Volume 10 Issue 01 Published, February 15, 2006 ISSN 1535-864X DOI: 10.1535/itj.1001.08

Intel® Technology Journal

Converged Communications

Seamless Collaboration – Enabling Best in Class VoIP Experience on Intel® Centrino® Mobile Technology

More information, including current and past issues of Intel Technology Journal, can be found at:
http://developer.intel.com/technology/itj/index.htm

Seamless Collaboration–Enabling Best-in-Class VoIP Experience on Intel® Centrino® Mobile Technology

Hani Elgebaly, Mobility Group, Intel Corporation
Farid Adrangi, Mobility Group, Intel Corporation
Rajeev Muralidhar, Mobility Group, Intel Corporation
Lakshmi Ramachandran, Mobility Group, Intel Corporation
Manish Nair, Mobility Group, Intel Corporation

Index words: Voice over Internet Protocol (VoIP), seamless collaboration, wireless LANs, Quality of Service (QoS)

ABSTRACT

With the introduction of Intel® Centrino® mobile technology, today's mobile notebook is an advanced communication platform with integrated wired and wireless technologies. One of the important communication technologies that has been making excellent progress in the industry is Voice over IP (VoIP). The Centrino notebook augments ingredient technologies to improve the VoIP experience for the end user, resulting in enhanced productivity.

Recent advances in VoIP applications, protocols, infrastructure, equipment and peripherals have resulted in IT departments of medium and large businesses becoming interested in deploying this technology. The goal is to have the business road-warrior communicate in a seamless way both inside and outside the enterprise. Also with the advancements in the Internet Protocol Private Branch Exchange (IP PBX) infrastructure, IT shops have been racing to replace their legacy, cumbersome Time Division Multiplexing (TDM) PBXs with easy-to-manage and more flexible IP PBXs. The impediments to such deployment have been the ability to deliver the expected Quality of Service (QoS) for voice calls, call privacy, and ease of use of VoIP applications.

In this paper, we present the platform ingredient technologies that Centrino offers this year (2006) in order to deliver a best-in-class VoIP experience over wireless and wired networks. This includes Intel's QoS solution for enhanced VoIP experience over Wireless Local Area Networks (WLAN), Intel's Array Microphone for better open audio, Intel's Integrated Performance Primitives (IPP) libraries for improved overall VoIP experience, and enabling the use of Bluetooth* headsets for hands-free VoIP.

INTRODUCTION

VoIP is being increasingly deployed by enterprise IT departments as a cost-effective replacement for PBX-based telephone networks. The single biggest advantage of deploying VoIP is that IT administrators have to maintain just one kind of network for both data and voice applications resulting in significant cost savings. Moreover, the quality of VoIP calls has improved significantly over the past few years as voice quality-related issues have been addressed in data networks as well as in VoIP devices. VoIP is also gaining popularity in the consumer space, in homes and Wi-Fi public hotspots like hotels, airports, and coffee shops.

The popularity of WLANs has tremendously increased among notebook users with ubiquitous Wi-Fi deployments in enterprise networks and popular places like airports, hotels, and coffee shops. Enterprises of all sizes are adopting WLANs as connectivity extensions of their corporate network beyond office spaces and cubicles, or in some cases as a replacement for the wired network.

As VoIP and WLAN deployments continue to gain momentum in the enterprise, VoIP usage over WLAN is also expected to gain popularity, especially among notebook users with softphone applications. Supporting adequate quality for VoIP over WLAN has some unique challenges arising from the characteristics of WLANs. Typically, the bandwidth available over a WLAN is an order of magnitude less than that available over a wired LAN. Moreover, a WLAN signal is susceptible to interference from adjacent WLANs as well as other radio devices. This causes increased delays and sometimes losses for packets transmitted over WLAN. Another big challenge is that of Quality of Service (QoS) for VoIP traffic. Intel's Seamless Collaboration architecture, described in this paper, addresses WLAN QoS challenges and introduces other platform technology ingredients to improve the VoIP experience on Intel Centrino mobile technology notebooks.

In this paper, we first provide an overview of Intel Centrino 2006 Seamless Collaboration scenarios. We then describe the network topology and landscape for VoIP and its deployment challenges in enterprise networks, and for VoIP QoS over wired and wireless networks. In the following sections, we present Intel's QoS architecture for VoIP over WLAN, and we describe the Intel Integrated Performance Primitives (IPPs) that are available for softphone optimization. We then provide an overview of Intel's Array Microphone architecture and Bluetooth-Wireless Coexistence solution for Bluetooth headset enabling.

VOIP SEAMLESS COLLABORATION USAGE SCENARIOS

Seamless Collaboration brings some new usage models to Intel Centrino mobile technology 2006 platforms with the focus on enhancements for integrated communication devices. These new usage models allow the notebook to become an effective communication device for voice anytime/anywhere. There are several hardware/firmware, software/drivers, and optimization components that support the Seamless Collaboration architecture.

There are three key usage scenarios that will be enabled by Seamless Collaboration: "Office on the Go," "Multi-party VoIP conferencing," and "Integrated communication experience."

Office on the Go

This enables users to use their Intel Centrino mobile technology notebooks as a central communication device for voice anytime/anywhere. The user can make and receive voice calls anywhere as long as he/she has IP network connectivity to his/her enterprise network, as if he/she were physically present at the office. Many users will be connecting to their respective IP networks via WLAN and will use wireless headsets/handsets using Bluetooth. This usage model also allows the user's calls to be automatically routed to the user's current location. For example, the user could have incoming calls to both his/her desk and cell phone be automatically forwarded to his/her notebook softphone, whenever connected to the enterprise network. The user would be able to configure which phones and calls are automatically forwarded to his/her current location. This way the user can effectively create a virtual office experience.

Figure 1: Office-on-the-Go scenarios

Multi-Party VoIP Conferencing

This feature enables users to use VoIP softphones for multi-party audio conferencing (including adhoc and bridge-based conferences) with optimized narrow-band and wide-band audio codecs and Intel Microphone Array open audio technology. This usage scenario allows a user to conference in more than two parties using the softphone application on the notebook, thus saving on expensive conference bridges. Dual core and other platform ingredient technologies enable hosting N multi-party conferencing sessions on Intel Centrino mobile technology, where N is greater than 8.

Integrated Communication Experience

One instance of this usage scenario provides users the facility to receive their voice messages as e-mail in their Outlook* inbox. This allows users to access their voice mail stored as e-mail in their inboxes and to listen to them in any order. Users can also choose to reply to the voice mail via a text message. Other instances of this usage scenario include integrated optimized softphones with productivity and e-business enterprise applications (Click-to-Dial scenarios). For example, if an employee in the company's Enterprise Resource Planning (ERP) system has a question or inquiry about a transaction, he/she can click to dial customer support and instantly talk with a support representative.

* Other names and brands may be claimed as the property of others.

Key VoIP Client Components

Softphone Application

A softphone is a software VoIP application that provides a phone-like user interface to a notebook user complete with a dialing pad on the screen that can be clicked with a mouse. The softphone application may interface with external peripheral devices such as Bluetooth or a USB microphone and/or headset. Some softphone applications may also allow a user to configure VoIP parameters such as voice codec, call characteristics, or data rate.

Most softphone applications require a call manager on the network that authenticates softphone users, maps phone extensions to IP addresses, and routes VoIP call signaling messages between softphone users. A call manager can run on a server for PC-to-PC calling. Alternatively, a call manager can be integrated with a hybrid PBX, in which case it can also provide a connection to the enterprise and external PSTNs.

Audio Peripherals

A softphone application uses notebook audio peripherals such as codecs, built-in speakers, and microphones to convert user speech into VoIP packets and vice versa. Alternatively, USB or Bluetooth peripherals, such as handsets, headsets, speakers, and microphones, can also be used as audio peripherals.

Audio Codecs

A software audio coder-decoder (codec) is needed to sample and encode audio from a microphone to bits that can be sent in a VoIP packet and to decode bits from incoming VoIP packets to audio signals that can be played through the speakers or headphones. A variety of codecs is available for different conditions such as available bandwidth, audio quality, protection from lost packets, etc. Most audio codecs developed in the last few years were either designed to provide quality comparable to PSTNs (e.g., G.711 [8]) or to allow the use of VoIP with low-bandwidth networks such as dial-up (e.g., G.729 [8]). These codecs typically sample human voice signals at 8 kHz to allow representation of audio frequencies up to 4 kHz. Codecs with sampling rates up to 8 kHz are known as narrowband codecs. More recently, wideband codecs have been developed with sampling rates as high as 16 kHz, which typically result in better perceived audio quality than is possible with narrowband codecs, because audio frequencies up to 7-8 kHz can be adequately represented. The use of such codecs is however limited to PC-to-PC VoIP calls, because a PSTN typically cannot transport or reproduce audio frequencies higher than 4 kHz.

Each codec can work with one or more frame rates, e.g., 20 ms, 30 ms, etc. A frame rate represents the time interval for which converted audio bits are encapsulated in an RTP packet to be sent to the other VoIP endpoint. Depending on the encoding algorithm, a codec produces either a fixed size or a variable size (in bits) audio sample. The data packet transmitted over a LAN or WLAN also consists of a User Datagram Protocol (UDP), an Internet Protocol (IP), a Logical Link Control (LLC), and Media Access Control (MAC) headers.

VOIP DEPLOYMENT IN ENTERPRISE NETWORKS

Traditionally, enterprises have maintained two separate networks: a voice network, based on PBX and PSTN and an IP network for data applications, such as email, Web, VPN, etc. VoIP technology allows the separate voice and data networks to be merged into a single network because voice can be treated as just another application running over the IP (data) network. Traditional PBX and phone extensions can still be supported using hybrid PBXs that can manage traditional (analog) phones as well as VoIP phones connected via Ethernet. Many enterprises have chosen to deploy a single IP network in their new office buildings where there is no need to support existing analog (PSTN) phone extensions. In addition, many enterprises have replaced their analog phone networks with VoIP [9].

Figure 2 shows an example enterprise network with desktop and notebook clients connected to LAN and WLAN segments. The WLAN segment is operated by an IEEE 802.11 Access Point (AP). These segments are connected to the corporate network via a wired backbone. Other components on the corporate network shown are enterprise servers, a hybrid PBX, legacy phones, and VoIP phones. In addition to these components, call managers are responsible for registration and authentication of VoIP clients and for routing VoIP signaling messages to the correct VoIP endpoints. Call managers also manage phone extensions mapped to VoIP phones (soft and hardphones). A media gateway, on the other hand, converts voice samples from VoIP packets on the VoIP network to a modulated analog signal used by the PSTN.

Beyond the AP, the enterprise network comprises a WLAN Distribution System (DS)–typically implemented as an Ethernet-based LAN, Layer 3 routers, switches, IP / Multi-Protocol Label Switching (MPLS)/Differentiated Services (DiffServ) domains, etc. Networks at multiple sites may be interconnected via VPN/MPLS tunnels/frame relay to give the appearance of a single corporate network. A VoIP call initiated from a WLAN client might terminate at an end host that is outside the Extended Basic Service

Set (BSS) serviced by the WLAN DS. In order to provide service differentiation across the entire end-to-end VoIP data flow path, the AP and each router/switch on the way must be configured to identify and prioritize VoIP flows. Details of QoS support required in various enterprise network components are given in a later section.

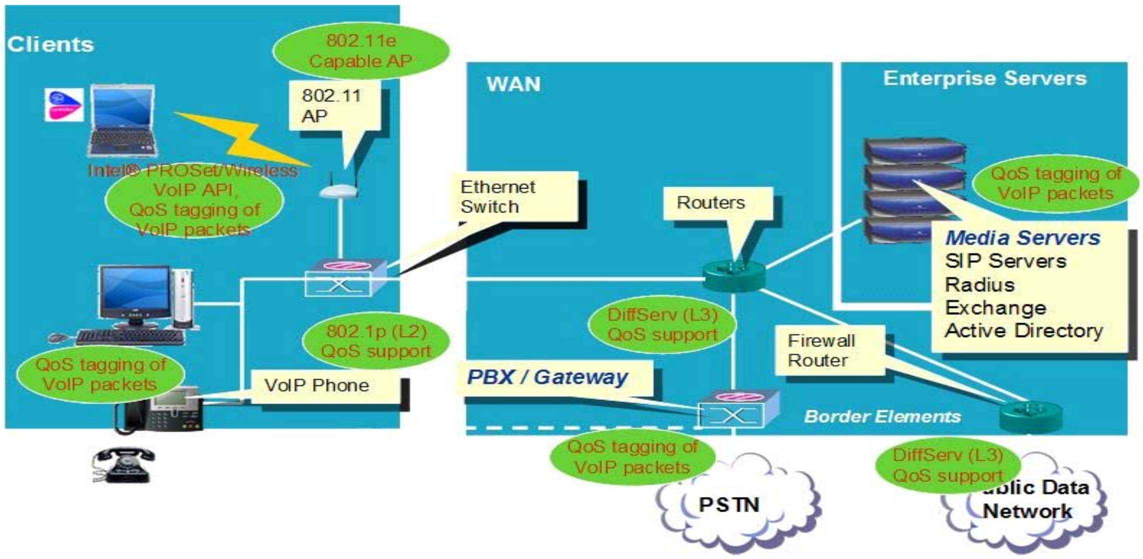

Figure 2: Enterprise VoIP deployment across wired and wireless networks

VoIP QoS over Wired Network

In order to provide service differentiation over the end-to-end VoIP data flow path, each router/switch on the way must be configured to identify and prioritize VoIP flows. This can be achieved as follows using Differentiated Services Code Point (DSCP), and it is also shown in Figure 2.

1. VoIP client endpoints must mark the VoIP packets that they send out with a pre-configured DSCP value.
2. In a large enterprise network consisting of multiple campuses/sites, when a VoIP packet arrives at a border router in the VoIP sender's network, it is identified as a VoIP packet from the DSCP in the IP header. The border router may use another DSCP value to ensure that this packet is given a prioritized treatment in the core enterprise network or VPN/Frame Relay tunnels that connect multiple sites.
3. When the VoIP packet reaches a border router in the receiver's network, it is identified as a VoIP packet based on the DSCP marking used in the core enterprise network. The border router may change the DSCP to a value that is reserved for VoIP flows in the receiver's network. The packet is then forwarded towards the receiver.

An IP PBX can use DSCP tags on its signaling packets so that these are forwarded with a higher priority over best-effort traffic in the enterprise LAN. If a SIP server is used for signaling, it can also send SIP signaling packets with the right QoS tags. A media gateway acts like another VoIP endpoint as far as the data network is concerned and hence it also tags the VoIP packets it sends with the appropriate DSCP values.

VOIP QOS OVER WLAN

The first published IEEE 802.11 WLAN standard [1] did not have any provision for supporting QoS over the wireless medium. This problem was subsequently addressed by the draft IEEE 802.11e standard [2]. In addition, an industry body known as Wi-Fi Alliance, formed for certification of IEEE 802.11 standards, published its own interoperability specification called Wireless Multimedia (WMM) [3], which is based on an earlier draft version of IEEE 802.11e. QoS support in IEEE 802.11e comes in two flavors:

1. **Prioritized QoS:** This allows classification of WLAN traffic into different categories based on their priorities. Higher priority traffic is given preferred access to the WLAN over lower priority traffic. This is achieved by a channel access function known as Enhanced Distributed Coordination Function (EDCF), or Enhanced Distributed Channel Access (EDCA). Prioritized QoS only provides a statistical guarantee. Voice and network control traffic are given the highest priority followed by video, best-effort data, and background data.
2. **Parameterized QoS:** Support for parameterized QoS is provided using a centralized Hybrid Coordination

Function (HCF) at the WLAN AP. HCF allows the wireless medium to be alternately used for contention-based access (using EDCF) and contention-free access (using HCF controlled channel access or HCCA). In the HCCA scheme, the AP grants opportunities for the WLAN Stations (STAs) to transmit by polling them based on their traffic requirements.

Admission control can be used with both the above schemes to limit the type and mix of calls on the AP. Parameterized QoS is optional and is not widely supported by AP and WLAN client vendors. This paper refers to only the EDCA scheme of IEEE 802.11e/WMM.

To achieve priority-based packet processing, a VoIP client needs to mark the VoIP packets that they send out with a pre-configured DSCP value. This will ensure that VoIP packets are processed with high priority according to the WMM/IEEE 802.11e specification implemented by the MAC layers of the WLAN client and the AP.

If the receiver is also a WLAN client associated with an AP, IP routing and Layer 2 switching will deliver the VoIP packet to the AP. The AP identifies the packet as a VoIP packet from its DSCP marking. Additional classification based on IEEE 802.11e Traffic Classification (TCLAS) may be performed by the AP in its WLAN MAC layer to classify the packet to a particular Traffic Stream (TS) for an associated WLAN client. The VoIP packet is then scheduled for delivery to the WLAN client by the AP's MAC layer.

Priority-based packet processing using EDCA mechanisms definitely helps real-time traffic like VoIP, but this is not adequate to provide hard guarantees. For example, an EDCA-based WLAN cannot support an infinite number of VoIP calls even though VoIP traffic can be treated with higher priority. Based on channel conditions, existing load on the AP, etc. there is a limit on how many calls can be supported. Once this limit is reached, MAC collisions, retransmissions, etc. cause delays to all the VoIP calls, thereby degrading the call quality. Hence, call admission control is needed. Once the AP enforces a limit on the number of VoIP calls that can be supported as high-priority calls, all subsequent requests must be rejected in order to maintain the delay/bandwidth guarantees expected for the already admitted calls. More details on how prioritization and admission control work in WLANs are available in IEEE 802.11e/WMM.

Our simulations show that both priority packet processing and call admission control are essential for providing call capacity and bandwidth guarantees to VoIP calls. Figure 3 shows the improvement in VoIP capacity with EDCF, over plain IEEE 802.11 WLANs, for a single AP. The effect of admission control on the quality of the VoIP calls can be seen in Figure 4, where the quality of admitted calls does not go down when traffic exceeds the capacity limit. Calls that are denied admission are treated as best effort, so that admitted calls continue to get the reserved QoS.

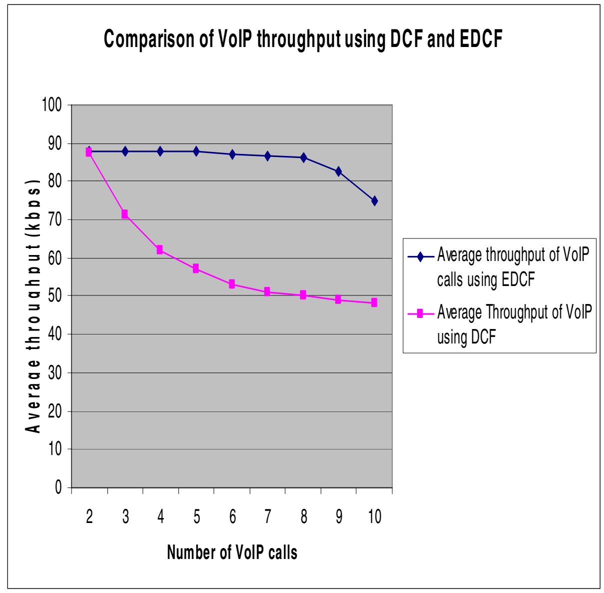

Figure 3: Comparison of performance of VoIP call in DCF vs. EDCF

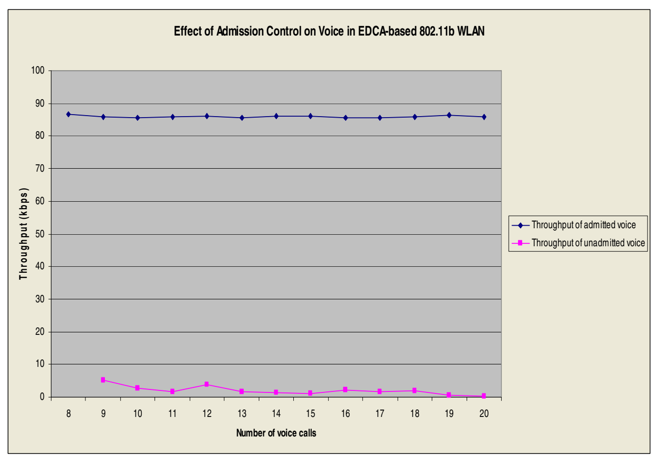

Figure 4: Effect of call admission control on the performance of VoIP calls

VOIP OVER WLAN CLIENT ARCHITECTURE

Figure 5 shows the client side of Intel VoIP over WLAN architecture, which consists of the following key components:

1. Softphone application.

2. Intel WLAN Softphone Application Programming Interface (API): Softphone applications use this interface to convey VoIP profile information to the WLAN QoS Manager. A *VoIP profile* is a collection of application-level parameters that defines the QoS requirements of the application. This API provides profile and flow management functionality. In addition, this can also be used to collect statistics for a VoIP flow such as packet loss, delay, and jitter, which can be used by the application for troubleshooting or for renegotiating the call.

3. WLAN QoS Manager: The WLAN QoS Manager is the interface between the softphone application and the WLAN driver. This is represented by the "Winsock Layer/GQoS API" block in Figure 5. It provides translation of a VoIP profile into appropriate parameters as understood by the WLAN MAC layer, i.e., traffic specification (TSPEC) and traffic classifier (TCLAS) information as defined in IEEE 802.11e/WMM.

4. WLAN MAC Layer: The Client WLAN MAC layer is split across the WLAN driver and the WLAN NIC. It contains both standards-based implementations of IEEE 802.11/.11e as well as any optimizations that Intel VoIP over WLAN architecture provides for improving the VoIP experience.

5. Intel Management Application: Intel management application can be used by IT managers to configure new VoIP profiles and to delete/modify existing VoIP profiles.

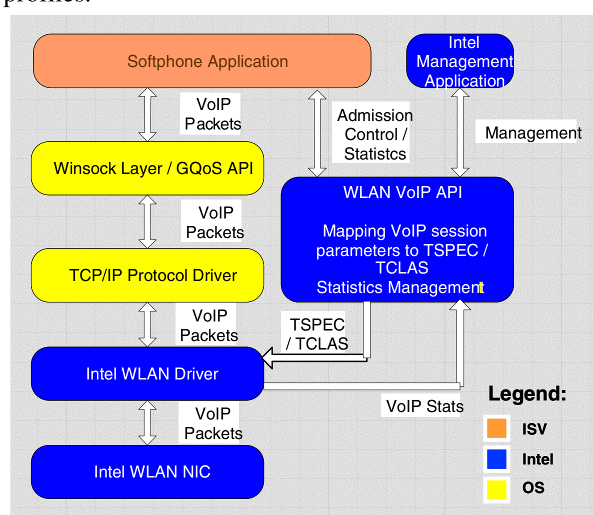

Figure 5: VoIP over WLAN client architecture

Note that VoIP and signaling packets generated by the softphone application (and its peer application/call manager) traverse the path shown through the Winsock [7] layer, TCP/IP stack, and the WLAN subsystem. The control path shown on the right side of the diagram is used to set up a VoIP flow over the WLAN network between the client and the AP.

Figure 6 shows an example sequence of VoIP flow creation in the Intel VoIP over WLAN architecture. When the softphone (VoIP) application is ready to initiate or accept a VoIP call, it first queries the list of VoIP profiles supported by the Intel WLAN VoIP API. From this list, the application picks a VoIP profile, depending on the codec and frame rate it has selected for the call, and it issues a *Create Flow* request to the VoIP API with the selected profile and a classifier, that describes the VoIP traffic. The VoIP API and QoS Manager map the VoIP profile to a WLAN Traffic Specification (TSPEC) and the classifier to WLAN TCLAS and request the WLAN driver to add a traffic stream over the air using these parameters. The WLAN driver sends the request to the AP according to the protocol specified in WMM/IEEE 802.11e. If the request for Create-Flow is accepted, the AP allocates medium time to the VoIP call and returns a successful response to the client. Otherwise, the AP simply returns a failure response to the client. The WLAN driver then processes the response and returns a success or a failure to the QoS Manager for the Created-Flow request. Lastly, the QoS Manager makes a callback into the softphone application via the VoIP API to inform it of the result of the call creation.

In case of a failure response, the softphone application may generate a busy signal indicating to the user that resources for the call are not available. The user may try the call again at a later time.

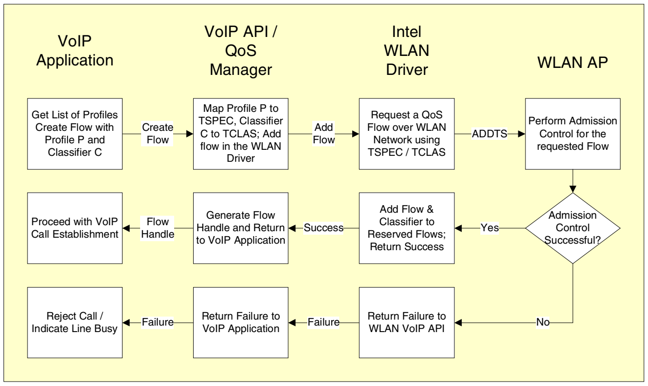

Figure 6: Example of VoIP flow creation over WLAN

VPN Implications

Most VPN technologies rely on "tunneling" the original data traffic from one VPN endpoint to another using some form of encapsulation. Since the original data packet is encrypted before it is transported over a VPN tunnel, header attributes from the original packet such as UDP/TCP port numbers, IP addresses etc. are not visible to a WLAN driver when it examines a VPN packet. This limits the ability of WLAN drivers to correctly identify and prioritize a flow. As an example, if an application creates a VoIP flow reservation in the WLAN driver and provides a classifier containing source/destination IP addresses, port numbers and protocol, the driver will be able to create the flow but will fail to identify packets that belong to this flow. Thus the VoIP packets will not be assigned to the correct flow reservation and will not get the priority over WLAN that they deserve. To avoid this, when a VPN connection is in use, the application should use only a DSCP value as a classifier instead of using a full transport layer classifier. When the application detects the presence of a VPN, it can use a classifier consisting only of tunnel source and destination IP addresses and/or DSCP value in the IP header instead of using a full transport layer classifier.

Today's VPN products give more importance to security than performance. Both L2TP and PPTP (two popular VPN tunneling protocols) bring in performance issues due to the processing and overhead involved in encrypting and encapsulating the packets. In PPTP, the packet is encapsulated inside a Generic Routing Encapsulation (GRE) packet, which is then encapsulated inside an IP packet before being sent across the tunnel. In L2TP, packets are encapsulated 4 to 6 times depending on the IPSec policy used. It also provides additional levels of security through the use of DES/3DES encryption which impacts performance.

To avoid performance and classification problems in the WLAN driver, it is highly recommended that enterprises use the IEEE 802.11i/WPA2 method instead of VPNs to encrypt WLAN traffic within the enterprise.

INTEL INTEGRATED PERFORMANCE PRIMITIVES

Intel IPPs are a set of optimized cross-platform software functions that boost media processing performance by taking advantage of processor microarchitecture, algorithmic techniques, and instruction sets such as Streaming SIMD Extensions (SSE3). Figure 7 depicts IPP components that could significantly enhance PC applications such as Audio, Video, Image, Graphics, Speech, Math, Signal, and Cryptography and save developers from costly and time-consuming hand-coding and optimization.

Table 1 below depicts a list of speech coding samples built with Intel IPPs as the building blocks that are bit-exact with the standard.

Table 1: List of codecs supported by Intel IPPs

Speech Coding Samples	Windows*	Linux*
G.722.1	✓	✓
GSM/WMR WB / G.722.2	✓	✓
G.723.1	✓	✓
G.726	✓	✓
G.728	✓	✓
G.729	✓	✓
GSM-AMR	✓	✓
GSM-FR	✓	✓

For more information on the Intel IPP Application Programming Interface (API), please visit http://www.intel.com/cd/ids/developer/asmo-na/eng/dc/mobile/242763.htm.

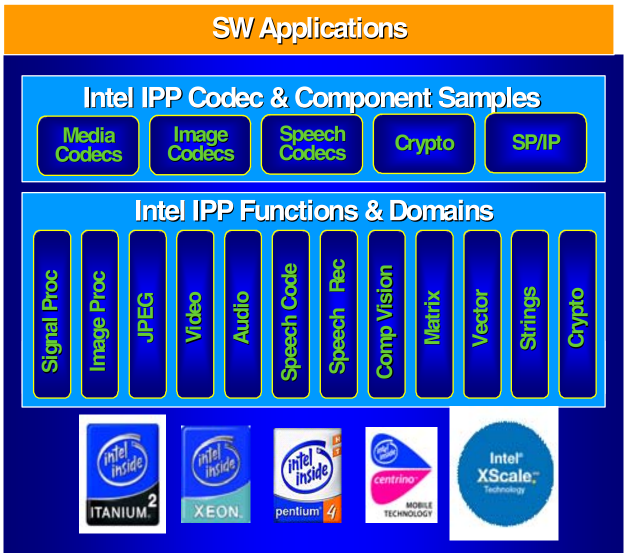

Figure 7: Software development with IPP

Figure 8 shows the performance gains of the highly optimized Intel IPP functions over compiled C code. See [Intel IPP Performance] for more information regarding Intel IPP performance. For more information about Intel IPP 5.0 including code samples, free evaluation copies, or to purchase copies visit [Intel® Integrated Performance Primitives 5.0].

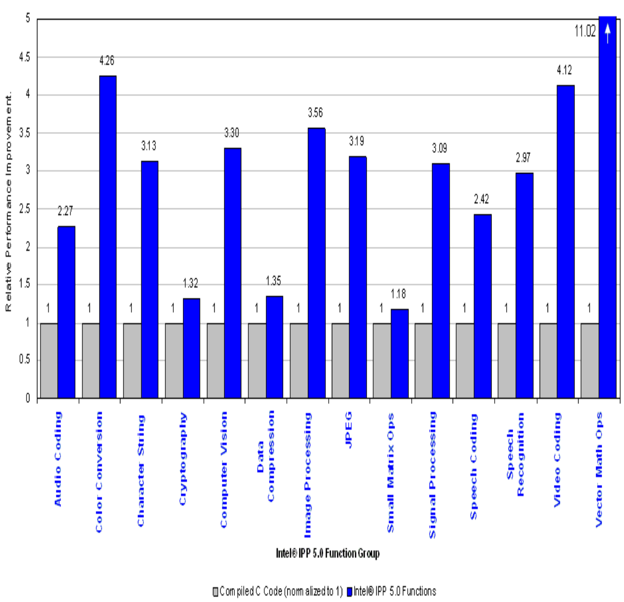

Figure 8: Intel IPP 5.0 Performance Primitives

HIGH-DEFINITION AUDIO AND ARRAY MICROPHONE FOR OPEN AUDIO

While other audio implementations have limited support for simple array microphones, Intel HD Audio supports larger array microphones. By increasing the size of the array microphone, users get incredibly clean input through better noise cancellation and beam forming. This produces a higher-quality input to voice recognition, VoIP, and other voice-driven activities.

An array microphone provides a speakerphone-like "open audio" usage solution for softphone VoIP calls. It does not require the user to wear headphones or a headset microphone; rather the user can speak directly into the device and listen to the platform speaker at the same time. The array microphone provides higher-quality audio and an improved user experience.

An array microphone is a set of multiple microphone elements integrated on the mobile platform to provide better microphone input quality and features, such as Acoustic Echo Cancellation (AEC) and repetitive noise attenuation and filtering or Noise Canceling .

Figure 9 depicts Intel Array Microphone architecture with four microphone elements.

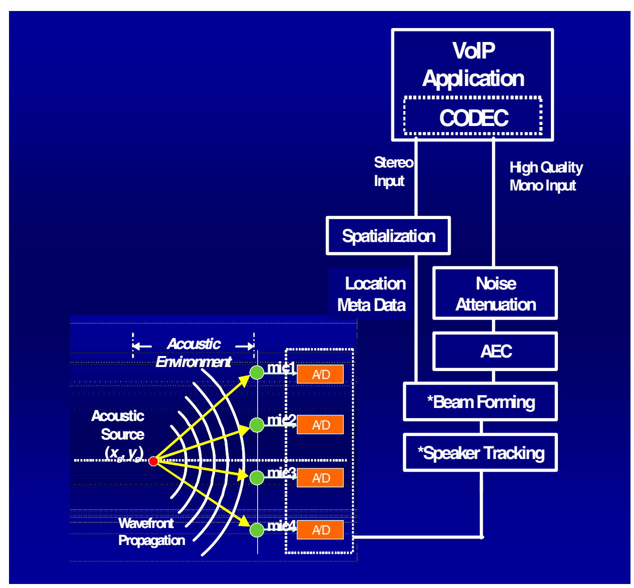

Figure 9: Intel Array Microphone architecture

Figure 9 depicts four key processes associated with Array Microphone processing. Speaker Tracking allows tracking of speaker source and direction to be able to replicate directionality at the rear endpoint if needed. Beam Forming implements speaker localization and sound attenuation. AEC ensures isolation between the speaker and the microphone for elimination of annoying echoes during the conversation. Finally, Repetitive Noise

Attenuation and Filtering ensures effective noise cancellation for clearer audio speech.

For more information about Intel High Definition Audio design and implementation please refer to http://www.intel.com/design/chipsets/hdaudio.htm.

Bluetooth Headset for Hands-Free VoIP

Bluetooth headsets are commonly used as an audio device with softphones. Bluetooth headsets enable users to have a hands-free VoIP conversation and to manage their connection, for example, accept/terminate a call, adjust volume, etc. A few softphone vendors have chosen to integrate Bluetooth headsets to enable ease of use with wireless headsets. The softphone and Bluetooth stack integration architecture on Microsoft Windows* operating system is depicted in Figure 10.

The Bluetooth stack is composed of the API interface; middleware protocols, such as the Service Discovery Protocol (SDP); Logical Link Control and Adaptation Protocol (L2CAP); RFComm (for serial communication); and Bluetooth profiles such as the headset profile, which is used for wireless headsets.

The softphones need to be integrated with the Bluetooth API to achieve service discovery and to interact with Bluetooth headset devices, i.e., be able to send and receive commands to and from Bluetooth headsets. Intel has worked closely with the Bluetooth stack and softphone vendors to define a common profile for VoIP control commands for headsets/handsets. Going forward, Intel will work with the Bluetooth stack vendors to define a common API for both headsets and handsets to enable softphones to interact with Bluetooth stacks from different vendors in the same way. This will significantly simplify the softphone integration efforts to support Bluetooth headsets for VoIP.

The coexistence of Bluetooth and WLAN radios on a notebook could impact the VoIP quality due to the interference between the two radios. Intel Centrino 2006 platforms implement the Wireless Co-existence Solution (WCS) to solve the problem of Bluetooth-WLAN coexistence by sending WLAN channels to Bluetooth. The WCS describes the interface between two radios running simultaneously in the same laptop to transfer information on the channel occupied by WLAN to a Bluetooth NIC. As a result, Bluetooth will skip the

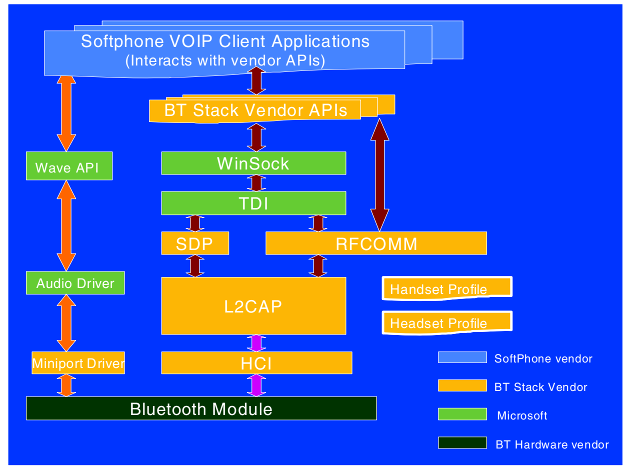

Figure 10: Softphone and Bluetooth stack integration

WLAN channel for non-critical events, and WLAN performance will be recovered in such cases. This solution significantly increases the throughput of WLAN when Bluetooth is presented. Additionally, the WCS specification provides better protection of critical Bluetooth communications (voice for example). This is accomplished by the following method:

- When Bluetooth expects to receive or transmit a high-priority packet on a Bluetooth channel that is located inside the WLAN channel, Bluetooth raises a BT_Priority signal.
- WLAN radio will then defer or kill its transmission if a BT_Priority signal is raised to avoid collision with Bluetooth packets.

For more information about Intel WCS solutions, please refer to [6].

CONCLUSIONS AND FUTURE WORK

In this paper, we presented Intel's Seamless Collaboration Architecture for VoIP on WLANs. We described our QoS architecture for VoIP that enables softphone applications to take advantage of QoS features on the Intel Centrino mobile technology notebook over WLANs. We also described other hardware/software solution ingredients, such as array microphones, IPPs, and the Bluetooth wireless coexistence solution, all of which make Intel Centrino mobile technology notebooks provide a best-in-class experience for VoIP over WLANs.

In addition to VoIP, collaboration usages that involve simultaneous audio, video, and data conferencing are becoming common in enterprises. Our future work is directed at extending this architecture to support multimedia collaboration applications. This includes hardware and software solutions for video peripherals like embedded cameras and remote handheld devices. Other usage models like multiparty ad hoc conferencing and IPP

* Other names and brands may be claimed as the property of others.

extensions for video conferencing are also part of this effort.

ACKNOWLEDGMENTS

We thank Ajay Bakre for his invaluable contributions to this work.

We also thank Duncan Glendinning, Selim Aissi, Ajay Gupta, Venkat Gokulrangan, Vijay Rao, Charlie Lasswell, and Sanjay Bakshi for the quality time they spent on reviewing this document and providing invaluable feedback.

REFERENCES

[1] *IEEE. Part 11*, "Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications," ANSI/IEEE Std 802.11, 1999 Edition, (ISO/IEC 8802-11:1999(E)).

[2] *IEEE. Part 11*, "Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications: Medium Access Control (MAC) Enhancements for Quality of Service (QoS), ANSI/IEEE Std 802.11e," Draft 13.0, January 2005.

[3] *Wi-Fi Alliance*, "WMM including APSD Specification," *WMM-Specification-draft-050720c.doc*.

[4] Rajeev Muralidhar, Ajay Bakre, Hani Elgebaly, "On the Performance of VoIP in EDCA-based IEEE 802.11 WLANs," in *Proceedings of World Wireless Congress, 2005*.

[5] http://www.skype.com.*

[6] Gordon Chinn et al., "Next-Generation PC Platform Built on Intel® Centrino® Mobile Technology New Usage Models," *Intel Technology Journal*, Q1'05.

[7] http://www.microsoft.com/msj/0599/LayeredService/LayeredService.aspx*

[8] http://www.itu.int/rec/recommendation.asp?type=products&parent=T-REC-g*

[9] Bernstein Research Call, Market Survey, "U.S. Telecom: Corporate Shifts to IP Driving Losing Battle for Carriers Within IT Budgets but Price Decays Moderating." May 2005.

AUTHORS' BIOGRAPHIES

Hani Elgebaly is the enterprise software architecture Manager in Intel's Mobile Platform Group. Hani's team is responsible for wireless and IP networking software technologies for VoIP, manageability, security, and mobile connectivity. Hani has been with Intel for ten years where he held multiple engineering management and lead architect positions at Intel labs (Corporate Technology Group) and with videoconferencing and IP telephony product groups. He has over 14 pending US patents and 15 journal and conference publications. He is also a senior IEEE member. He received his Ph.D. degree in Computer Science from the University of Victoria, British Columbia, Canada, his M.S. degree in Computational Science from the University of Saskatchewan, Canada, and his B.Sc. degree in Electrical and Computer Engineering with high honors from Cairo University, Egypt. His e-mail is hani.elgebaly at intel.com.

Farid Adrangi is an architecture manager in Intel's Mobile Platforms Architecture and Development Group where he leads architecture and development of Seamless Collaboration, one of the pillars in digital office initiatives. Farid's team is responsible for platform architecture and ingredients to deliver best-in-class audio/video/data conferencing experiences over wireless and wired networks. Prior to this, he worked in Intel's Research Communication Technology Lab as a lead architect on mobility and WLAN/WiMAX architecture and standardization efforts in 3GPP, IETF, and GSMA. Farid has two issued and over 15 pending US patents and 3 IETF RFC publications, and he has earned a GSMA Outstanding Award for his contribution to the Interworking WLAN Roaming architecture with 3GPP. He received a B.S. degree in Computer Science with high honors from Portland State University. He has been with Intel since 1994. His e-mail address is farid.adrangi at intel.com.

Rajeev Muralidhar is a senior software engineer in Intel's Mobile Platforms Architecture and Development Group where he works on architecture, research, and development of next-generation wireless software for Intel Centrino mobile technology platforms. Prior to this, he worked in Intel Research's Network Architecture Lab on architectures for network processor-based platforms for the Internet, including standardization efforts at the Network Processing Forum. Rajeev has over 6 pending US patents and over 10 journal and conference publications. He has an M.S. degree from Rutgers University, and a B.E. degree from NIT, India; both in Computer Engineering. He has been with Intel since 2000. His e-mail address is Rajeev.D.Muralidhar at intel.com.

Lakshmi Ramachandran is a staff architect in Intel's Mobile Platform Architecture and Development Group where she leads architecture, research, and development of next-generation wireless software for notebooks based on Intel Centrino mobile technology. In her 5 years at Intel, her areas of focus have included network management for SS7 stacks and network processors. Lakshmi has several research publications and over 5 pending US patents. She is also a co-author of a handbook on Mobile Wireless Networking, published by John Wiley and Sons. Her current focus areas are multimedia QoS for WLANs and new usage models for mobile platforms. She holds a Masters degree in Computer Science from the

Indian Institute of Science. Her e-mail address is Lakshmi.Ramachandran at intel.com.

Manish Nair is a senior software engineer in Intel's Mobile Platforms Architecture and Development Group where he works on architecture, research, and development of next-generation wireless software for Intel Centrino mobile technology platforms. Manish received his B.Tech degree from the College of Engineering, Trivandrum, India in Computer Science and Engineering. He is currently pursuing research in IP Security at the Indian Institute of Science, Bangalore. His e-mail is manish.nair at intel.com.

THIS PAGE INTENTIONALLY LEFT BLANK

For further information visit:

developer.intel.com/technology/itj/index.htm