PUSH TECHNOLOGY:
A KEY INGREDIENT OF APPLICATION INTERACTIVITY

February 2011

Introduction

Since the late 1990s, push technology has risen to prominence twice as the way to sync data between a server and an end device, first for Internet applications and now for mobile applications on smartphones. In 1997, *Wired* magazine predicted the death of the web browser, being replaced with push-based screensavers like PointCast, that flash news, press releases and commercial messages. Thus began Web-based push notification. Flash-forward 13 years and mobile synchronization technologies such as polling and push notifications now enable today's smartphone applications to do amazing things from providing up-to-the-minute news to helping locate your car.

But the rapid growth of smartphones and the explosion of app stores have led to an uncontrolled proliferation of "chatty" network-inefficient applications on a carrier's network. Left unchecked, these applications are starting to cause network congestion, network access failure and poor battery life, all of which contribute to a poor user experience and increased delivery costs for the carriers. Luckily, carriers can solve these problems by switching to an optimized wireless push protocol. This whitepaper will examine today's push technology landscape with a focus on wireless smartphone devices to help readers understand how this technology can drive interactivity with improved user experience.

Application Interactivity: Current demand and future trends

The proliferation of application-rich mobile devices, spearheaded by the introduction of the iPhone in 2007, has caused a culture-changing phenomenon not only in the way people communicate, but, more importantly, in the way they seek information.

Increasingly, mobile phones are being used for data as much, if not more than, for voice communication. The Apple iPhone is leading the way in mobile innovation and impact with its depth of applications and user experience— but Android phones have caught up in popularity. In 2008, email was the single standalone data application, but today the average cell phone is used about 70% for voice, while the iPhone is used only about 45% for voice. Gartner, Inc. forecasts that sales of media tablets will reach 19.5 million units in 2010 and will surpass 208 million units by 2014. With the smartphone revolution well under way and the tablet revolution ramping up, there is the potential for billions of embedded connections that could severely impact wireless networks in the next few years.

Five trends are converging to cause this overwhelming demand for interactivity: 3G/4G network technology, social networking, streaming Internet video, Internet voice (VoIP) and mobile devices. Social networking communications, enabled by mobile phone apps, are mushrooming and e-commerce platforms are rapidly emerging.

Facebook has experienced more than 500 million app downloads and the iPhone has more than 10 billion—this averages out to about 47 apps per user, many of which offer constant contact and updates to users via the Internet. In addition to personal data use, e-commerce is becoming increasingly prevalent and adding to the data traffic with constant push notifications on status (eBay), pricing and sales, as well as cell phone GPS location-based services. New innovative businesses have emerged combining social networking with other businesses like gaming or ecommerce to create a new wave of consumer facing applications that are specifically targeted at a real-time mobile lifestyle and very well adapted to smartphones and mobile devices. As the deployment of new interactive consumer applications on smart mobile devices accelerates, the resulting traffic on wireless networks is bound to grow exponentially.

Impact of interactivity on the network

Clearly, although wireless data and mobile Internet access has long been sought after by the industry and the mobile operator, this exponential growth in demand for interactive data has exceeded all expectations and has placed an unanticipated burden on the existing network infrastructure. For the network, too many applications, too much background data, streaming video and connected social media has resulted in inefficient data utilization that congests the network far beyond capacity.

Now that mobile data is overtaking voice as the primary use of mobile devices, the signal traffic resulting from mobile data use is now outpacing mobile data traffic itself by 30 to 50 percent, if not higher, says a recent Signals Research study. Signaling is how mobile devices request data channels that enable apps to communicate with the Internet in order to check for data updates. The rise in traffic has impacted signaling channels and has limited user access to data channels. Too many applications combined with inefficient push implementations for sending updated data to mobile devices have resulted in high congestion that impacts the user experience and drains battery life. Especially in high use areas such as malls, conventions and events, there is simply not enough bandwidth to service all the users at the same time. The result is slow access times, dropped calls, and sometimes no access at all. Have you ever attended a professional football game at a stadium and been unable to make a call on your mobile phone? That is because all those thousands of people around you are trying to connect to the network as well, not simply to make phone calls, but to watch video and get updates of other games, text message their friends, check email, etc.

> "Signal traffic resulting from mobile data use is now outpacing mobile data traffic itself by 30 to 50 percent"
>
> -Signals Research 2010

In the future, whole new groups of connected devices are anticipated, such as smart meters, in-vehicle telematic systems, environmental sensors and health monitors. These embedded mobile or machine-to-machine connects will lead to an order of magnitude increase in total mobile connections worldwide. The inefficient data utilization comes from each app establishing a separate connection with the cloud and communicating independently from others.

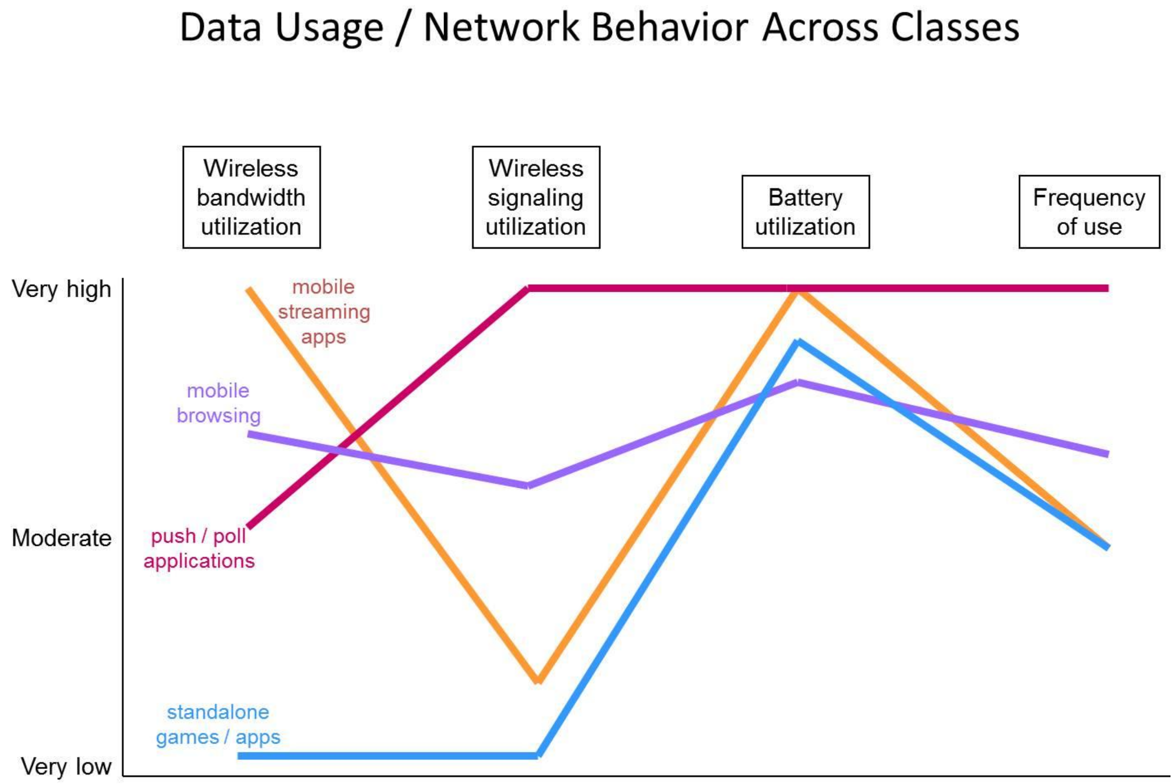

How Polling Works

This automatic check at regular intervals requires a mobile device to move through four states from being idle (state 1) to having a dedicated network channel (state 4) for voice or data communications. Each time the device changes its state, it signals the local radio network controller (RNC) to communicate the status change. Going from idle to dedicated channel, for instance, requires 30 signaling messages. In addition, in several of these states the phone may be in connected mode even when it appears to be just doing nothing. The RNC is designed to handle a certain number of connections based on a statistical analysis of the average number of users in an area and the number of times they use their phones for calls, messaging or Internet access.

Many applications, though, skew this average by using synchronization technology to automatically poll (check for new data) the network on a regular basis without user initiation. This consumes RNC resources, creates massive data traffic and drains the device's battery. In summary, many chatty phones consuming small amounts of bandwidth can negatively impact network loading far more than a few heavy data users consuming more bandwidth but less often.

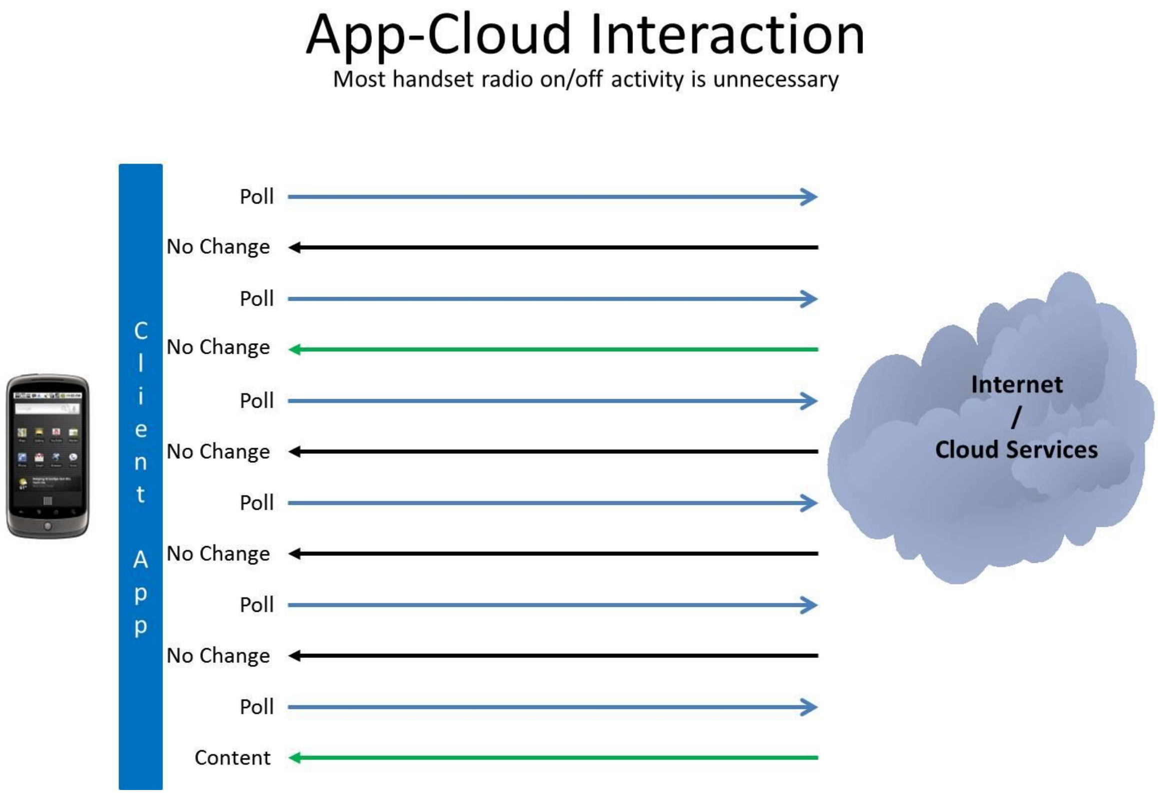

Figure 2. Each app establishes a separate connection with the Cloud and communicates independently from other apps, automatically polling for new information at regular intervals.

Increased signaling traffic puts some pressure on the already congested radio access network (RAN), but most of the pressure is exerted on the core network, which lacks the resources to deal with the increase in signaling traffic. In GSM and HSPA networks, where the same channel transports both voice and data, the congestion affects both services. The problem is exacerbated when users, unable to connect to the network for a call, messaging, or Internet surfing, keep trying, which only increases the amount of signaling, and further increases the congestion.

Analyst Monica Paolini writing in *Fierce Wireless* says that it is unlikely that the 4G LTE and WiMAX deployments rolling out this year will solve the problem, and, in fact, may make it worse as they are likely to generate even more signal traffic through the enablement of VoIP, video calling and viewing and other applications with substantial background processes. 4G technologies are also likely to promote the adoption of much higher numbers of terminal consumer devices such as digital players, automotive modules, mobile health care devices and machine-to-machine (M2M) devices. These devices generate little data traffic, but high levels of signaling.

Survey of mobile synchronization technologies

Three types of mobile synchronization technologies are in use today by carriers and mobile device manufacturers: basic polling, web-oriented push software and optimized wireless push protocol.

Basic Polling

Push protocols evolved to support mobile email and are now being adopted for other mobile content. Regardless of whether the receiver uses polling or push, outgoing data is generally pushed from the sender to the delivery agent using Simple Mail Transfer Protocol (SMTP). However, if the receiver uses a polling protocol for incoming mail, delivery from the user's mail server to the mobile device is done using a poll. Post Office Protocol (POP3) is an example of a polling email delivery protocol. At login and then at regular intervals, the user device polls the server to see if there are new mail or data updates.

Polling is used because even though the server would normally be permanently connected to the device, it does not necessarily always know how to locate the user's device, which may only be connected occasionally and also change network addresses quite often. For example, a user with a laptop on a WiFi connection may be assigned different addresses from the network DHCP server periodically and have no persistent network name. When updates are available on the server to be delivered to the end user on the device, the system does not know what address the client is currently assigned.

Push solutions such as push email provide an "always-on" or "instant notification" experience, in which notification or actual content updates are actively transferred, or pushed, as it arrives on the server, to the user's mobile device. Because push-enabled services provide this "instant notification" experience, they are proven to increase user engagement with the mobile application and thus lower service churn compared to services that are based on a traditional time-based polling mechanism.

> *"[With Push] apps are opened up to 228% more and weekly session times have increased by 103%"*
>
> *-Mobclix 2010*

Web-oriented push software

Several services provide push notifications to mobile devices for email and other applications via web-based push software.

Apple's Push Notification service (APNs) enables applications to integrate with the service and to transport their push notifications to devices and users. Add-ons to APNs such as iLime, Urban Airship, Boxcar and others, augment the basic capabilities of APNs and provide an easy path for application developers to deliver push notifications and in-app purchasing features on iPhone, Android and BlackBerry platforms.

These types of services are key enablers for the explosion of smartphone apps that keep users constantly updated on anything and everything. These web-based services, however, are not designed to coordinate the use of the radio on the device. Figure 3 demonstrates how each

application establishes an individual connection which might be considered optimized for the application but, results in massive amounts of signaling that overburden networks.

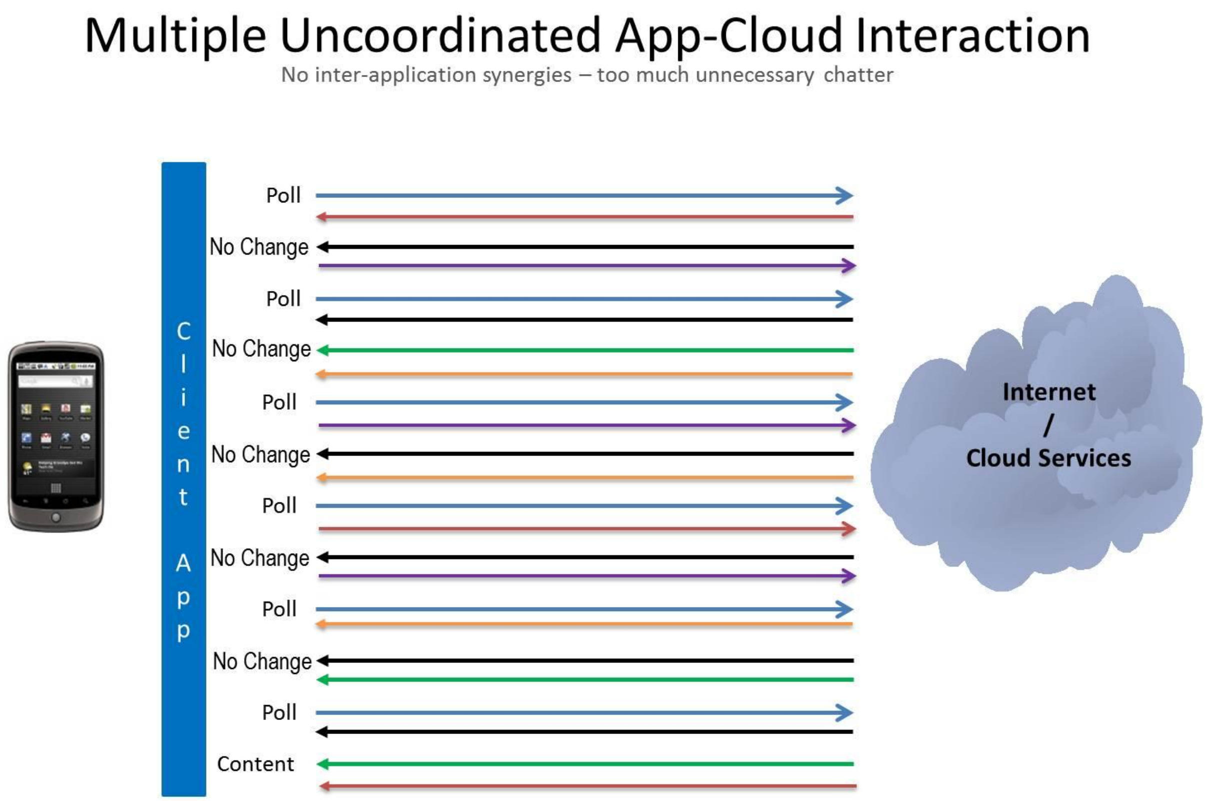

Wireless Push Protocol

The Apple iPhone and Nokia devices with Nokia Messaging, as well as devices accessing Microsoft Exchange with ActiveSync operate with a protocol designed to deliver real-time updates but not always optimized for all aspects of a wireless networks. One basic approach to optimizing synchronization over wireless networks is to only transport data when there is actually real information and updates to be exchanged. This cuts down on the amount of data pushed to each user and the network bandwidth used but is not, in most cases, optimized for other limited network resources such as signaling and device battery.

APNs use push technology through a constantly open TCP connection to forward application notifications and updates from the application servers to Apple devices. The Apple approach is to kill data connection whenever it's not in use as quickly as possible, which can result in requesting new data connections, overloading the network with a huge amount of signaling and requests to get an open data channel on the network. This technique is not unique to the APNS— phones using the Google Android platform use the same kind of battery-saving trick.

The Apple iOS4.2 has taken steps to solve this problem by supporting a new state transition feature developed by Nokia Siemens Networks called network controlled fast dormancy (NCFD), which sets parameters on how, and how often, a smartphone switches between idle and active modes while also preserving device battery life. This could help operators to better manage signaling traffic loads.

Microsoft's Direct Push, more familiarly known as Microsoft Exchange ActiveSync, a built-in feature of the Microsoft Exchange Server, keeps devices current over a cellular network connection. It notifies the mobile device when new content is ready to be synchronized to the phone. When synchronization is complete, a new request is generated to start the process again. This guarantees that Exchange updates are delivered quickly to the mobile device which, in turn, is always synchronized with the Exchange server. However, this requires keeping a data channel constantly open which can be very taxing on the network and the device battery.

Nokia Messaging enables users of Nokia phones to quickly set up, access, and use their email on their devices through a service that is optimized for mobile email. The service automatically delivers push email using the Internet Message Access Protocol (IMAP). When a user's device receives a notification from a server, it may choose to fetch the new data from the server. This makes retrieval of new messages and data more flexible than a purely-push system, because the client can choose whether or not to download the new data.

SEVEN Optimized Push Technology

"SEVEN push reduces the amount of data transferred and uses an optimized transport protocol to decrease data transferred by 8X"

-Mobclix 2010

SEVEN has taken a unique approach to its push technology where it delivers on the need for instant notification and content while taking into account the limited bandwidth, signaling and battery resources available on the network and the device. The SEVEN Optimized Push technology leverages both the data channel and the SMS connection in a hybrid fashion to wake up the device when dormant, open a data connection when the server application signals that new data is available, and to tear down the connection within a timing interval that is optimized for the user's behavior and the various applications in use on the device. The SEVEN Optimized Push technology both reduces the amount of data transferred and uses an optimized transport protocol to decrease data transferred by 8X when compared to non-push synchronization. This reduction in radio activity is a major factor in increasing device battery life, as well as decreasing the amount of signaling and data traffic on the network.

SEVEN's push technology is distinguished from the competition by very low power consumption and very fast update pushing. Some tests show that SEVEN will push messages with one minute in most cases and it only consumes several kilo bytes data traffic per day and about 0.01W power.

SEVEN Push Notifications – the future of push services

While the SEVEN push technology has been optimized for SEVEN's own applications, such as Mobile Email and Mobile IM, SEVEN has open API's for its Push Notification service that allow application developers, operators and device manufactures access to the company's multi-service, multi-device, and provide them a means by which to leverage the company's push technology—specifically notifications and content synchronization capabilities—for any messaging services and mobile applications. With SEVEN Optimized Push Notifications, application developers and content providers can directly leverage SEVEN's proven synchronization technology rather than having to redevelop their own.

Up until now, the bulk of these services have been available only on very expensive smartphones, which ignores the nearly 70% of mobile users. SEVEN Optimized Push Notifications both builds on SEVEN's success with bringing engaging services to all mobile devices and optimizing data transport for carriers.

With SEVEN, users get content on their device when they become available. Delays in data updates and the need for users to check for updates manually are eliminated. With SEVEN's Optimized Push Technology, data is exchanged on the network only when needed - this significantly decreases the radio activity on the device, decreasing signaling on the network and extending device battery life.

Network Testing Results

The best mobile synchronization technology is optimized for the constraints of a wireless network with software on both a device and a server that has the intelligence to communicate only when needed. That's directly opposite of the polling model that is being commonly used by most mobile applications (and some 'push' models) with a network signaling level that is far beyond the typical or even common push model. It's also in contrast with any web-based technology where bandwidth and signaling constraints are not as prevalent. A comparison test, featuring SEVEN's Optimized Push Notifications vs. a polling protocol, shows how dramatic the difference can be: the optimized protocol reduced data transfer by up to 80 percent, and reduced battery power consumption by 50 percent, an important end-user benefit.

Conclusion

Network signaling can only become an even greater problem as smartphones, tablets and apps become more sophisticated and popular. In order to solve the core problem, carriers, device manufacturers and developers have to rethink the way mobile devices and applications interact with the network. The right technology exists today, but a carrier might have deployed push for messaging, the handset manufacturer might have a different protocol for social media, and an application provider could use yet another protocol for its own application, all of them operating independently and unsynchronized. Carriers, manufacturers and app developers will need to work together on careful and deliberate deployment of highly efficient push technology in order to mitigate network congestion, while also allowing carriers to cost-effectively scale their networks to meet the expanding demand.