Airvana
Mobile Broadband. **Everywhere**.

Mobile VoIP over 1xEV-DO

A Technical Whitepaper

Paul Callahan
VP, Business Development
Airvana, Inc.
pcallahan@airvana.com
+1 978 244 1877

The Business Case for Mobile VoIP

The combination of high performance mobile broadband and Voice over IP (VoIP) stands to usher in a new era for the wireless industry. This paper makes a direct comparison between VoIP over EV-DO and 1xRTT circuit-based networks. To date, CDMA2000 1xRTT has been the most spectrally efficient mobile voice network available, but the combination of VoIP and EV-DO – or broadband mobile VoIP – will represent a dramatic improvement over all previous generations of mobile voice technology.

- **Higher voice capacity –** Mobile VoIP promises to deliver up to 120% more voice capacity than circuit-based mobile voice.
- **Lower costs –** Because they are packet networks end-to-end, VoIP-based mobile networks are easier to manage, more technology efficient, and lower cost to operate than circuit-based mobile voice networks.
- **Faster development of applications and services –** Instead of taking years to develop and deliver new services, operators can create new applications and services in weeks or months on mobile VoIP platforms.
- **Wider range of access networks –** Instead of being restricted to a single access network, VoIP applications enable operators to deliver services over multiple access networks: EV-DO, wireline, home Wi-Fi, and public Wi-Fi.

But these gains are dependent on technology developments in two critical areas: the emergence of a new class of mobile VoIP phones and enhancements to the radio access network (RAN) to specifically support VoIP and multimedia traffic.

Mobile VoIP Phones & 1xEV-DO Rev A

VoIP has become mass-market technology because of the widespread availability of always-on cable and DSL broadband Internet access. Though these are wired connections, ubiquitous Internet access has driven the development of packetized telephony and multimedia over IP. As a result, broadband-connected PCs support a combination of soft phones, instant messaging, collaboration tools, and video telephony.

But two stumbling blocks have prevented the VoIP from entering fully mobile access networks. Unlike mobile phones, laptops are not ideal for true mobility – they are portable at best. In addition, previous generations of mobile data networks did not support the stringent quality-of-service requirements that VoIP and multimedia applications demand. The development of mobile VoIP handsets and a RAN designed to carry delay-sensitive multimedia traffic will change that picture.

- **Mobile VoIP phones emerge –** A new class of mobile phones will soon be delivered – based on packet VoIP instead of circuit technology – offering a wide range of integrated voice and multimedia applications. Connected to high-quality mobile broadband, wireless VoIP phones will fuel rapid application and service innovation, and radically change the economics of operator networks.
- **1xEV-DO Revision A (Rev A) has been designed for VoIP –** Rev A has been standardized by the Third Generation Partnership Project 2 (3GPP2) and specifically enhanced to support VoIP and multimedia applications with quality of service (QoS), high capacity reverse link, low latency air link, fast call setup, reduced packet overhead, and rapid handoff. The Rev A RAN will be faster, higher quality, and lower-cost than any previous generation of mobile network.

Shifting to Packet

Rev A's support for VoIP and multimedia services has tipped the balance between circuit and packet-based infrastructures. Because VoIP offers higher capacity voice and better performing data at lower operating costs, operators should begin capping investments in additional 1xRTT circuit capacity.

Capping Growth of Circuit-Switched Technologies

As mobile VoIP handsets come to market, 1xRTT will become a less profitable means to offer voice services. Why?

- **Voice average revenue per user (ARPU) continues to drop** – Operators continue to report the decline of voice revenue per user. For example, KDDI reports strong growth in data service ARPU, but a decline in voice[1] (see Figure 1). As a result, dollars invested in 1xRTT equipment will result in lower profit margin for operators.
- **Voice minutes continue to climb –** As mobile operators continue to offer flat-rate and bulk voice pricing plans, mobile minutes of use (MoU) per month are increasing. Leap Wireless reports that 52% of their customers use mobile phones as their only phone service and talk an average 1,500 MoU per month[2]. Verizon Wireless customers increased their MoU 16% in 2005 and T-Mobile reports that their customers talked an average of 1013 minutes per month in Q1 2006[3].
- **Data is attracting higher ARPU customers** – Newer, higher paying customers are attracted to broadband data. In the overall operator revenue mix, voice will increasingly give way to data. Verizon Wireless, for example, reports a relatively flat ARPU for combined voice and data, but rapid growth in data ARPU[4] (see Figure 2).

---

[1] KDDI Financial Results

[2] Leap Wireless Financial Results

[3] T-Mobile Financial Results

[4] Verizon Wireless Financial Results

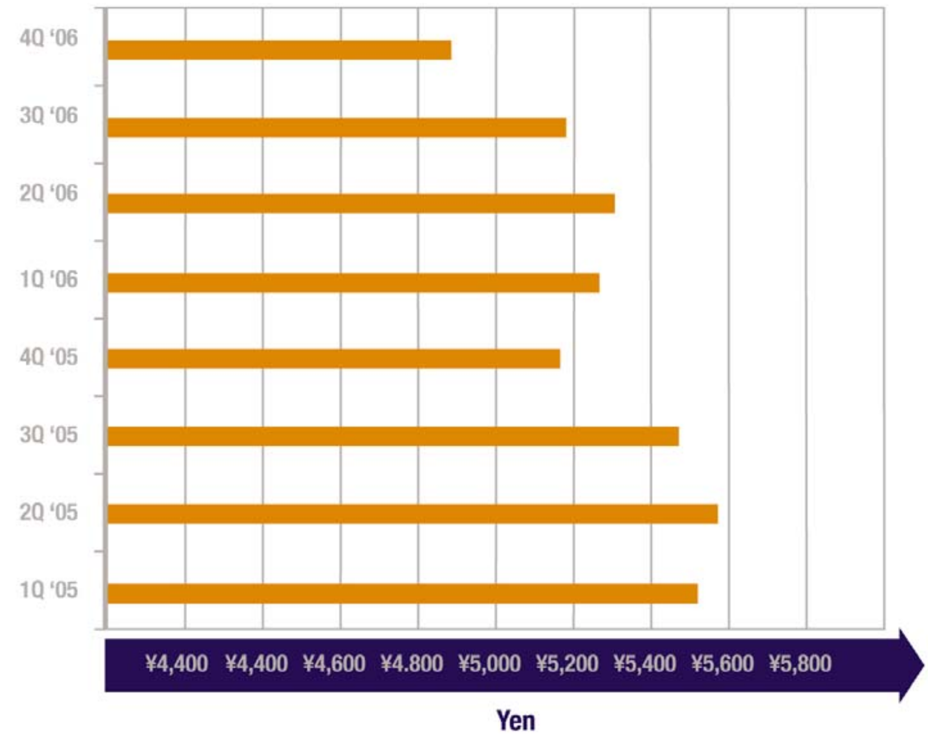

Figure 1: KDDI's declining 3G mobile voice ARPU, Source: KDDI

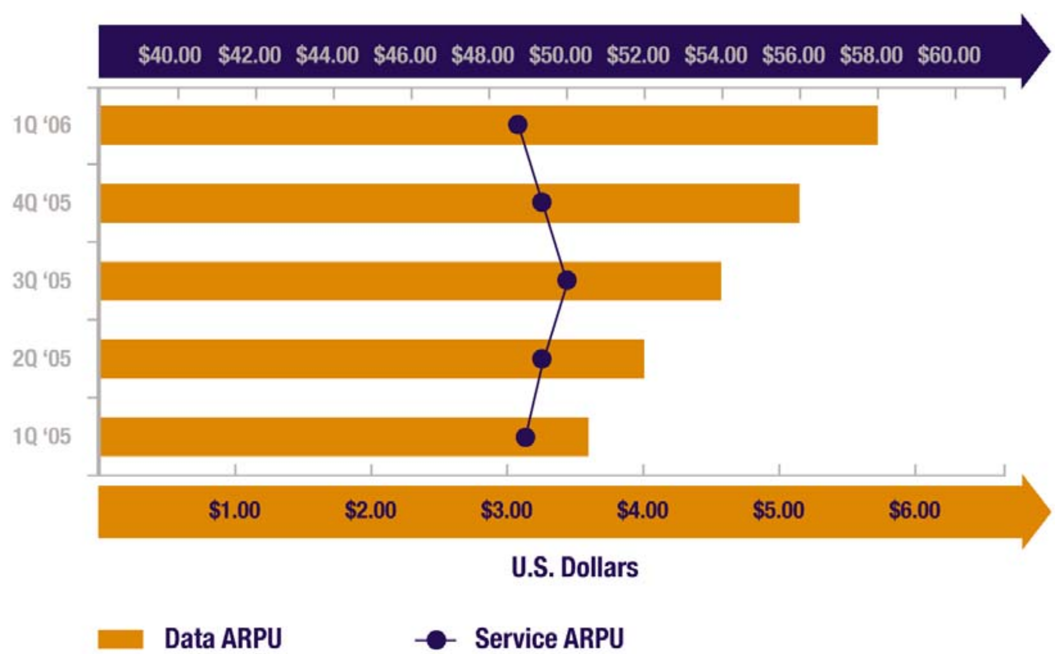

Figure 2: Verizon Wireless increasing data ARPU, Source: Verizon Wireless

Higher Capacity with Mobile VoIP

VoIP over Rev A offers a way to increase voice capacity more cost effectively than continuing to add 1xRTT circuit modules. Technology advances in CDMA have continually lowered the overall price-per-bit for RAN transport. This has allowed manufacturers to offer higher capacity Rev A modules at a unit cost that is roughly equivalent to older, lower capacity 1xRTT. In addition, VoIP over Rev A offers:

- **Higher voice capacity than 1xRTT** – Today's circuit-based 1xRTT supports a practical maximum of 27 simultaneous voice users per sector (per 1.25 MHz of channel) with an acceptable mouth-to-ear delay of 290 ms. VoIP over Rev A supports 40-60 simultaneous users[5], a 48-120% increase in voice capacity[6] over 1xRTT.
- **Higher quality than 1xRTT** – Circuit-based 1xRTT delivers 290 ms ear-to-mouth delay, but Rev A supports better quality overall under a variety of conditions. For example, with 27 simultaneous users – the equivalent of 1xRTT's maximum capacity – Rev A provides voice latency as low as 215 ms (see Table 1).

Simultaneous VoIP Users	20	30	40
End-to-End Latency (msec) (2 - Handset Antennas)	215	218	222

Table 1: Latency vs. number of simultaneous VoIP users, Source: Airvana, Inc.

Simultaneous Data + VoIP Performance

Even with guaranteed voice quality of 270 ms mouth-to-ear delay for 40-60 simultaneous VoIP users, Rev A makes 400 Kbps per sector available on the forward link for data users (see Figure 3). This is in sharp contrast to 1xRTT, which cannot deliver simultaneous voice and broadband data under any circumstances.

[5] "VoIP over cdma2000 1xEV-DO Revision A," Yavuz et al., IEEE Communications Magazine, February, 2006

[6] "QUALCOMM Successfully Demonstrates Fully Mobile VoIP Calls Across a Number of Field Test Environments," QUALCOMM Press Release, June 07, 2006

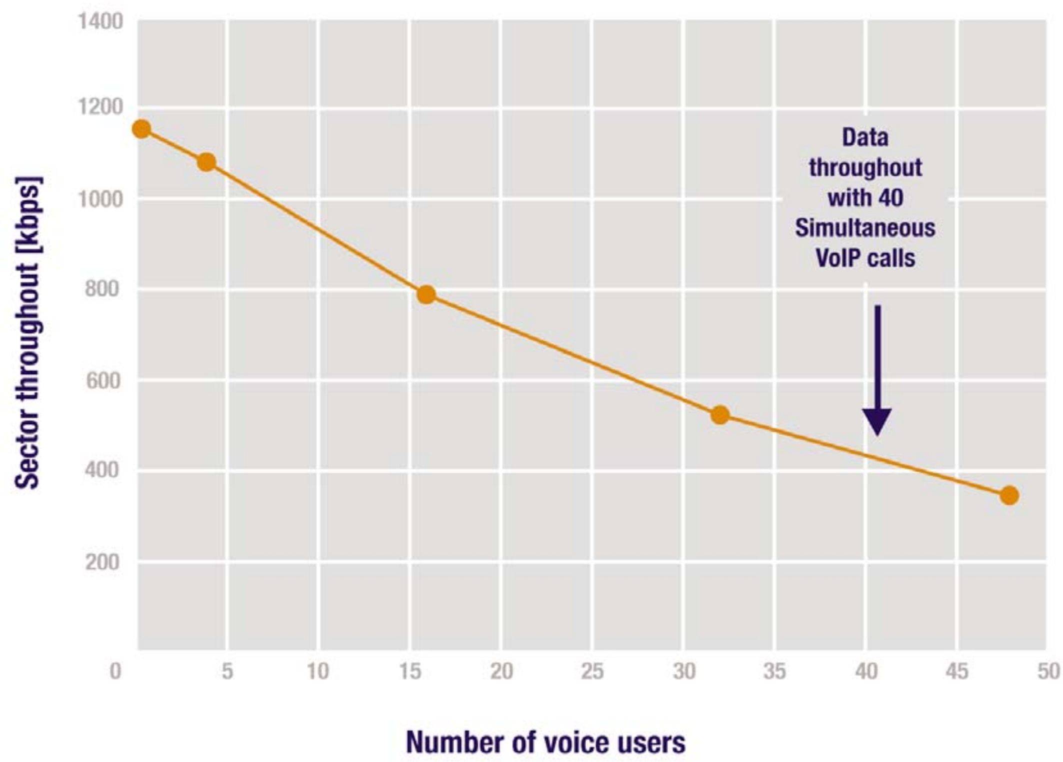

Figure 3: Simultaneous VoIP and sustained data sector throughput, Source: Airvana, Inc.

Lowering Costs with Mobile VoIP

VoIP implementations will lower costs for mobile operators in two areas. First, mobile VoIP will make more efficient use of spectrum by using the same channels or carriers as data, and by dynamically sharing the load between the two. Second, because mobile VoIP will utilize the same all-packet core with data, management complexity, power consumption, capital expenditures, and maintenance costs will be dramatically less than for two separate networks. Industry consensus is that by shifting to a common packet core, operators can experience significant cost savings overall when compared to the cost of building and operating separate circuit and packet networks.

Spectrum Efficiency

Current CDMA networks require one radio channel for 1xRTT circuit voice, and one radio channel for broadband data (EV-DO). This separation of voice and broadband, data at the radio level, results in inefficient use of operator's spectrum. Given the cost of licensed spectrum acquisition, setting aside spectrum for one use, as opposed to another, can be inefficient and expensive.

In the past, hybrid (circuit voice + packet data) air interface technologies such as 1xEV-DV were considered to overcome this inefficiency. However, such hybrid schemes were difficult to implement because of the complexities involved in mixing circuit and packet traffic on the same channel.

Now with VoIP over EV-DO, spectrum efficiency can be enhanced by “trunking” or combining the use of channels. In current CDMA networks, when voice usage is high but broadband data usage is low, it is not possible to offload voice traffic to the broadband data channel and vice versa. Migrating to mobile VoIP over EV-DO allows voice and broadband data to be simultaneously carried over the same channel or carrier. This spectrum versatility provides efficient carrier load balancing for VoIP and data, further increasing overall efficiency. For example, if a sector is experiencing light data traffic but heavy voice traffic, VoIP calls can dynamically consume the bulk of the available channels. In a scenario where one carrier is dedicated for voice using 1xRTT, and one carrier for data using EV-DO, voice capacity tops out at 27 simultaneous users – even if the data carrier is lightly loaded (see Figure 4). By comparison, two carriers running EV-DO and VoIP would support a much higher voice capacity because of the combined use of the two carriers for voice and data.

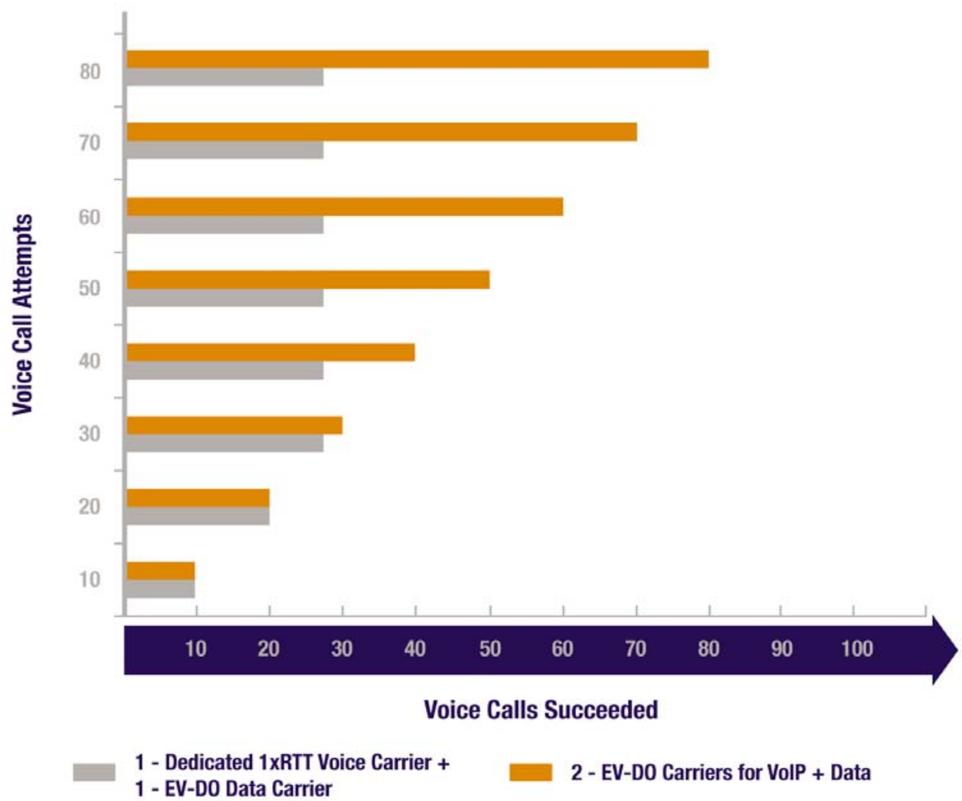

Figure 4: Trunking efficiency via mobile VoIP on Rev A

Core Network Cost Savings

Today’s mobile networks involve deployment and maintenance of two parallel networks, one for voice and one for packet data (see Figure 5). Mobile VoIP gives operators efficiency and cost savings through a single packet-based core network. Because mobile VoIP uses one core IP network to handle both voice and data, the infrastructure becomes:

Streamlined – When separate networks are used for voice and data, there are many highly redundant elements, such as the base station radios (BTS and RN) and the controllers (BSC and RNC). The additional cost for the carrier is not only in the capital cost of the equipment, but also in the ongoing maintenance and management expenses associated with these redundant elements. Figure 6 shows how the infrastructure is streamlined once the circuit network is eliminated.

Lower cost – In the transition to a packet core, all circuit-oriented network elements disappear and costs are reduced. Other circuit elements are replaced by lower cost VoIP elements. For example, the circuit-based Mobile Switching Center (MSC) is replaced by a Media Gateway (MGW) and packet voice switch (softswitch) or SIP server.

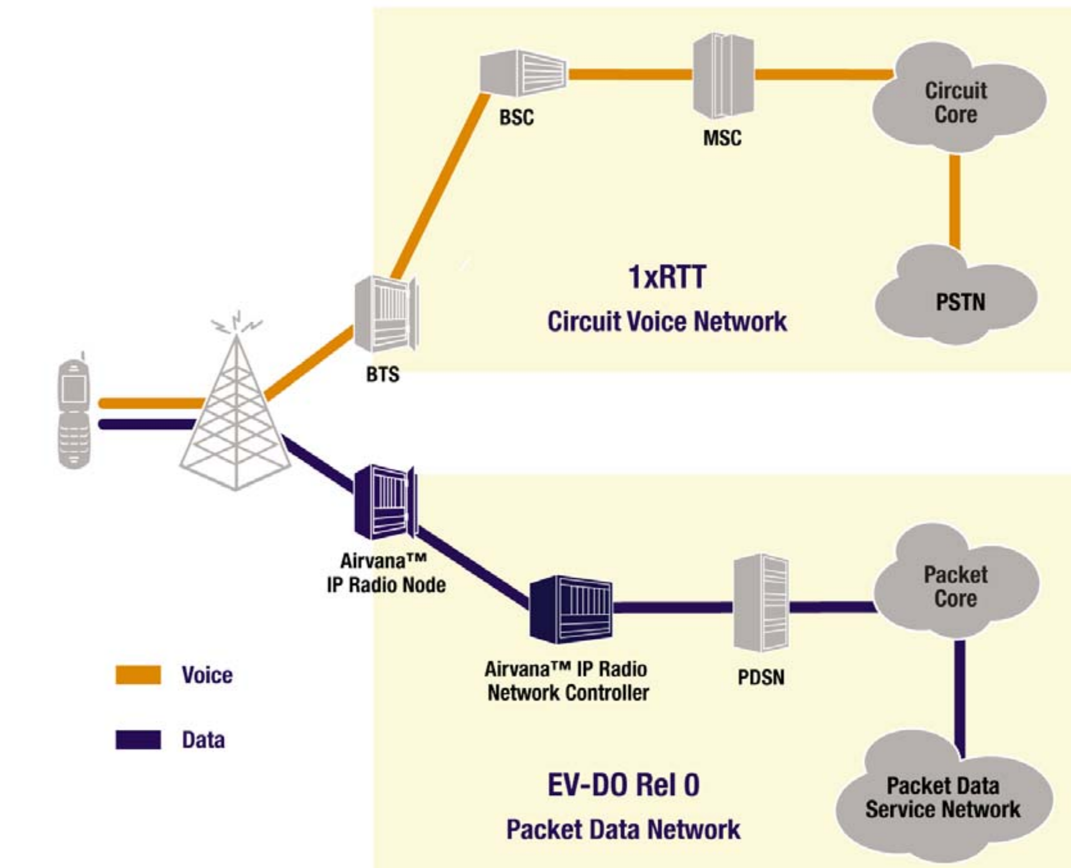

Figure 5: Separate mobile voice and data networks

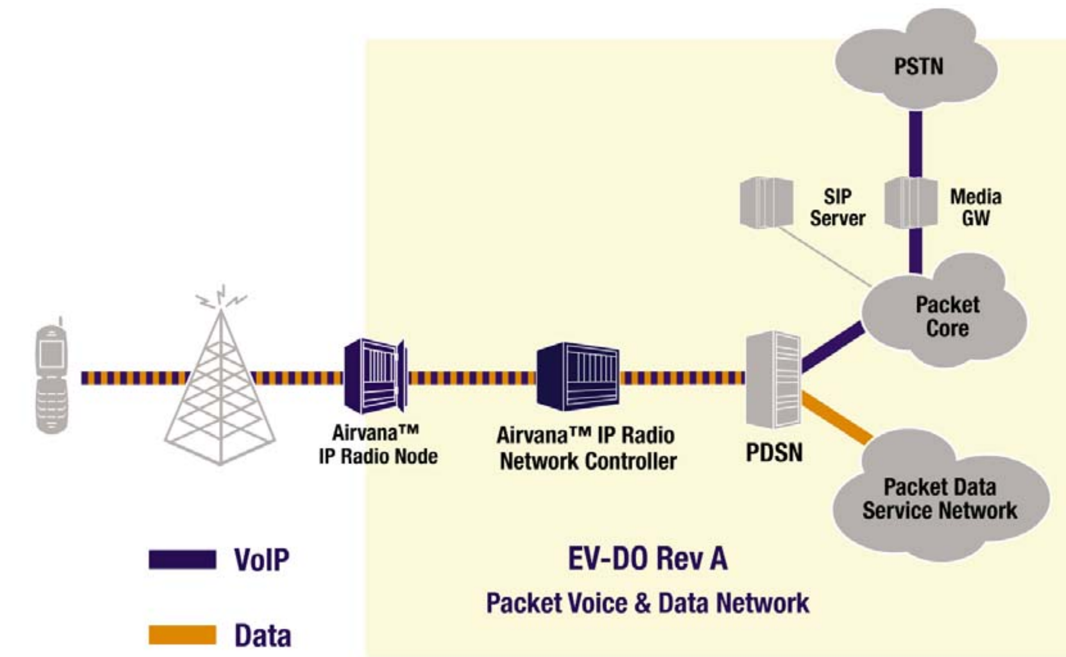

Figure 6: Packet RAN and packet core network

Mobile VoIP - New Applications & Services

Just as widespread availability of wired broadband propelled VoIP adoption, so the large scale rollout of 1xEV-DO Rev A will become the driving force for mobile VoIP. Service providers seek to provide services anytime, anywhere, running over any device. As an increasing number of service providers roll out 1xEV-DO networks around the world, Rev A will become an ideal platform to provide ubiquitous communications and entertainment experiences on the move – and on multiple IP-enabled devices (cell phone, laptop, or PDA). New mobile VoIP applications include:

- **Push to Talk (PTT)** – Push-to-Talk with sub-second call setup and Push-to-Conference with full duplex Push-to-Talk
- **IM** – Instant messaging with integrated voice and video, application sharing, and file transfer
- **Video Telephony** – Simultaneous voice and video, mobile-to-mobile, with low delay

The first widespread mobile VoIP application will be PTT. Sprint has chosen Rev A as the technology to migrate millions of iDEN subscribers to a more scalable version of PTT based on IP. Sprint's offering will use an optimized version of PTT, Qualcomm's QCHAT, with very fast call set-up times. Coupled with QCHAT, Rev A will be the first cellular network to deliver sub-second call set-up for Push-to-Talk. Because PTT operates on an end-to-end basis, network-wide adoption of Rev A is required and will establish Rev A as the first proven network for scalable multimedia services. PTT and other IP-based services require a number of underlying capabilities which are only available via VoIP:

Context Awareness – VoIP allows users to display their context. Users can choose to publish whether they are busy, off-line, not taking calls, have calls forwarded to another device, in a meeting, driving, only reachable by certain people, what device they are currently reachable on, their location, or who else is on a conference call (see Figure 7). This context awareness enables and enhances a whole class of emerging next generation of mobile applications.

Application Integration – The VoIP environment allows subscribers to have certain applications appear as buddies in IM, or make their presence or identity known in other ways. For example, this capability could allow a purchasing agent to get an IM notification when supplies are low.

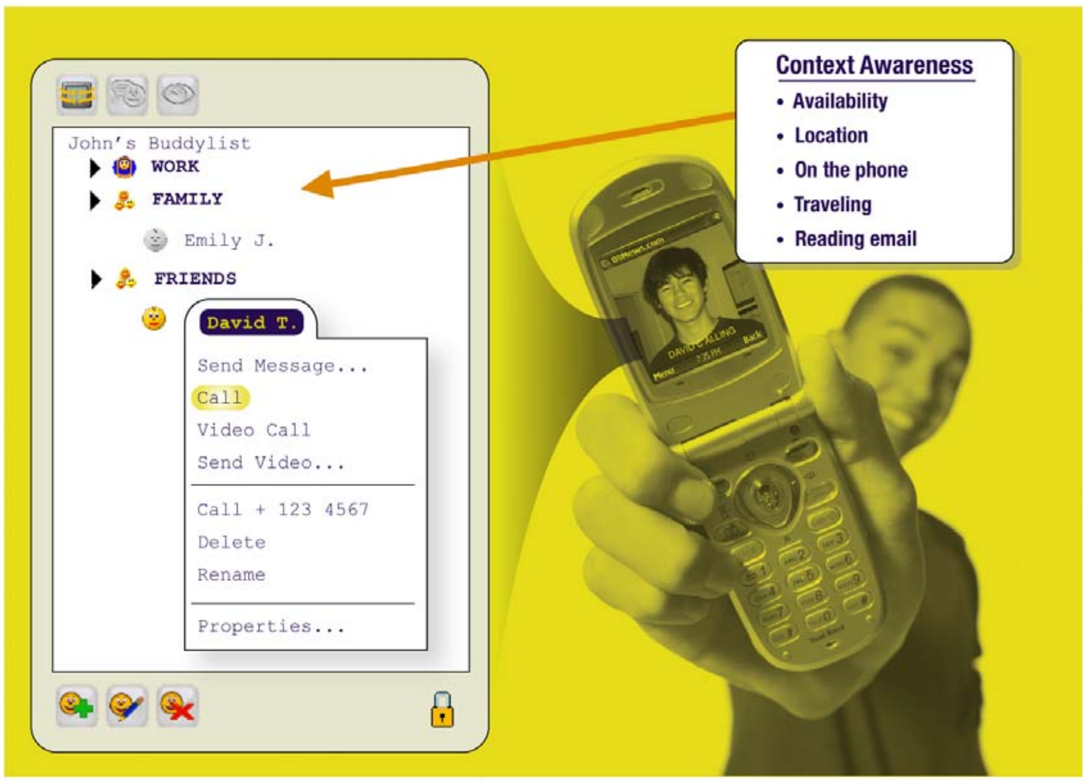

Figure 7: Context awareness via mobile VoIP

Mobile VoIP Support in Rev A

EV-DO users are projected to grow from 53 million by the end of 2006 to 342 million by 2010[7]. This rapid growth will increase demand for Rev A's increased capacity and higher data rates. Additionally, Rev A has been designed to meet operator requirements for a RAN that can transport delay-sensitive applications such as VoIP, Push-to-Talk, and video telephony together with bursty, delay-tolerant Web traffic. VoIP-related features of Rev A include:

- **Flow-Oriented (QoS)** – Differentiated QoS on multiple flows provides higher priority for a VoIP call than a Web session, even on the same mobile phone.
- **High Capacity Reverse Link** – Higher uplink speeds enable VoIP and other symmetric services such as Push-to-Talk and video telephony.
- **Low Latency Airlink** – Reduced latency improves the quality of VoIP and Push-to-Talk.
- **Fast Call Setup** – Short call setup improves performance of Push-to-Talk and VoIP.
- **Reduced Packet Overhead** – IP header compression improves the VoIP and Push-to-Talk capacity of the RAN.
- **Rapid Handoff** – Early indication of base station switching facilitates non-disruptive handoff for VoIP calls.

Flow-Oriented QoS

Rev A has adopted a flow-oriented QoS architecture that allows the network to provide differentiated treatment to separate application flows, even when the applications are associated with the same handset or smartphone. To deliver such differentiated treatment, EV-DO Rev A utilizes higher layer protocols between the network and the handset to set up and configure these flows based on the needs of the applications. By doing so, Rev A can handle separate flow types with differentiated levels of service, across the network – a critical requirement in a mixed application environment (see Table 2). This allows Rev A to assure the delivery of critical delay-sensitive traffic flows like VoIP even in the presence of bursty Web and e-mail traffic.

[7] Strategy Analytics, 2006

IP Flow Type	Requirement	Flow Handling
Application Signaling for Push-to-Talk, VoIP, Video	Low-latency, Shortest call setup	High priority packet scheduling, air link retransmissions enabled
Push-to-Talk, VoIP & Video Telephony Audio Bearer	Low-latency, low-rate	High priority packet scheduling, air link retransmissions disabled
Video Telephony & Video Bearer	Higher-rate	Medium priority packet scheduling, retransmissions may be enabled
Web, FTP	Latency insensitive, high-rate	Best effort scheduling, air link retransmissions enabled

Table 2: IP flow types and handling

Centrally Controlled Handset QoS – The access network determines resource allocation for each application flow, and sets the rules for the handsets. This unique interaction between the access network and the handset enables access network resources to be centrally managed while distributed handsets handle applications and flows. This is especially important for VoIP and other delay-sensitive applications that require QoS.

Multi-Flow at Handset – Unlike EV-DO Release 0, 1xRTT, and HSDPA, which only support a single flow per handset, Rev A allows separate applications, such as VoIP vs. Web browsing, to handle traffic differently on a flow by flow basis. For example, the transmission of delay-sensitive VoIP packet may be delayed when it waits in the scheduler queue in the base station behind a large delay-tolerant packet destined for a Web browser. With a Rev A QoS enabled handset or smart phone, the smaller delay-sensitive packet would be sent to the VoIP application ahead of the larger delay-tolerant Web packet. Using such QoS treatment, Rev A delivers the low latency and low jitter attributes necessary for multimedia applications.

High Capacity Reverse Link

Increased data rate and throughput for the EV-DO uplink is also critical to the support of VoIP, Push-to-Talk, and video telephony applications. To support these latency-sensitive and symmetric applications, EV-DO Rev A increases the uplink peak rate to 1.8 Mbps, and improves overall sector throughput to 700-800 Mbps. More importantly, Rev A also reduces the uplink's overall latency characteristics, and allows the data rate to be selected on an application-by-application basis and based on sector load.

Higher data rate and throughput – Hybrid Automatic Repeat Request (H-ARQ) combines rapid error detection and time-slotted, interlaced sub-packets to improve data rate and throughput. Because sub-packets carry redundant information about the packet being delivered, H-ARQ provides fast recovery from lost packets. As a result, H-ARQ recovers from lost packets without forcing retries in higher-layer protocols, even at markedly higher data rates. In addition, uplink packets are power-boosted to allow early termination of packets with high reliability. This early termination, or "guess" as to the packet's composition before it has completely arrived, also boosts reliability and performance.

Dynamic Rate Control – Unlike the fixed rates provided by 1xRTT, Rev A dynamically adjusts data rates based on a combination of available handset transmit power and the loading of the cell sector. Each handset receives and processes instantaneous sector load information sent to it by multiple cell sites nearby. The handset then factors in the transmit power of its own packets as they have been received at the cell site to dynamically pick a data rate. The handset determines this data rate on a flow-by-flow basis – thus allowing a VoIP call to operate at a different data rate from a Web data session. This flow-by-flow handling greatly increases the throughput for a sector by allowing VoIP to operate at lower data rates without taxing higher speed applications like FTP or Web access.

Low Latency Airlink

In Rev A, a broad range of system-wide latency reductions and fast call setup techniques have been implemented. The bulk of these reductions are to support VoIP, Push-to-Talk, and other latency-sensitive applications.

Low Latency Handset Flows – Handsets can transmit packets in two modes: low-latency or high-capacity. For example, a mobile phone may transmit VoIP or other latency-sensitive traffic in a low-latency mode and best effort data traffic in a high-capacity mode. Each flow is then assigned a low-latency or high-capacity attribute as the preference for transmission. Additionally, low-latency packets send smaller, and redundant, power-boosted sub-packets to ensure they can hit their latency targets. Larger high-capacity packets are transmitted at normal power levels.

Short and Multi-User Packets on Forward Link – Rev A's forward link, or downlink, supports short packets (128, 256, and 512 bits) to further reduce latency for applications like Push-to-Talk and VoIP. In addition, Rev A can deliver multi-user physical layer packets that allow higher packing efficiencies and thus better capacity.

Lower Latency on Access and Control Channels – Rev A also offers enhanced access channel rates of 19.2 Kbps and 38.4 Kbps to facilitate faster network access. This combined with shorter packets on the control channel, reduces connection set-up time or "push-to-beep" time for Push-to-Talk applications and call set up for VoIP. Push-to-beep time is the most critical user satisfaction metric for Push-to-Talk services.

Fast Call Setup

While short call setup is most critical for Push-to-Talk, it is also important for VoIP. Both Push-to-Talk and VoIP rely heavily on paging messages – short messages to setup calls to idle or sleeping handsets. Rev A uses the following enhancements to improve call setup time, system-wide.

Data over Signaling (DoS) – allows the handset to exchange initial call set up messages before a traffic connection is established. This enables the call set up time to be reduced by as much as 800 msec. Without DoS, Push-to-Talk users would experience significant and unacceptable push-to-beep delay.

Enhanced Idle State Protocol (EISP) – enables shorter waking intervals for handsets operating in stand-by mode. In Push-to-Talk applications, it is desirable to configure handsets to wake up from sleep mode more frequently to look for incoming calls, thus reducing push-to-beep time. Rev A's new sub-synchronous control channel facilitates the configuration of very short waking intervals (less than 429 ms).

Distance-based paging – allows the RAN to selectively send Push-to-Talk and VoIP pages to specific cells, instead of flooding the networks with page messages over the control channel. Distance-based paging allows large scale Push-to-Talk networks to achieve greater paging efficiency by delivering an optimal utilization of shared channel resources by idle handsets.

Reduced Packet Overhead

Robust Header Compression (RoHC) is a compression technology standardized by the Internet Engineering Task Force (IETF, RFC 3095) to further optimize IP for VoIP. Because VoIP frames represent a small payload (22 bytes), standard RTP/UDP/IP packets that carry these frames would add significant header overhead if left uncompressed. RoHC reduces the header overhead to 2-4 bytes, by eliminating most of the unnecessary redundancy built into packet headers.

Rapid Handoff

When mobile data-only handsets cross cell boundaries and experience handoffs, they normally experience a 100-200 ms delay. For a mobile VoIP user, this would be unacceptable. Most de-jitter buffers on mobile VoIP handsets will be able to handle 20-30 ms delay and make it undetectable to the voice user on a call, but 100-200 ms delay is beyond the capability of most de-jitter buffers and voice quality would suffer. To eliminate this delay, forward sector switching in EV-DO has been enhanced with the introduction of a Data Source Control (DSC) channel. The DSC channel allows the handset to send an early indication of an impending handoff to the base stations.

Virtual Soft Handoff – Just before a mobile VoIP handset performs a handoff to another cell, it sends a message to its current cell and the next nearest cell indicated that it is about handoff using the DSC channel. This message is then passed up to the Radio Node Controller (RNC) which then starts sending traffic to both cells simultaneously. This facilitates a low-latency, non-disruptive soft handoff for mobile VoIP.

Conclusion

1xEV-DO Rev A has been designed to deliver broadband everywhere with unprecedented performance and reliability. Rev A will enable mobile networks to offer a wide range of multimedia applications and rapid service creation – all with the highest capacity and lowest cost.

Mobile VoIP’s low-cost handsets and efficient network will alter the economics of broadband wireless networks and create significant competitive and economic advantage for the operators who adopt it. The combination of mobile VoIP and EV-DO Rev A will provide solutions for both ends of the market: low cost solutions for developing markets, as well as high-performance systems for urban environments.

Glossary

For a glossary of technical and marketing terms used in this paper, go to: http://www.airvananet.com/technology/technology_glossary.htm .

About Airvana

Airvana builds mobile broadband infrastructure solutions deploying an innovative IP-Radio Access Network architecture that delivers carrier-class performance and reliability. Networks based on Airvana's technology service millions of consumer and business subscribers; and a wide range of mobile broadband information, communications, and entertainment applications.

Worldwide, Airvana's mobile broadband systems based on advanced CDMA2000 1xEV-DO are deployed on six continents in 16 major networks by industry-leading service providers who demand high standards of carrier-class performance. In North America, Airvana's largest deployments are Sprint PCS (USA), Verizon Wireless (USA), Bell Mobility (Canada), and Telus (Canada). The company's products are also deployed in networks including Telstra (Australia), Vesper (Brazil), Eurotel (Czech Republic), and Pelephone (Israel).

In addition, Airvana builds solutions for other wireless markets including in-flight, air-to-ground communications; and fixed-mobile convergence (FMC) for the integration of cellular and IP-based access technologies. Airvana is headquartered in Chelmsford, MA, USA.

For more information, please visit the company's Web site at http://www.airvana.com