mixed urban, suburban, and rural areas. The data set covers all call and data sessions over seven days and represents the behavior of 1.6 million unique subscribers. Data sessions were broken into over half a billion segments and the data consumption and location calculated for each one. Locations were then grouped into 50 m by 50 m bins/map tiles.

Analysis was carried out to understand two aspects of user behavior: the distribution of data usage among users, the extreme users; geographic distribution of usage, the distribution and nature of hotspots. The top 100 hotspot tiles were analyzed to determine the type of location within broad categories, and the data consumed within them grouped within these categories.

Part I: The Extreme Data Users

As noted in previous reports, data consumption is dominated by a relatively small number of users. Understanding the behavior of these extreme users is very valuable to an operator looking to optimize and plan the growth of a mobile network. These users form a significant proportion of the demand, and they are also early indicators of the behavior of much larger numbers of customers further along the technology adoption lifecycle.

The distribution of data consumption by user fraction is shown in Figure 1. It shows that 50% of the data is consumed by just 0.67% of the customers and only 8.58% of customers consume 90% of the data. This finding is broadly consistent with those from studies over the previous four years, allowing for different populations under study.

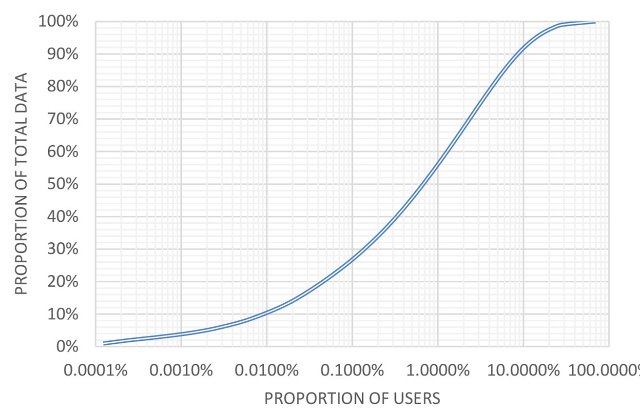

Figure 1. Data consumption by proportion of customers

In previous studies, the ratio of data sent in the downlink (data consumed by the device such as web-page content) to the ratio of the uplink (data sent from the device such as photos and e-mails) has been observed to be 6:1 for UMTS and 5:1 for LTE. In this study, the ratio was examined across the customers from the highest data consumers to the lowest. Figure 2 shows the results.

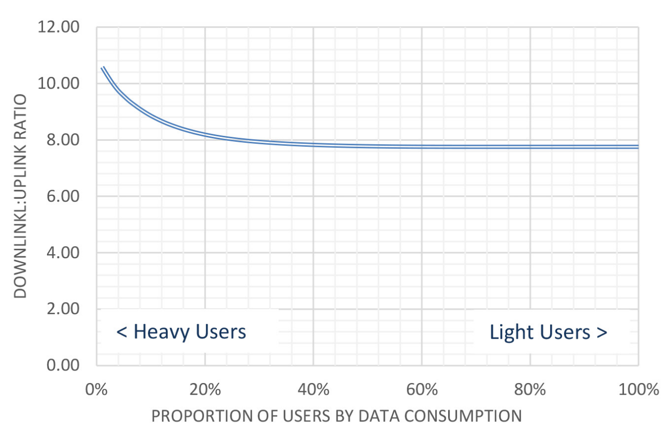

Figure 2. Downlink/uplink ratio by proportion of users

Customers with the highest data consumption use more downlink data relative to uplink at a ratio of 10.58:1 as opposed to the average for this network of 7.75. While this finding is perhaps to be expected, it is generally believed that extreme users will include significant numbers who stream large amounts of video—the fact that the difference is not much larger is in itself remarkable. If the extreme data use was driven purely by streamed content, we would expect the ratio for this group to be in the 100s to 1. No study has been carried out to explain this ratio but we can speculate that it may be the impact of cloud-based services such as Dropbox and Google Drive which have a much more symmetrical profile than streamed content.

Part II: The Extreme Location

New for this study is an analysis of the consumption of data by location. Each data connection is broken into segments, and the location for each segment calculated and grouped into 'bins' or 'tiles' measuring 50 x 50 m. The segmentation of connections is vital as a single data connection may last for many hours, with activity spread unevenly between them. Some other techniques are limited to identifying only the start or end location of the connection or context, leading to a very poor representation of where the data was consumed. The geolocation techniques used for this study allow for a true representation of the location the device was in at the time the data was passed over the air interface.

For each of the 7 million geographic bins identified, the fraction of the total data consumed in that bin for the 7 days was calculated and the bins sorted by the total amount of data. This produces the view in Figure 3 similar to the user view, showing how the data is distributed over the area under study.

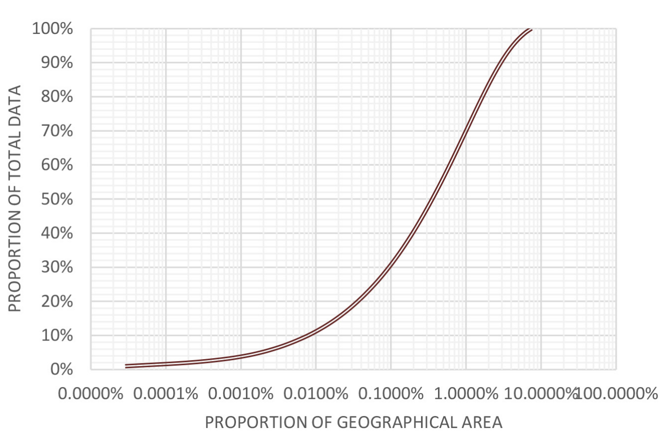

Figure 3. Data consumption by proportion of geographical area

This is a remarkable finding. 50% of the data is consumed in just 0.35% of the 17,461 km2 geographical area covered by this network. 4.47% of the total data is consumed in just the 100 busiest of these 50 x 50 m bins. Looked at in reverse, over 90% of the area generates less than 1% of the traffic.

Typically, it is possible to assume a number of likely candidates for these 100 'extreme locations' such as airports and other major transport hubs. For the purpose of this study, these assumptions were tested. The top 100 bins were identified on a map and the areas classified using available map data and manual verification. Figure 4 shows the distribution of data by area classification.

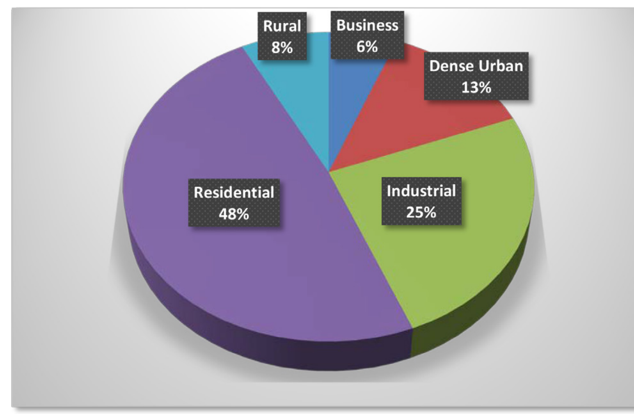

Figure 4. Data by area classification for top 100 50 x 50 m bins

Whereas some of these findings are to be expected, the dominance of residential is surprising. It is recognized that the use of mobile broadband as a substitute for fixed residential broadband varies by operator and even by region within an operator, and this would have a large impact on this percentage. Also surprising is the relative size of industrial and dense urban. In many countries, dense urban areas have been the focus of large deployments of either carrier-controlled or public-access WiFi deployments, and offload onto these may well account for the relative low ranking of these areas. Again, this would be expected to vary by operator and region. That rural areas should appear in this ranking at all is surprising but this illustrates that without good data it is impossible to predict where demand will come from. As with the study of extreme users, it is possible to study the downlink/uplink ratio for these 100 extreme locations and group them by the same categories. Figure 5 shows the results.

Classification	DL:UL Ratio
Business	9.93:1
Dense Urban	7.09:1
Industrial	11.56:1
Residential	5.69:1
Rural	1.93:1

Figure 5. Downlink/uplink ratio by area classification for top 100 bins

It is possible to speculate on the nature of the applications driving usage in these extreme areas, but once again, the results do not align with expectations. It would be expected that traffic in residential areas would be dominated by media streaming, and that may well be the case for the majority of residential areas and the entire area viewed as an average. However, in these extreme locations, we observe a relatively low DL:UL ratio more typically associated with either cloud storage or perhaps private file sharing.

Figure 6 shows a more detailed breakdown of the areas.

Classification	Data Fraction	DL/UL
Dense Urban	11.88%	6.53:1
Farm	0.58%	53.28:1
Industrial	20.39%	11.15:1
Light Industrial	2.23%	9.13:1
Main road	0.72%	26.20:1
Manufacturing	2.18%	27.73:1
Offices	0.46%	38.46:1
Park/Residential	0.66%	18.20:1
Residential	45.02%	5.42:1
Rural	4.55%	2.07:1
Rural/light industrial	1.12%	0.05:1
Rural/Village	0.63%	52.30:1
School	1.33%	18.01:1
Shops	5.46%	9.31:1
University Campus	1.51%	10.13:1
Urban	1.30%	24.68:1

Figure 6. Detailed area classification

Part III: Network Operator Response

The Extreme Response

The ever-increasing demand for data requires a measured response by network operators to keep the services offered to their customers competitive. There are three main ways to increase capacity:

- New air interfaces
- New spectrum
- New cell sites

New air interfaces, including UMTS-R99, UMTS-HSPA, LTE, and LTE-A, have been the focus of considerable research over the past two decades. However, the introduction of each new air interface presents two challenges to the network operator: 1) the time for standardization, manufacture and deployment and 2) the Shannon Limit. The latency associated with standardization is several years or longer, which is considerably slower than the time for data demand to double. Increasingly sophisticated coding and modulation schemes bring each new air interface closer to the Shannon Limit so that the capacity gains due to improved spectral efficiency are increasingly smaller and smaller (see the blue curve that approaches the red curve in Figure 5).

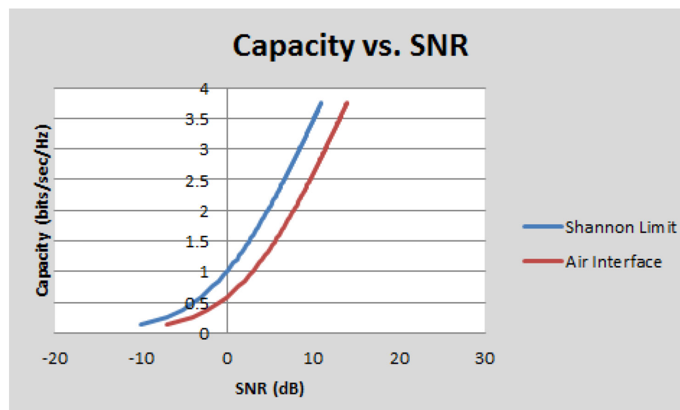

Figure 7. Air interface capacity and the Shannon Limit

The availability of new spectrum provides capacity relief but at a considerable cost and delay. Total air interface capacity is proportional to the product of the available spectrum and the efficiency (in bits/second/Hertz). New spectrum is either auctioned by governments or re-sold by network operators or re-farmed from older technologies. As such, it is a scarce commodity made available with time constraints on the order of years.

There are two types of new cell sites. The first type involves traditional macrocell sites with very large cell footprints. These can take a long time to deploy due to zoning requirements and the fact that most locations that can accommodate macrocell sites already have macrocell sites. The second type of new cell site involves small cells. Small cells can be deployed in a considerably shorter amount of time due to relaxed zoning requirements and due to the reduced complexity of a small cell (often requiring only power and backhaul, aided by self-organizing network capabilities). It is well understood in the industry that small cells will be needed to satisfy increasing mobile data demand.

As well as small cells using cellular technology, operators are looking at using carrier-managed WiFi as a data-offload technology. This can enable service delivery at a lower cost per bit but introduces an additional complexity of managing and optimizing such a heterogeneous network.

Deployment Approaches

For the first time, this report provides empirical evidence to support the anecdotal view that planning small-cell deployments using traditional methods challenges the basic business case for the deployment. If a blanket deployment strategy is adopted, costs will be relatively high, and the opportunity to focus the investment in locations with the most data will be lost. Even a blanket deployment in a dense urban area is unlikely to maximize the opportunity. Targeting small cells, or even macro-expansion in areas expected to be high traffic, based on only on macrocell statistics and local knowledge, is equally likely to be inefficient. The industry anecdotal evidence that this is only 30% effective is supported by the far-from-obvious distribution of traffic shown in Figure 4.

To effectively plan deployments that 'personalize' the network to meet customer needs requires detailed data about several aspects of data consumption, only some of which have been included in this report. In particular, networks need to be planned to consider peak data usage as well as total and typical usage, and studying this aspect will undoubtedly reveal additional short-lived hotspots. Whether traffic is indoor or outdoor and moving or stationary should also be considered when making investment choices.

Summary

The data capacity crunch in the wireless industry started by the launch of the iPhone is not abating. Consistent with prior reports, the hungriest 1% of UMTS users consume approximately half of all UMTS Data.

As well as the data consumption being limited to a relatively small fraction of customers, this usage is concentrated in very small geographical areas. Whereas this in itself may be intuitive, the nature of these extreme locations is not intuitive. While global and national trends driving increases in data consumption are well understood, examining the traffic for the most extreme users and locations indicates that factors other than these must be at work as the usage patterns are not always as expected.

It is well-known that small cells will play an important role in addressing the demands of extreme users and future network evolution (especially in the context of several SON use cases). While the introduction of small cells and the migration to LTE are crucial to satisfying future worldwide data demand, these tasks must be done in a manner that focuses on specific customer demands and the locations where the demand (and attendant performance issues) occur.

Contact Us **+1 844 GO VIAVI**
(+1 844 468 4284)

To reach the Viavi office nearest you,
visit viavisolutions.com/contacts.

mobile-data-trends-wp-li-nse-ae
30179618 900 0116