Final Project: Skype

Frank Bulk
IST 656
Instructor: Peter Morrissey
May 5, 2004

Section 1: Background Analysis

Skype is both a product and the name of the company. As a product, it is the software that "gives people new power to affordably stay in touch with their friends and family by taking advantage of their technology and connectivity investments"[1], according to the company's CEO, Niklas Zennström. Simply stated, Skype allows people to make phone calls over the internet, similar to existing solutions offered by Net2Phone and MSN Messenger. As a company, Skype's mission is to provide "a simple, reliable and friendly communications tool that just works"[2]. Unlike VoIP companies such as Vonage, or services such as AT&T's Call-Vantage which require authentication servers and managed points of presence around the country, Skype provides free software and service requiring an extremely minimal support infrastructure in spite of the Internet's dynamic environment.

Skype released its product in the fall of 2003, but it started as company and product development in the spring, 6 months earlier, according to a CNET news.com article.[3] According to the company's website, Niklas Zennström is CEO and co-founder of Skype and Janus Friis acts as VP of Strategy and co-founder. Their previous company, KaZaA, had taken the peer to peer (P2P) file-sharing world by storm in 2000 when the previous incumbent Napster was shut down for legal reasons. Zennström is the elder of the two by ten years, and acquired dual degrees in Business Administration and Engineering Physics from Uppsala University in Sweden. Friis, on the other hand, dropped out of high school, according to the Fortune Magazine's article "Catch us if you can".[4] Friis gained his technical experience working for several ISPs. Zennström has had extensive business experience working 8 years for a Tele2, a telecom company in his native Sweden. Zennström first met Friis in 1997 when he hired him to work at Cybercity, one of Denmark's first ISPs. In late 1999 Friis convinced Zennström to leave Tele2 to start a company together, a project which eventually turned into a technology called FastTrack. Rather than be the only user of this P2P file-sharing technology, they decided to license it, two of which became services call Morpheus and Grokster. Zennström and Friis hired other programmers to write their own client, KaZaA.

---

[1] *The Company*. 20 Apr. 2004 <http://ui.skype.com/company.html>
[2] *The Company*. 20 Apr. 2004 <http://ui.skype.com/company.html>
[3] Charny, Ben. "Why VoIP is music to Kazaa's ear." *CNET news.com* 2 Dec. 2003. 26 Apr 2004 <http://news.com.com/2102-1082_3-5074558.html>
[4] Roth, Daniel. "Catch us if you can" *Fortune* 9 Feb. 2004. Business Source Elite. EBSCOhost. 19 Apr. 2004

KaZaA was released in March 2001,[5] a few months before RIAA (Recording Industry Association of America) shut down Napster in July. The timing was very fortunate for KaZaA's founders, and their products became phenomenally successful. But the public's use of their software to exchange music, programs, and other more dubious material, made it impossible for these two young entrepreneurs to broker any licensing deals with American groups such as RIAA or the Dutch group, Buma/Stemra. Eventually a Dutch court ordered KaZaA shut down in early 2002. As a result, Zennström and Friis sold KaZaA and kazaa.com, and transferred the P2P technology called FastTrack to a company called Joltid. As expected, RIAA continues to pursue the case, attempting to serve subpoenas wherever their legal jurisdiction extended. In order to avoid the long arm of the American law, the two founders have tried to keep their whereabouts somewhat hidden.

Out of this notoriety and history the two turned to another project: Skype. Instead of transferring files, Skype was designed to make voice calls over the Internet. Their expired P2P competitor, Napster, had received funding from Bertelsmann to the tune of $50 million.[6] When the pair began KaZaA they weren't able to get any venture capital funding, forcing Zennström to fund it with his savings own savings, according to the Fortune article. Well aware that funding allows for expansion of products and services, once Skype's beta client became available in the fall of 2003, the company quickly courted several venture capital (VC) firms. Bessemer Venture Partners led the 'A' round of financing for an estimated USD $25 million of money in mid-December 2003, along with Mangrove Capital Partners. Draper Investment Company, a company focused on early-stage European investments also contributed, as well as one of its key directors, Tim Draper[7] [8]. In Skype's 'B' round of funding for USD $18.8 million, the Draper family supported Skype once again using another one of their established VC firms, Draper Fisher Jurvetson, for USD $9 million[9]. Several other individuals and companies, as well as Index Ventures, joined in that

---

[5] "KaZaA;" *Wikipedia*.com Wikipedia: the free encyclopedia. 1 May 2004 <http://en.wikipedia.org/wiki/Kazaa>
[6] Gibney, Frank. "Middelhoff's Vision." *Time Europe* 13 Nov. 2000. 1 May 2004 <http://www.time.com/time/europe/magazine/2000/1113/bertelsmann.html>
[7] Press release. "Bessemer Venture Partners Funds Skype Technologies, S.A". 22 Dec. 2003. 30 Apr 2004 <http://www.bvp.com/news/pr.asp?id=622>
[8] AltAssets Monthly Roundup. 28 Feb. 2002. 30 Apr. 2004 <http://www.altassets.com/roundup/arc/2002/nz666.php>
[9] Charny, Ben. "Money brings Skype closer to mainstream." *CNET news.com* 15 Mar. 2004. 1 May 2004 <http://news.com.com/2100-7352-5173238.html>

‘B’ round. With over $40 million in total funding, Skype is on the way to be able to invest major funds in future business development. Because their current software has been distributed for free and costs nothing to use, Skype will turn to additional services such as voice mail and gateway connectivity to the PSTN. Their website also refers to new connectivity platforms such as mobile devices and handsets. They have made their first move into this space with their recent release of PocketSkype that runs on faster version of the Pocket PC 2003. According to a CNET news.com article, they chose Plantronics as their preferred provider of headsets. In a joint marketing arrangement, Plantronics will give away two months of premium service (those details still unknown) to those who purchase their headset.[10] In this way Skype receives additional exposure as well as another set of users.

Until this time Skype’s product has focused clearly on the consumer market, at least in the initial releases. According to a CNET news.com interview with Zennström: “We’re starting with individuals. We’re doing this bottom up. It’s grassroots for businesses.”[11] According to Friis in a different interview with the same electronics news source, Skype provides good sound quality, easy setup and configuration, and no cost[12]. These are all legitimate concerns that consumers of other products such as Net2Phone or NetMeeting have experienced.

The challenge with the consumer market is that it will not generate the sums of money that an agreement with a large multi-national company can. Skype’s internet communications infrastructure is based on a peer-to-peer model for the best voice quality, but doing so makes call progress monitoring much more difficult. Although aware of the issues, Skype’s business model, according to the same interview with Zennström, is to sell value-added services such as voicemail and gateway connections to the PSTN.[13] In the

---

[10] Charny, Ben. “Skype plugs in Plantronics headsets.” *CNET news.com* 18 Mar. 2004. 1 May 2004 <http://news.com.com/2110-7352-5175515.html>
[11] McCullagh, Declan. “Skype's VoIP ambitions.” *CNET news.com* 2 Dec. 2004. 24 Apr. 2004 <http://news.com.com/2008-7352-5112783.html>
[12] Charny, Ben. “Why VoIP is music to Kazaa's ear.” *CNET news.com* 2 Dec. 2003. 26 Apr. 2004 <http://news.com.com/2102-1082_3-5074558.html>
[13] McCullagh, Declan. “Skype's VoIP ambitions.” *CNET news.com* 2 Dec. 2004. 24 Apr. 2004 <http://news.com.com/2008-7352-5112783.html>

latest releases they have added 5-way conference calling, according to the San Francisco Chronicle.[14] This is their first small step into the business space.

Skype's future work will be expanding the current client and backend infrastructure to include services such as voice mail and PSTN gateways. Thought has been given to connecting to SIP services, and Jeff Pulver mentions in a blog that he has had discussions with Skype on how his FWD service can interconnect with Skype. Skype will be able to use VC money to achieve their goals of hiring people for their London office and investing in the necessary equipment, mentions the Fortune magazine article. This would put their functionality on par with services such as Free World Dialup, LibreTel, and Vonage, although each of these has their own unique characteristics, and none resembles Skype's peer to peer communication technology. Skype is also working on Linux and Mac versions of their software, according to an interview with Zennström.[15] [16] Agreements have also been made with several manufacturers of USB phones that will allow a telephone-like device to be plugged into the computer and interact with Skype client hosted on the computer to make and receive calls. The same San Francisco Chronicle article refers to an USB-based DECT cordless phone adapter, specifically the Siemens Gigaset M34, which will allow any DECT cordless phone to use Skype.

Section 2: Practical Functionality

The 7 MB download installed very easily. Much like any other instant-messaging program, there is a dialog box to login or create an account. Once I created an account on their site I easily logged in. Unlike AOL Instant Messenger, which has sold the application's GUI real estate for advertising, Skype's website points out that that they have made an explicit commitment to stay away from such ad-ware.

Although I made an informal call on the local LAN, my two official local calls were from the offices of CENT to two Road Runner customers. Calls Number 1 & 3 both were 16 hops away from me, and both were behind NAT gateways. In spite of the fact that we

---

[14] Sandred, Jan. "Skype's not yet up to all the hype" *San Francisco Chronicle on the Web* 15 Apr. 2004. 20 Apr 2004 <http://www.sfgate.com/cgi-bin/article.cgi?file=/chronicle/archive/2004/04/05/BUGF75VUFI1.DTL>

[15] "Skype Offers Free Conference Calling" *theMAURITIUS.COM* 25 Feb. 2004. 1 May 2004 <http://www.themauritius.com/article5421.html>

[16] McCullagh, Declan. "Skype's VoIP ambitions." *CNET news.com* 2 Dec. 2004. 24 Apr. 2004 <http://news.com.com/2008-7352-5112783.html>

both were behind NAT gateways, the calls setup occurred very quickly and without any problems. The sound quality on both calls was excellent and used the higher bandwidth iSAC codec. For Call Number 1 my wife used our cheap handheld microphone, so her voice level was low, but everything was still very understandable. I couldn't perceive any latency and the low ping times, about 25 milliseconds, bore that out. When I performed a little more analysis on the call using NetPeeker's software,[17] I found out that it was not really a direct call. Instead, her call was being routed from Syracuse University to France! Because I did not analyze the other side of the call I do not know if her traffic came in directly from France, or perhaps via another hop. Call Number 3 to Frank Robinson also enjoyed good voice quality, contributed in part by his quality headset. Our communication was direct (substantiated by NetPeeker's software), but I could not perceive an increase in quality or decrease in latency as compared to Call Number 1. I suspect it was using the "UDP Hole Punching Algorithm" successfully with Frank Robinson's SOHO router, but not with my home one. We found that when we talked simultaneously to each other, one reciting the alphabet and the other counting, that the call sputtered. It seems that it is not as full-duplex as we thought. Further testing would be required to find out if it was the sound card or Skype program that caused this glitch.

Call Number 2 was to my brother in northwest Iowa. He was behind a firewall, and I behind a NAT box and a firewall, but again, we had no problems setting up the call. He had a medium quality headset on, while I used a cheap boom microphone with my laptop's speakers. The quality was good, even better than a landline connection – again, an iSAC codec was in use. We had no problems throughout the call, and even though a few packets were dropped, both in the formal part of the call, as well as later in other conversations, there was no perceivable loss in quality.

It was a little more difficult to find someone with whom I could make an international call and who would share their IP address with me. The first was an informal call to someone in the Netherlands. He was not using a headset but a simple microphone and computer speakers. The call quality was very good. The second, more formal call, was to someone in Jamaica. He was using a DSL line running 64 kbps symmetrically. He would not share his IP address with me (he called his ISP to find out if it was safe to do so).

[17] Go to http://www.net-peeker.com to download the utility

His call was not direct but proxied. That call did exhibit the worst quality I had heard so far, and it still used the higher-bandwidth ISAC codec. I made several attempts to speak to others in foreign countries, but many used Skype just for chat and were in the process of getting headphones, or just ignored me.

Presence was updated very quickly. I observed presence changes within a second in the lab and at home. This is replicable to WAN situations, as well, because presence is monitored by supernodes, so even two Skype clients in the same IP network will need to be informed by a connecting supernode.

Frank Robinson and I attempted to use NetPeeker's abilities to shape traffic and reduce it to just a 33.6 kbps flow, in order to simulate a dial-up connection, but unfortunately NetPeeker seems to shape only TCP, and does not do very well with UDP. Because a lot of encryption and codification was going on, I also tracked my system's CPU usage. The Skype.exe process is assigned High priority (13), two steps above Normal (8). It's like QoS (quality of service) for the CPU! One call on my processor (Pentium M 1.4 GHz) used at around 20 to 25% of the CPU, according to Windows TaskManger and SysInternal's Process Explorer. When I added a third person into the conversation for a 3-way conference call the utilization increased to about 40 to 45%, suggesting that the CPU utilization is approximately linear. It's reasonable to assume that I would barely be able to make a 5-way conference, in which I would be hosting 4 of the conversations. Sending an instant message took about 2-3% of the processor. It also used about 12 to 14 MB of memory, and that stayed stable, in spite of the numbers of calls or activity.

Section 3: Architecture

Skype is based on a P2P infrastructure, so little or no data is stored on a central server. According to a CNET news.com interview with co-founder, Friis, "True P2P software creates a network through which all clients join together dynamically to help each other route traffic and store information."[18] There is the client or node, supernodes, and Global Index server. As Figure 1 shows, the clients are the gray blocks connected to red

[18] McCullagh, Declan. "Skype's VoIP ambitions." *CNET news.com* 2 Dec. 2004. 24 Apr. 2004 <http://news.com.com/2008-7352-5112783.html>

supernodes. The supernodes are all connected to each other, and as a group, they can be considered a Global Index server (the surrounding blue enclosure).

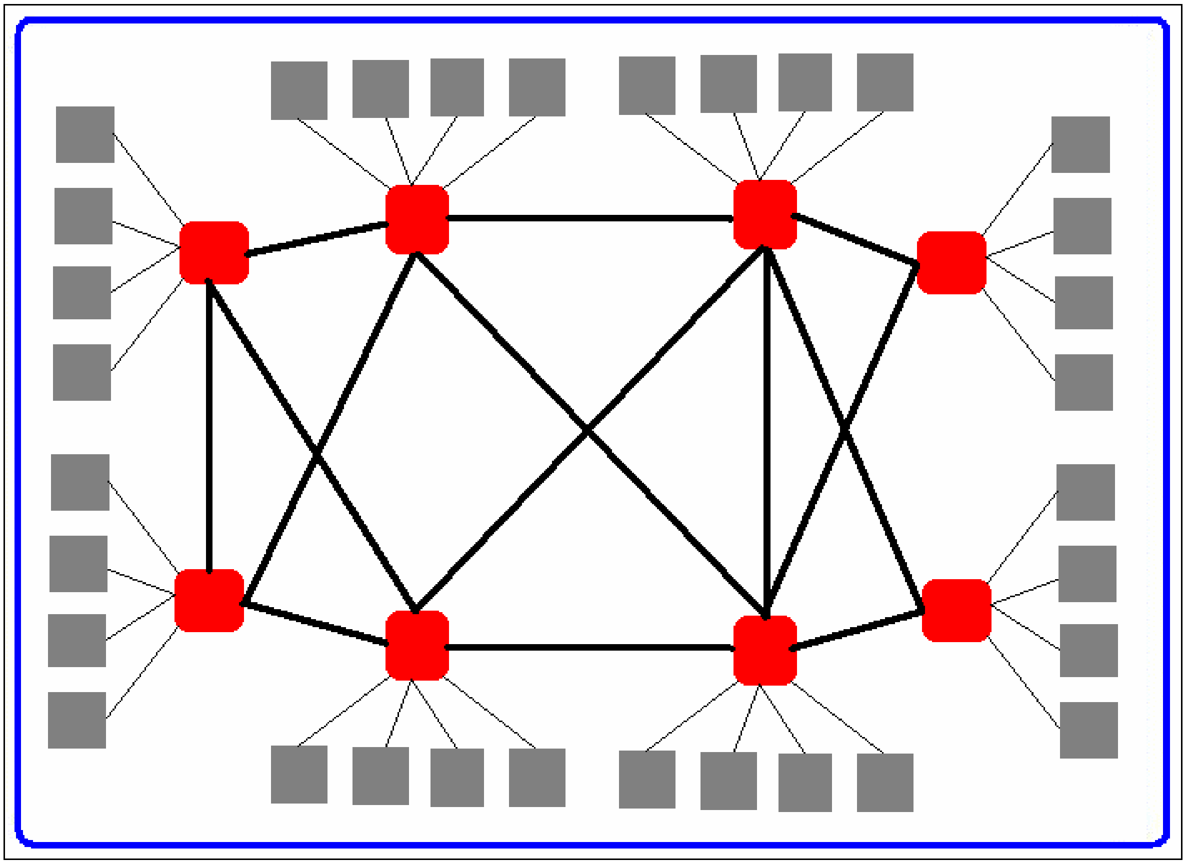

Figure 1: Model of Skype Infrastructure

Each additional client that logs on keeps track of several 'adjacent' supernodes. Just as there are core routers on the internet, there is equivalent functionality in the supernodes of the Skype network. According to a CNET news.com interview, each supernode handles several hundred clients, based on their memory and bandwidth availability, as well as their 'uptime'. These supernodes are always on open, routable IP addresses. In other words, Skype users with a cable/DSL router-NAT box will have no hope of ever becoming a supernode.

The Global Index (GI) server has been referred in a disparaging blog entitled "Why Skype is No Different"[19] as a physical device, but Skype's more detailed FAQ[20] suggests that it is more a hierarchical network of supernodes. I have diagrammed it as partial mesh,

[19] Rao, Aswath "Why Skype is No Different" *VoIP Daily* 27 Mar. 2004. 30 Apr. 2004 <http://gigaom.com/thevoipdaily/archives/2004/03/why_skype_is_no_different.html>

[20] *P2P Telephony Explained*. 30 Apr. 2004 <http://www.skype.com/skype_p2pexplained.html>

because there is nothing in my research that suggests that some supernodes only communicate to other supernodes. The GI handles communication between supernodes in such a way that each node can quickly search the entire network and keep track of presence with minimal latency.

The first time a newly installed client runs it connects to some hard coded IP addresses, probably servers run by Skype or their associates, as at least one blog about a Skype researcher suggests.[21] Those server(s) provide a dynamic and active list of supernodes which the client stores for future reference and connections. In this way, Skype, the company, does not need to host many machines for directory services, but just seeds the initial connection this P2P mesh.

According to a CNET news.com interview with creator Niklas Zennström, each client sets up a connection to four or five standby paths, dynamically choosing the one with the best bandwidth or lowest latency.[22] According to a separate interview, each client is always establishes a TCP/IP connection to a supernode [23]. This was verified in my analysis. What that means is that if a connection can't be setup via one path, an alternative path will be used. This is done to circumvent the problems that are introduced when one or both sides are behind a firewall or NAT device.

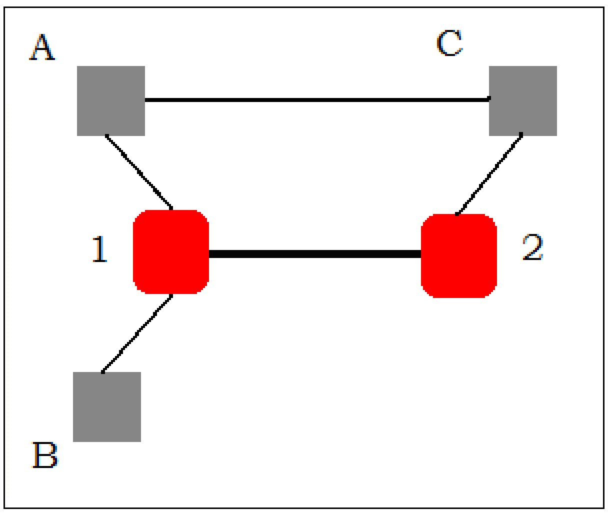

Figure 2: Calling diagrams

[21] Schulzrinne, Henning "Henning Schulzrinne on Skype" *techblurbs* 21 Jan. 2004. 1 May 2004 <http://stefan.blogs.com/techblurbs/2004/01/henning_schulzr.html>
[22] McCullagh, Declan. "Skype's VoIP ambitions." *CNET news.com* 2 Dec. 2004. 24 Apr. 2004 <http://news.com.com/2008-7352-5112783.html>
[23] Peter. "How does Skype get through Firewalls and NAT Routers?" *The Register*. 3 Oct. 2003. 30 Apr. 2004 <http://www.theregister.co.uk/2003/10/08/how_does_skype_get_through/>

In a typical point to point setup, if both clients are behind NAT and no port forwarding has been configured, node A will not be able to communicate to node B because the NAT device that's on the receiving end of the call is not expecting an incoming connection and so there is no translation from public to private socket to be made. But in the case of Skype, node A might talk to supernode 1, as does node B, so supernode 1 can mediate the call, or, as Zennström calls it, synchronized. In fact, he says that most of the times there isn't a direct connection between nodes. If the nodes don't talk to the same supernode, as is the case between nodes B and C, then the supernodes 1 and 2 communicate between each other. If the firewall does allow incoming connections the alternate end will initiate an UDP connection to the open client and communication will proceed directly, as is the case between nodes A and C.

My packet analysis showed that Skype's chat support uses the supernodes and does not use direct communication. The text traffic is likely so low in comparison to voice traffic that it's not worth the effort to create a direct link. The instant messaging interface is a handy way to see if someone is online, and to initiate a text dialog before moving over to audio.

While account names, passwords, and email addresses may be centralized by Skype, according to a blog,[24] little else is. Contacts lists aren't stored on the network, so installing the client on another computer requires requesting authorizations with your friends each time again. Each Skype user's profile (which includes information such as location, language, birthdate, and contact information) must be re-uploaded to the directory if the 72 hour expiry period has passed, according to the Skype FAQ.[25] In the same way that calls are connected across via supernodes, there is speculation that searching for contacts is accomplished by performing a broadcast search within a fully distributed index, namely the supernodes and Global Index.[26] My search tests showed that the same query returned different results each time, even when executed seconds apart. This verifies the distributed and 'laziness' of the search algorithm, and that the search is not a comprehensive result set, but just what could be returned within a specific time period.

---

[24] Henshall, Stuart. "Experimenting on Skype reveals more" *Unbound Spiral* 26 Sep. 2003. 30 Apr. 2004 <http://www.henshall.com/blog/archives/000427.html>
[25] *Frequently Asked Questions*. 30 Apr. 2004 <http://www.skype.com/help_faq.html>
[26] http://zgp.org/pipermail/p2p-hackers/2004-March/001785.html

In spite of Skype's ability to cross most NAT and firewall boundaries, it can not traverse proxy servers or authenticated firewalls. Although it seems to be functionality that could easily be built into the client, adding it might divert from Skype's philosophy that once the software is installed, it should be "dead-simple to use".[27]

There is currently no capability to communicate with other SIP-based devices. When asked about SIP support during a BoardWatch interview, Zennström affirmed that Skype would not use SIP.[28] [29] In fact, the Skype development team specifically avoided SIP because of NAT and firewalls problems. In a previously mentioned CNET news.com interview, Friis says that SIP "could not do what we wanted it to do."[30] Nevertheless, Friis does say that the future of Skype includes possible interoperability with Skype.

A connection to the 'legacy' PSTN is high on the list of Skype's development goals, according to a CNET news.com interview,[31] and an eWEEK article confirms that they "believe in interoperabilitiy… and are open to discussion with other companies."[32]

The reason that Skype conversations have such good sound quality is because it uses voice codecs such as iSAC that have an 8 kHz frequency range,[33] as compared to roughly 3.4 kHz range provided by the PSTN.[34] Skype also uses the same iLBC codec that Swedish-based Global IP sound freely provides other VoIP clients such as X-Ten, according to Aswath Rao,[35] and used by Asterisk in their soft PBX. iLBC is a royalty-free, low-bit rate codec that is currently a work item (draft-ietf-avt-rtp-ilbc-04.txt) for the IETF's audio/video transport Working Group. According to the wiki at German website, this codec

---

[27] *P2P Telephony Explained*. 30 Apr. 2004 <http://www.skype.com/skype_p2pexplained.html>
[28] Maitland, Jo. "KaZaA Founder, Niklas Zennström." *Boardwatch* 16 Jul. 2003. 24 Apr. 2004 <http://www.boardwatch.com/document.asp?doc_id=36886>
[29] http://www.lightreading.com/document.asp?site=lightreading&doc_id=36886&page_number=5
[30] Charny, Ben. "Why VoIP is music to Kazaa's ear." *CNET news.com* 2 Dec. 2003. 26 Apr 2004 <http://news.com.com/2102-1082_3-5074558.html>
[31] McCullagh, Declan. "Skype's VoIP ambitions." *CNET news.com* 2 Dec. 2003. 24 Apr. 2004 <http://news.com.com/2008-7352-5112783.html>
[32] Livingston, Brian. "Beware Skype's Hype; focus on sip-compliant internet calling instead." *eWeek* 1 Dec. 2003. Lexis-Nexis. 19 Apr. 2004 <http://www.lexis-nexis.com>
[33] "GIPS iSAC."Global IP Sound 1 May 2004 <http://www.globalipsound.com/pdf/gips_iSAC.pdf>
[34] Isomäki, Markus. "Peer-to-Peer Communication Services in the Internet", Nokia Research Center <http://keskus.hut.fi/opetus/s38030/F03/Report-p2p-spam-2003.pdf>
[35] Rao, Aswath "Why Skype is No Different" *VoIP Daily* 27 Mar. 2004. 30 Apr. 2004 <http://gigaom.com/thevoipdaily/archives/2004/03/why_skype_is_no_different.html>

has been designed for internet telephony applications and degrades gracefully when frames are lost.[36]

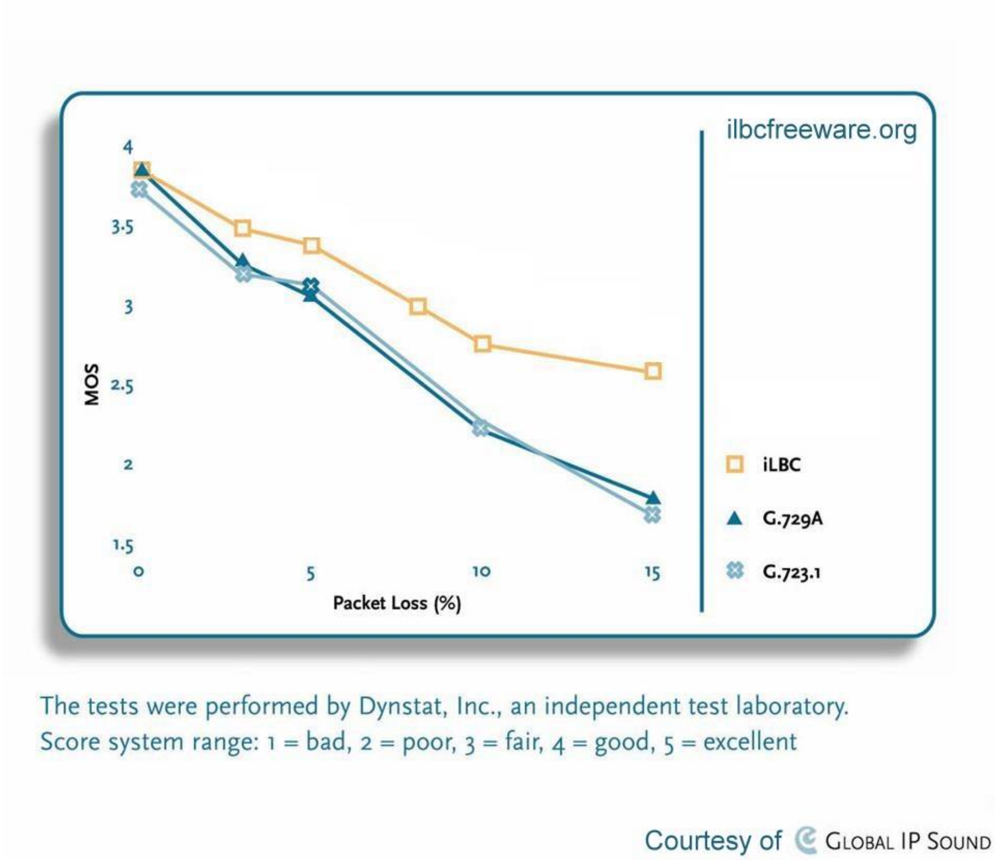

Figure 3: Codec Comparison

As the graph above shows, iLBC's beginning MOS (mean opinion score) starts above that of G.729A and G.723.1, codecs commonly used for internet telephony, and does not degrade as quickly as packet loss increases. When using a voice frame size of 30 ms it generates a 13.33 kbps stream (399 bits, or 50 bytes), and for a frame size of 20 ms it generates a 15.2 kbps stream (303 bits, or 38 bytes). The higher bit rate for the smaller audio frame can be attributed to the fact that more bits are used, resulting in better sound quality and making it more resilient to packet loss (missing one 30 ms frame is worse than missing a 20 ms frame) with lower computational delay.[37]

---

36 "Global IP Sound iLBC low bit rate codec." 10 Apr. 2004. 1 May 2004 <http://www.my-voip.de/wiki/pmwiki.php/Main/ILBC?>

37 "News: Announcements." 26 Feb. 2004. 1 May 2004 <http://www.ilbcfreeware.org/news.html>

Section 4: Protocol Analysis

The resulting packet capture from a Skype call was unique. I captured a call setup during testing with Joshua Bulk. The first few packets are TCP and seem to talk to the supernodes. After that there were a burst of five UDP packets to IP addresses that were seen only once more. They seemed to be like Skype ping packets, used to check the quality and bandwidth of the connection.[38] There were a few TCP responses from the first connections, and then a group of three UDP packets, also ping-like. After that there are a few more TCP packet exchanges with the same supernodes that were used at the beginning of the packet trace. Presumably the call was being negotiated. Subsequently the call began with a long stream of small UDP packets, averaging 150 bytes in size, directly to Joshua's computer. Each UDP packet encapsulated about 30 ms of voice, because each packet was spaced in intervals of about 30 ms. The call was set up took about 3.3 seconds.

For the incoming call I had Joshua call me back. We synchronized our computer clocks and I started capturing packets three seconds before he initiated the call. A stream of TCP connections came in from my associated supernodes, and about 3 seconds later the UDP stream from Joshua came in. Because there weren't any UDP pings adding more traffic pairs, it was easy to see in Packetyzer's connections tab that my computer was talking to 4 different supernodes.

Using NetPeeker v2.5 I also analyzed a live call. This program shows the live connections for each program, including IP address and port number, as well as the rate of traffic traversing over each session. When I reviewed the data, I was amazed by number of active UDP and TCP sessions generated by Skype. I placed a call from the CENT office at Syracuse University to my home PC that's connected to RoadRunner/TimeWarner. Because both sides of the connected are NATed, Skype used a supernode to connect the two end points. Over the following 135 minutes my proxy changed four times, from France, Taiwan, Japan, and then back to Taiwan. I found this to be incredible, considering the two computers were only a mile away from each other. That shows that Skype's call routing algorithms are probably not very sophisticated and optimized, or, that the numbers of supernodes are so few that the options are limited. I also found that Skype was continually

[38] Isomäki, Markus. "Peer-to-Peer Communication Services in the Internet", Nokia Research Center, p. 37 http://keskus.hut.fi/opetus/s38030/F03/Report-p2p-spam-2003.pdf>

opening additional low or no traffic UDP sessions to various IP address, usually nine of them. Three or four of these were added every few minutes, and an equal amount is dropped. From the Figure 4 it's also clear that the Skype client maintains TCP in addition to the existing UDP connections to about half of the supernodes. There is also a TCP connection to ui.skype.com, which is likely a user interface port. This web-specific port probably relates to Skype's use of Internet Explorer temporary files, as mentioned in Isomaki's article. It was not clear to my why Skype would need to write to the browser's temporary files.

The use of several simultaneous connections always open to supernodes allows the Skype client to quickly move away from worsening connection to a better one, without requiring the time to find a better route while the call degrades.

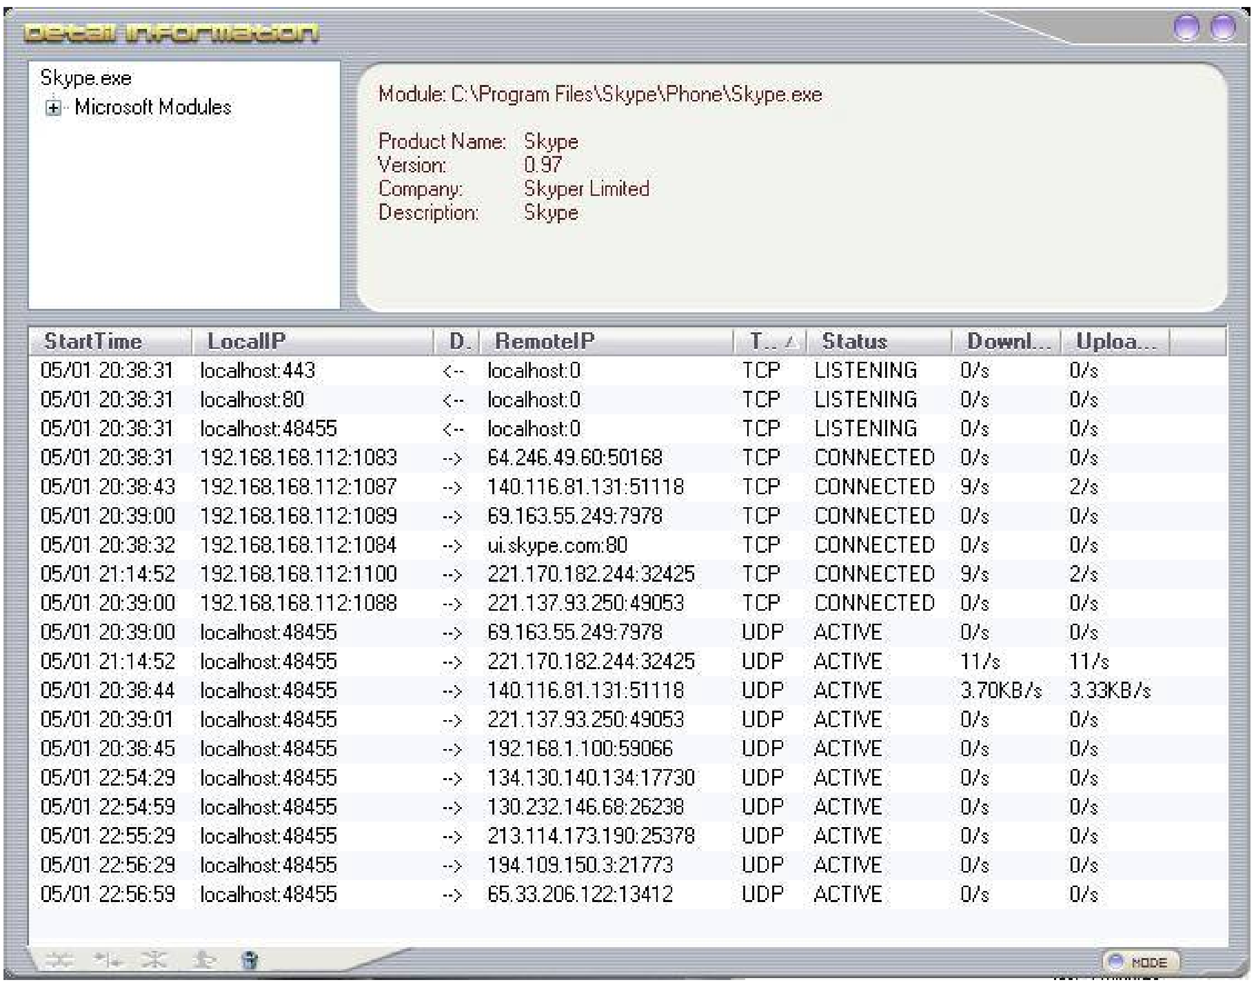

StartTime	LocalIP	D..	RemoteIP	T..	Status	Downl...	Uploa...
05/01 20:38:31	localhost:443	<--	localhost:0	TCP	LISTENING	0/s	0/s
05/01 20:38:31	localhost:80	<--	localhost:0	TCP	LISTENING	0/s	0/s
05/01 20:38:31	localhost:48455	<--	localhost:0	TCP	LISTENING	0/s	0/s
05/01 20:38:31	192.168.168.112:1083	-->	64.246.49.60:50168	TCP	CONNECTED	0/s	0/s
05/01 20:38:43	192.168.168.112:1087	-->	140.116.81.131:51118	TCP	CONNECTED	9/s	2/s
05/01 20:39:00	192.168.168.112:1089	-->	69.163.55.249:7978	TCP	CONNECTED	0/s	0/s
05/01 20:38:32	192.168.168.112:1084	-->	ui.skype.com:80	TCP	CONNECTED	0/s	0/s
05/01 21:14:52	192.168.168.112:1100	-->	221.170.182.244:32425	TCP	CONNECTED	9/s	2/s
05/01 20:39:00	192.168.168.112:1088	-->	221.137.93.250:49053	TCP	CONNECTED	0/s	0/s
05/01 20:39:00	localhost:48455	-->	69.163.55.249:7978	UDP	ACTIVE	0/s	0/s
05/01 21:14:52	localhost:48455	-->	221.170.182.244:32425	UDP	ACTIVE	11/s	11/s
05/01 20:38:44	localhost:48455	-->	140.116.81.131:51118	UDP	ACTIVE	3.70KB/s	3.33KB/s
05/01 20:39:01	localhost:48455	-->	221.137.93.250:49053	UDP	ACTIVE	0/s	0/s
05/01 20:38:45	localhost:48455	-->	192.168.1.100:59066	UDP	ACTIVE	0/s	0/s
05/01 22:54:29	localhost:48455	-->	134.130.140.134:17730	UDP	ACTIVE	0/s	0/s
05/01 22:54:59	localhost:48455	-->	130.232.146.68:26238	UDP	ACTIVE	0/s	0/s
05/01 22:55:29	localhost:48455	-->	213.114.173.190:25378	UDP	ACTIVE	0/s	0/s
05/01 22:56:29	localhost:48455	-->	194.109.150.3:21773	UDP	ACTIVE	0/s	0/s
05/01 22:56:59	localhost:48455	-->	65.33.206.122:13412	UDP	ACTIVE	0/s	0/s

Figure 4: Skype Connection Table

All the calls that I saw used the iSAC codec and I was seeing rates between 80 and 85 kbps in Packetyzer. Knowing that the codec was supposed to be using something more like 32 kbps, this initially surprised me. But the numbers I had recorded were for both directions, so splitting it in half gives rates about 40 kbps. That speed seems very reasonable considering the 64 bit UDP frame size overhead. Those speeds were also confirmed with EtherPeek NX. Speeds between 27 and 30 kbps (converted from bytes per second) were identified with NetPeeker.

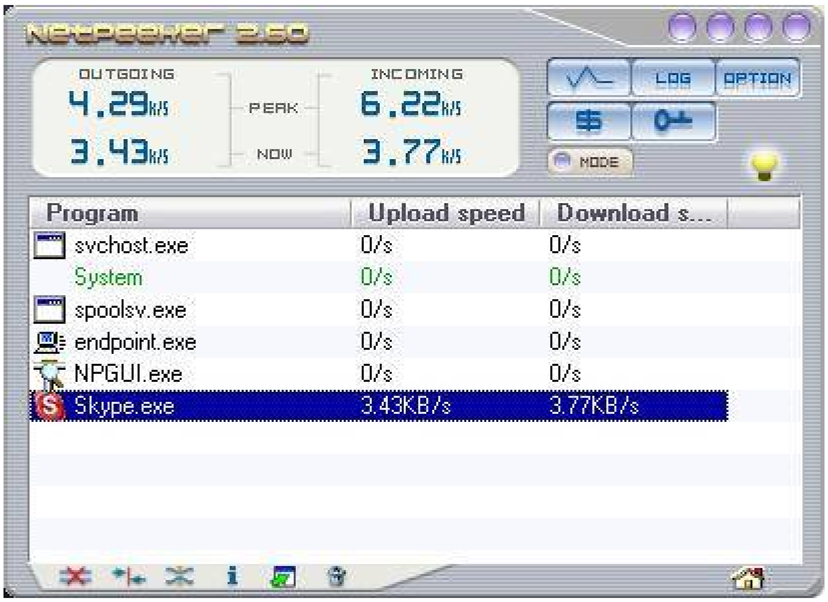

Figure 4: Skype Data Rates

This confirms that it is fairly reasonable for even a modem with dial-up speeds of 33.6 kbps up and 52 kbps down to use Skype with a more compressible codec such as iLBC. The consistency in speed also confirms that there is no sound suppression. Even if my wife turned off the microphone there was no noticeable decrease in throughput.

Section 5: Security

There are four standard elements to security: authentication, privacy, integrity, and availability. Skype addresses most of these issues.

Skype requires each user to *authenticate* into the Global Index server before they can communicate with other users. Because their credentials are secured using strong encryption, it is reasonable to assume that no logged in user can pretend to be another node.

All communications, both for signaling and voice traffic, are encrypted, which ensures *privacy*[39]. This means it is impossible to know what exactly is being transacted between two nodes, two supernodes, or between a node and supernode. According to their website, Skype uses AES (Advanced Encryption Standard) with 256 bits.[40] Because Skype doesn't use certificates, symmetric AES keys need to be shared before anything can be decrypted. To obtain the AES keys Skype uses another encryption mechanism: 1536 to 2048 bit RSA. During the RSA authentication sequence, Skype passes the public key which is then verified using the Skype server's private key.

One of the major reasons that calls and text messages are encrypted is to avoid leaking the conversation to users that are hosting supernodes. Because some calls are being relayed via supernodes, if it wasn't encrypted the user at the supernode could literally capture all the packets and play them back. In spite of the encryption, traffic flow can be analyzed. The government may be concerned that terrorists could use Skype to further their plans, but the IP communication flow is still known. If the call is literally P2P, as it was in the case for Call Number 2, a packet capture will show the 'from' and 'to' nodes. If the call is being proxied, as in the case for Call Number 1, the investigative agency would need to tap those connections and based on the content of the traffic (presuming that the information is not re-encrypted at each supernode, but rather, a new source and destination IP address is put onto the packet), and try to coordinate all the hops. In those situations the identification challenge is quite daunting.

The security is strong enough to protect the conversation, but the technology still makes it possible for the source or destination of the call to be identified. Privacy has never been a "design goal" of the system, according to Zennström in a CNET news.com interview.[41] Skype does not guarantee absolute anonymity.

---

[39] *P2P Telephony Explained*. 30 Apr. 2004 <http://www.skype.com/skype_p2pexplained.html>
[40] *Frequently Asked Questions*. 30 Apr. 2004 <http://www.skype.com/help_faq.html>
[41] McCullagh, Declan. "Skype's VoIP ambitions." *CNET news.com* 2 Dec. 2004. 24 Apr. 2004 <http://news.com.com/2008-7352-5112783.html>

Integrity of content is assured because any change or manipulation in packets will make it impossible for the receiving node to decrypt the now-damaged packet. Integrity of source or destination is not necessarily assured. The IP source and destination addresses could be manipulated, and because the Skype protocols are not open for review, it's possible that source and destination information is not also stored in the data field of the packets for comparison.

Skype deals very well with availability, at least on a large scale, because of its unique P2P communication model. If the performance of any one link would be reduced, the client will work with the supernodes to find another way around. Of course, if the link between the user's computer and their network connection comes under attack, there are no other available routes and the connection would be impaired.

Section 6: Compare and Contrast with FWD

Similarities

Both Free World Dialup (FWD) and Skype are the same in the sense that they are closed communities. Members of FWD generally can call only others using the FWD service, but there is some gateway support to other SIP services. In a similar fashion, Skype users can call only other Skype users. Both FWD and Skype negotiate the call in a comparable manner. They use a third-party server to coordinate the call setup between the two clients. With FWD it's the SIP proxy server, while Skype uses a supernode. Once the call has been set up they both use an UDP stream to send the voice packets.

Both businesses are also both for-profit and privately owned. Both launched a basic service with a free client, and both know that in order to succeed, they need to offer advanced services. That's why Skype is busy working on providing mobile handset, and creating voicemail and PSTN-gateway services. And that's why Jeff Pulver, one of the principals at FWD, is busy selling a wireless SIP phone and promoting his for-pay service, Librétel, that does offer the ability to call internationally. FWD also realizes that the benefit of consumer VoIP will increase exponentially has more people sign on, which is why they now have connections with Packet8 and are working on connecting with Vonage.[42]

Differences

Free World Dialup is based on SIP, an IETF standard that continues to receive attention and development. Skype's call setup and management protocols are closed, as well as encrypted, and the management team seems to have no intention to open the project up. This lack of openness will make it difficult, if not impossible, for other third-party clients using the Skype network to be made and be offered on the Linux and Mac platforms. Skype is also tied to certain codec support. FWD's system is codec independent (G.711, G.723, G.729, iLBC, GSM 6.10, etc), and two people can communicate as long as their client software supports the codec in use. Skype, on the other hand, seems to use just a few codecs, namely iLBC and iSAC, although these limitations do not seem to hurt voice quality. While I only tried X-Ten's default codec, a quick review of other FWD supported software such as SJ Lab's SJPhone and eyeP Media's eyeP Phone Lite show that there is no wide-band audo codec support. Skype's call quality is generally better than that provided by FWD's service. Skype's calls are always encrypted, while SIP generally does not use encryption. There are developing standards or drafts for SIP message body encryption using S/MIME,[43] and voice content via SRTP, in the context of MIKEY (Multimedia Internet KEYing.[44]

One of the most significant differences between the two products or companies is the fact that FWD requires registration at a centralized SIP registrar server, while Skype's presence and login capabilities are totally distributed. Skype voice calls are therefore not susceptible to server outages, though it presumes that Skype's certificate server is always up to provide for the RSA key exchange. This lack of centrality also lowers Skype's

---

[42] Pulver, Jeff "The Jeff Pulver Blog" 14 Feb. 2004. 1 May 2004 <http://192.246.69.231/jeff/personal/archives/000537.html>
[43] Chatterjee, Samir. "An Overview of SIP Security" Fall 2003. 1 May 2004 <http://www.internet2.edu/presentations/fall-03/20031014-SIP-Chatterjee.ppt>
[44] Arkko, J. et al "MIKEY: Multimedia Internet KEYing." Dec. 2003. 1 May 2004. <http://www.ietf.org/internet-drafts/draft-ietf-msec-mikey-08.txt>

operational costs for running their network. According to the Fortune Magazine article, it costs Vonage almost $400 to add a customer, but Skype only about a tenth of a cent.[45]

Skype also has an advantage because it can use supernodes to connect to previously unconnectable clients because of firewall or NAT issues. Much of the class struggled with getting the X-Ten client to work with their SOHO gateway. Although there are setup options that explicitly take care of NAT traversal or technical solutions such as STUN gateways, Skype avoids these problems altogether with an elegant solution: using a middleman to host and connect the call.

Section 7: Conclusion

Yes, Skype is for real. It offers excellent voice quality in an easy-to-use manner. The problems encountered by consumers and businesses alike with firewalls and/or NATs are largely avoided due to Skype's unique distributed and peer to peer infrastructure. Although Skype's product results from a closed software development and encrypted infrastructure, their execution until now has been superb. Their choice to make the client non-SIP based has received a lot of criticism, but in doing so they have overcome of the obstacles that SIP clients have faced. Standards-based technology is not always necessary with the new or disruptive technology if market traction can be gained, but once voice communication over the internet has significant market share standards will be necessary to create interoperability between vendors, which introduces competition and the resulting lower prices. If the company is able to expand services to include voicemail and provide gateway access to the PSTN and other SIP-based systems they will draw many more consumers, some who are willing to pay for premium services. Until that time, there is no way that this client can replace a regular phone.

The client lacks in two major aspects: there is no cross-platform support for the Mac and Linux operating systems. Second, contacts also aren't stored in the distributed so system, so each time the client is installed re-authorization is required from colleagues and friends. Allegedly both of these issues will be addressed in future releases.

Where Skype will really take off is if they are able to develop a product or appliance that can be housed in the corporate server room to authenticate corporate users to an existing identity directory, and integrate with business applications in such a way that its use is transparent, but

[45] Roth, Daniel. "Catch us if you can" *Fortune* 9 Feb. 2004. Business Source Elite. EBSCOhost. 19 Apr. 2004

that the financial benefits and increased productivity are compelling. The software needs to spread from the home as a fad to the board room.

Call Number 1: Cipactli Bulk

Your geographic location: City, State, Country	Syracuse, New York, USA
Partner's geographic location: City, State, Country	Syracuse, New York, USA
Type of connection, i.e. Dialup, DSL, Cable Modem	LAN
Date and Time of day of the call:	May 1, 2004, evening
Your IP address	Private NAT'ed IP address behind 128.230.97.222
Partner's IP address:	24.59.145.92
Your ISP:	Syracuse University
Partner's ISP:	TimeWarner/RoadRunner
Call quality rating and comments: (See below for explanation)	5 She could hear me very well, but her microphone and sound level is poor, so it sounded weak
Ping packets sent: (See Diagnostics for explanation)	13
Ping packets received:	13
Ping packets percent loss:	0%
Ping packets: Minimum/Maximum/Average Latency (round trip times)	21/45/26
Number of router hops between you and partner: (See tracert command below).	16

```
Tracing route to frnkblk.homeip.net [24.59.145.92]
over a maximum of 30 hops:

  1     3 ms     3 ms     3 ms  128.230.97.3
  2     4 ms     4 ms     4 ms  128.230.92.1
  3     4 ms     4 ms     4 ms  backbone.syr.edu [128.230.93.2]
  4     4 ms     4 ms     4 ms  syrgate3.syr.edu [128.230.85.2]
  5    10 ms     9 ms     9 ms  67.99.63.125
  6    13 ms    10 ms     9 ms  216.140.10.29
  7     9 ms     9 ms    10 ms  pop1-nye-P12-2.atdn.net [66.185.149.201]
  8    10 ms     9 ms     9 ms  bb2-nye-P0-0.atdn.net [66.185.151.50]
  9    13 ms    13 ms    13 ms  bb1-alb-P5-0.atdn.net [66.185.152.70]
 10    13 ms    13 ms    13 ms  pop1-alb-P0-1.atdn.net [66.185.148.101]
 11    14 ms    14 ms    14 ms  rr-albany.atdn.net [66.185.133.226]
 12    18 ms    18 ms    18 ms  pos15-0.syrcnyspp-rtr04.nyroc.rr.com [24.92.224.169]
 13    18 ms    18 ms    18 ms  srp5-0.syrcnyspp-rtr04.nyroc.rr.com [24.92.224.140]
 14    19 ms    19 ms    19 ms  rdc-24-92-228-36.nyroc.rr.com [24.92.228.36]
 15    19 ms    19 ms    24 ms  fas1-0.syrcnysyr-ubr03.nyroc.rr.com [24.92.228.57]
 16    26 ms    26 ms    23 ms  syr-24-59-145-92.twcny.rr.com [24.59.145.92]

Trace complete.
```

Call Number 2: Joshua Bulk

Your geographic location: City, State, Country	Syracuse, New York, USA
Partner's geographic location: City, State, Country	Sioux Center, Iowa, USA
Type of connection, i.e. Dialup, DSL, Cable Modem	LAN
Date and Time of day of the call:	May 1, 2004, midday
Your IP address	Private NAT'ed IP address behind 128.230.97.222
Partner's IP address:	24.59.145.92
Your ISP:	Syracuse University
Partner's ISP:	Dordt College, MTCNet, and then NetINS
Call quality rating and comments: (See below for explanation)	5 He was using a headset, while I was just using a cheap boom microphone with my laptop's speakers. The quality was really stunning. We never detected any voice quality drops of any kind
Ping packets sent: (See Diagnostics for explanation)	759
Ping packets received:	759
Ping packets percent loss:	0%
Ping packets: Minimum/Maximum/Average Latency (round trip times)	58/759/83
Number of router hops between you and partner: (See tracert command below).	19

```
Tracing route to PC10-105.dordt.edu [198.97.44.105]
over a maximum of 30 hops:

  1     3 ms     3 ms     3 ms  128.230.97.3
  2    14 ms     5 ms     3 ms  128.230.92.1
  3     3 ms     3 ms     3 ms  128.230.93.2
  4     4 ms     3 ms     4 ms  128.230.85.2
  5     9 ms     9 ms    10 ms  67.99.63.125
  6     9 ms     9 ms     9 ms  216.140.10.29
  7     9 ms    10 ms     9 ms  broadwing-gw.n54ny.ip.att.net [192.205.32.105]
  8    12 ms    12 ms    10 ms  tbr2-p012202.n54ny.ip.att.net [12.123.1.126]
  9    16 ms    17 ms    16 ms  tbr2-cl1.wswdc.ip.att.net [12.122.10.54]
 10    16 ms    16 ms    17 ms  tbr1-p013601.wswdc.ip.att.net [12.122.9.149]
 11    46 ms    46 ms    46 ms  tbr1-cl4.sl9mo.ip.att.net [12.122.10.30]
 12    44 ms    44 ms    44 ms  gbr6-p100.sl9mo.ip.att.net [12.122.11.110]
 13    45 ms    45 ms    45 ms  gar3-p370.sl9mo.ip.att.net [12.123.25.29]
 14    56 ms    56 ms    56 ms  12.124.129.114
 15    49 ms    49 ms    50 ms  ins-border10-at-0-3-0-un0.desm.netins.net [167.142.57.215]
 16    75 ms    63 ms    61 ms  sioux-center7.desm.netins.net [167.142.62.14]
 17    63 ms    62 ms    65 ms  router-outside.dordt.edu [216.248.99.206]
 18   103 ms    74 ms    68 ms  PC10-105.dordt.edu [198.97.44.105]
 19   128 ms   138 ms    60 ms  PC10-105.dordt.edu [198.97.44.105]

Trace complete.
```

Call Number 3: Frank Robinson (with Joshua Bulk at times)

Your geographic location: City, State, Country	Syracuse, New York, USA
Partner's geographic location: City, State, Country	Syracuse, New York, USA
Type of connection, i.e. Dialup, DSL, Cable Modem	LAN
Date and Time of day of the call:	May 1, 2004, midday
Your IP address	Private NAT'ed IP address behind 128.230.97.222
Partner's IP address:	24.59.146.221
Your ISP:	Syracuse University
Partner's ISP:	TimeWarner/RoadRunner
Call quality rating and comments: (See below for explanation)	5 The quality was excellent, with no blips or hiccups.
Ping packets sent: (See Diagnostics for explanation)	56
Ping packets received:	56
Ping packets percent loss:	0%
Ping packets: Minimum/Maximum/Average Latency (round trip times)	23/77/34
Number of router hops between you and partner: (See tracert command below).	16

```
Tracing route to syr-24-59-146-221.twcny.rr.com [24.59.146.221]
over a maximum of 30 hops:

  1     3 ms     3 ms     3 ms  128.230.97.3
  2     4 ms     4 ms     4 ms  128.230.92.1
  3     3 ms     3 ms     3 ms  backbone.syr.edu [128.230.93.2]
  4     4 ms     4 ms     4 ms  syrgate3.syr.edu [128.230.85.2]
  5    10 ms     9 ms     9 ms  67.99.63.125
  6     9 ms     9 ms    10 ms  216.140.10.29
  7    10 ms     9 ms     9 ms  pop1-nye-P12-2.atdn.net [66.185.149.201]
  8     9 ms    10 ms     9 ms  bb2-nye-P0-0.atdn.net [66.185.151.50]
  9    13 ms    13 ms    13 ms  bb1-alb-P5-0.atdn.net [66.185.152.70]
 10    13 ms    13 ms    13 ms  pop1-alb-P0-0.atdn.net [66.185.148.97]
 11    14 ms    14 ms    14 ms  rr-albany.atdn.net [66.185.133.230]
 12    19 ms    18 ms    18 ms  pos15-0.syrcnyspp-rtr04.nyroc.rr.com [24.92.224.169]
 13    18 ms    18 ms    18 ms  srp5-0.syrcnyspp-rtr04.nyroc.rr.com [24.92.224.140]
 14    18 ms    18 ms    18 ms  rdc-24-92-228-36.nyroc.rr.com [24.92.228.36]
 15    19 ms    19 ms    19 ms  fas2-0.syrcnysyr-ubr03.nyroc.rr.com [24.92.228.61]
 16    26 ms    53 ms    26 ms  syr-24-59-146-221.twcny.rr.com [24.59.146.221]

Trace complete.
```