

Chapter 9: Worrying about Skype

Introduction

Skype is a well-known free voice over IP service which has picked up millions of users within the last few years. Unlike many of its competitors, it has achieved almost household name status. Its iconic rank was made complete through its acquisition by eBay for \$2.6 billion in September 2005.

In its original, and still dominant mode, the Skype client runs on a personal computer. The Skype executable can be downloaded from www.skype.com and installed on a Windows or Linux machine. The client (figure 1) is protected by user-id and password (which can be stored by the client to avoid typing it on every use) and people you contact with Skype can be added to your ‘buddy list’ (‘buddy list’ is an AOL trademark) where you can double-click their entries to establish a call.

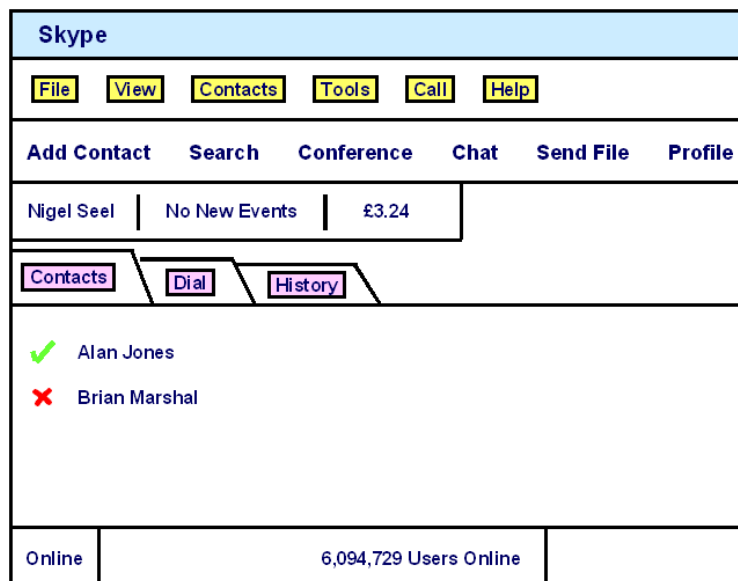


Figure 1. A schematic representation of the Skype client

To make or receive a Skype call, you need to purchase a headset with earphones and microphone. A PC’s built-in speakers and microphone would also work at some level of sound quality at the expense of privacy and/or annoyance to others. Skype can also do video-telephony (you need to purchase a web-cam), and provides facilities for instant messaging and file transfer. All communications are end-to-end encrypted. Up to five people, at time of writing, can collaborate in a voice conference call.

Calls between Skype clients anywhere in the world are free. In theory, you could set-up a videoconference session from London with a friend in Los Angeles and leave it open continuously - a virtual window. Since Skype provides neither the machine cycles to run the two clients, nor the bandwidth between the two machines, it costs Skype nothing.

How Skype works

Everyone will tell you that Skype’s essential differentiator from other VoIP providers is that it is peer-to-peer (P2P). We shall have a lot more to say about P2P in chapter 11, but note for now that unlike most VoIP service providers, Skype does not use its own servers to provide the bulk of its services - users host the Skype service on their own machines via the Skype client. To be quite precise, the key difference between Skype and more conventional VoIP providers lies in the signaling plane. Most everyone else uses protocols like SIP or H.323 which require network servers to maintain directories, locate users and forward the signalling messages which set-up and tear-down calls. Skype, however, relies on users’ PCs and Skype handsets to do these functions.

However, once a call session has been established, all VoIP providers, in the normal course of events, allow the two user terminals to communicate directly across an IP network, without the intrusion of intermediate servers. There are, of course, always exceptions: bridging for conference calls, dealing with NAT and firewall devices, the needs of lawful interception. But in the main, for *all* VoIP systems, media transport (i.e. the call itself) is peer-to-peer.

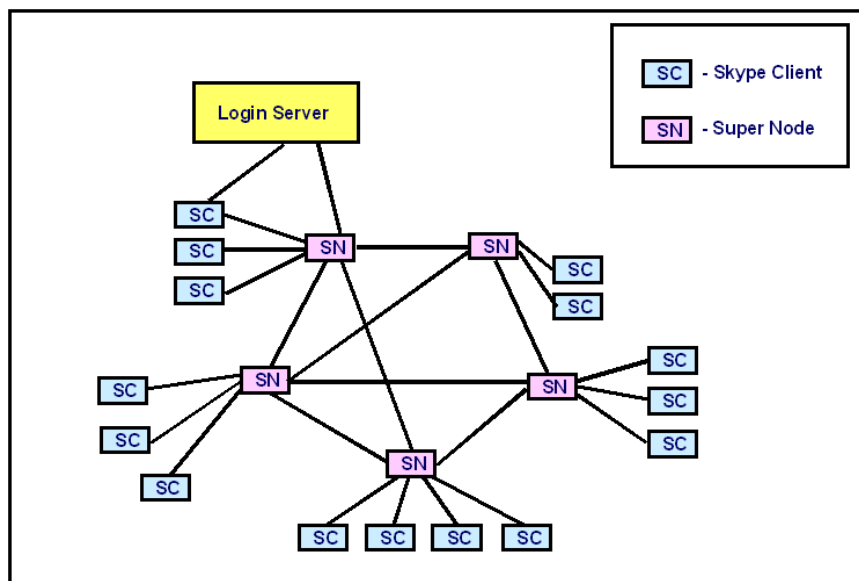


Figure 2. The Skype P2P architecture.

As shown in figure 2, there are two kinds of Skype nodes, an ordinary Skype client and a super node. Skype clients are the programs used by people who want to make and receive Skype calls. Super nodes are also skype clients, but running on these machines the client performs extra functions to help make the Skype service work as a whole. These extra super node functions can include:

- Acting as a proxy between a Skype client and the Skype login server.
- Helping a Skype client determine if it is behind a NAT/Firewall device.
- Helping Skype clients to find other Skype users.
- Proxying calls between Skype clients behind NAT devices and/or firewalls.
- Providing conferencing services (3+ way calling).

There is no direct way for a user to control whether their Skype client program becomes a super node or not, but promotion to super node does require that the machine running the client should have a public IP address and should not be behind a NAT device. Machines with more powerful CPUs, more memory and more network bandwidth are also preferred. Many home machines will, therefore, never become super nodes.

In experiments described in [1] a survey was made of super nodes on the global Internet. The US had 84%, Asia 9% and Europe 7% of identified super nodes. In more than 8,000 logins, 35% had a North American university suffix (.edu) and these comprised 102 universities. The top five universities were, in order, Harvard, Columbia, the University of Nebraska-Lincoln, the University of Pennsylvania and Boston University.

By allowing Skype clients to undertake super node functions if the machines they are running on have the right properties and capabilities, Skype has arranged that the clients are performing functions which in more conventional architectures are carried out by network servers. It seems clear that many machines have advertised their status as super nodes quite widely, and that the thousands of super nodes constitute a kind of global decentralised directory, describing both each other's existences, and that of connected Skype users.

When you log onto the Skype network, your Windows Skype client will look in a local cache on the 'C' drive to find a super node. It appears to make a random choice, as the local cache lists hundreds of super nodes and their IP addresses. This super node link is used to routes messages to the log-in server (which

is a dedicated Skype server, not a Skype client or super node) which authenticates you. Once logged in, your Skype client asks its connected super node for information about the IP addresses and connectivity status of people on your buddy list. If the super node has this information, it returns it, otherwise it sends back the IP addresses of other super nodes to which the queries should be sent. This process can be executed multiple times. If all else fails, the Skype client can access the login server (this is a last option to avoid this server being overwhelmed by traffic). The search process is remarkable fast, with the buddy list often being populated within 3-4 seconds.

To call someone on the buddy list, the entry is double-clicked. If both caller and callee have public IP addresses, call set-up occurs over a TCP connection set up directly between the caller and callee. If either or both Skype clients are behind NATs and/or firewalls, then TCP signalling is relayed through a super node. Once the call is established, the media transport (the voice call itself) is sent directly between caller and callee using UDP except in the case where both parties are behind firewalls which block UDP traffic. In this case the caller and callee communicate via a relaying super node using TCP with very small packet sizes. Call tear-down is similar to call set-up.

The extra functionality that Skype needs to employ to support calls to and from the public telephone network, the PSTN, cannot use the basic P2P architecture. Interfacing to the PSTN requires special media and signalling gateways, and these are additional pieces of equipment which Skype supports as part of its own infrastructure.

The establishment view of Skype

The setting is a conference in Barcelona on future networking. The speaker, Strato, is from ETSI, the European Telecommunications Standards Institute. He is talking about IMS standards, recalling that IMS is the IP Multimedia Subsystem, which handles call set-up and tear-down in the official ITU-T/ETSI version of the Next-Generation Network (we discussed this at length in chapter 2). But first Strato needs to talk about Skype.

Skype has been omnipresent at recent conferences. Many speakers admit to using it, and are impressed by its robustness, its voice quality, and of course by the fact that it is free to use. For others, it is a major danger to the industry - ('how do you compete with free?' asks Strato, rhetorically) - and a veritable road to hell.

Strato's case against Skype rests on five major assumptions.

- Skype cannot guarantee quality.
- Skype is a security risk.
- Skype is not using official standards.
- Skype undermines the carriers' business model.
- Skype is not as good as people think.

Strato speculates that IMS and other parts of the Next-Generation Network infrastructure will enable *respectable* operators to suppress such rogue traffic, and wonders aloud whether operators would be right to do so.

How worried should we be, I wonder. How *could* operators close Skype down, and would it be in their and/or the public interest for them to do so? After all, hundreds of millions of people have downloaded the Skype client, and millions are using Skype online at any one time. And there are many, many, other VoIP operators offering free on-net calls. Skype is unique only in using peer-to-peer (P2P) technology (invisible to the user) and being really easy to use.

Skype cannot guarantee quality

Skype uses the public Internet as its transport medium. Unlike traditional carriers, it sees the whole world as its marketplace, and this has been reflected in its rapid growth. Skype transmits voice frequencies between 50 Hz and 8 kHz, twice the bandwidth of ordinary PSTN calls, and the improved quality is quite noticeable over even a low-rate broadband connection.

Skype shows that today's Internet provides sufficient quality of service for voice calls. This comes as a perennial surprise to the advocates of Internet QoS add-ons such as bandwidth managers, QoS marking, Diffserv and RSVP, and all the other attempts to introduce multiple service classes to the Internet. However, on relatively uncongested links, Skype voice works just fine without all that stuff.

In fact as far as the Internet core is concerned, there are good reasons to believe that for a network with a rising traffic load, links will normally be uncongested. When a customer signs up with an ISP, they wish to connect to the whole Internet, not just to other users on that ISP's network. In a competitive market, it pays the ISP both to keep its own network uncongested, and also to ensure that peering links with other ISPs are also uncongested. By adopting this policy it cannot guarantee end-to-end performance (some remote ISP could still introduce congestion, although its immediate neighbours might want to 'have a

word’). However, by failing to adopt this policy it will absolutely *guarantee* congestion and thereby make itself uncompetitive. The market equilibrium is that a single class of service (‘best-effort’) Internet as a whole will be uncongested.

With an exponential model of traffic increase, most of the additional traffic comes at the end of any time interval. Assuming the network is never allowed to run in a chronically congested state, this implies that most of the time the network will be fairly empty as shown in figure 3.

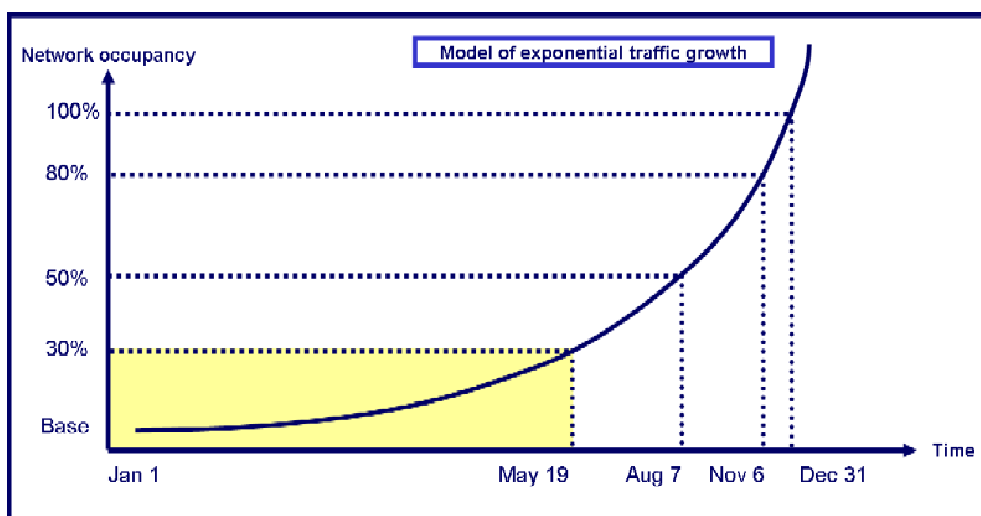


Figure 3. Network occupancy without congestion

There is a caveat of course, in the context of the net neutrality debate. If carriers manage to create a tiered Internet, with a non-congested QoS tier for which Service Providers such as Google and Yahoo! would pay premium prices, then the best-effort service might be allowed to periodically dip into congestion. As VoIP media transport, being peer-to-peer, does not rely upon a Service Provider being ‘in-the-loop’ of the connection for the call itself, then Skype users (and all other VoIP-using customers) might have to pay extra to have their calls marked-up and carried as high-priority packets. To be precise, end-users might receive a tiered bill, where they pay extra for the bulk amount of traffic carried at higher QoS classes, this could add an element of usage-based charging, depending on how the tariffing works. It is unlikely that carriers will enforce per-flow billing (too many flows), or be able to specifically identify Skype traffic per se (due to its encryption). Anyway, we are not in that future at the moment, but it is a strong possibility discussed further below. Chapter 13 also discusses these issues in more detail.

The access network is a different issue again. In a home environment, candidates for congestion might include any WiFi link between the Skype client and the home gateway, and then the broadband link itself. Skype voice calls consume from around 24 - 128 kbps of bandwidth in each direction. But this is well within the upload budget of today's broadband links, and is practically invisible in most users' downlink budgets. Even without class-of-service prioritization, few people today report service difficulties. The next generation of home gateways will support class of service traffic management as VoIP becomes more prevalent, and unless administrative actions are taken to downgrade Skype packets, this should improve the service further in the access part of the network.

So overall, the Internet, left to its natural devices, will not ruin Skype calls. But perhaps the carriers can artificially do it using their new IMS platforms or other packet-inspecting devices? There are several possible approaches.

Introduction of artificial impairments

Carriers *could* discriminate in favour of their own VoIP service and against non-approved VoIP providers as follows. They will arrange for their own clients, running on the PC or authorised terminal, to mark their approved VoIP packets with a real-time service class. Then, at the carrier edge device, they will add impairments such as jitter, delay and packet loss to any VoIP traffic not to/from the operator's own client, sufficient to disrupt voice quality.

But while this is technically possible, it seems too blunt a technique to work. The carrier will have as customers many other operators and ecommerce sites in the future which will want to support real-time traffic, and not all of them will use the operator's own VoIP system. And the regulator would have a field day in most advanced countries.

Lock-out Skype infrastructure or packets

Carriers could close access to the Skype web site or login server via a firewall, or drop Skype packets. There are reports of a few ISPs already having tried this.

In fact the Skype programming team is engaged in an arms war with ISPs trying to suppress Skype traffic. It used to be possible to stop Skype by blocking the IP address of the default login server, according to Baset and Schulzrinne [1]. However, they note that recent versions of Skype get around this by routing login requests via intermediate super nodes. In addition, it appears that the Skype executable has hard-coded a number of bootstrap super node IP addresses, which are not made public. Baset and Schulzrinne

were able to identify four standardized bytes in the Skype initial login message (0x16030100) which could be dropped by a firewall, but they point out that these values are hardly unusual in ordinary data traffic. Most Skype traffic, both control messages and user content, is encrypted, which ensures that there are no regularities which can be filtered.

Could a carrier get away with it?

We already see mobile network operators, as they introduce faster HSDPA data services, banning the use of VoIP. Their reasoning is understandable, but the position seems completely untenable in the longer term. The 3G mobile operators are in a complete bind: to exploit their new networks they have to provide generic high-speed (multimedia) connectivity. While consumers may be pushed into a walled-garden for a while, enterprise staff need access to their corporate networks and to the public Internet. And with non-usage-based tariffing (preferred by users) the marginal cost of third-party VoIP client usage becomes zero. At this point their existing and expensive circuit-voice revenues crumble.

And there seems no way out. For every 3G operator who attempts to use administrative means to stop VoIP, there will be a competitor who sees an advantage in permitting it. The average sales person on the road will run Skype (or similar) on their PC regardless of their company's contract with the MNO. Assuming the operator, for technical reasons, can't stop them, will this lead to the threatened contract cancellation? I don't think so.

Skype is a security risk?

The alleged security concern with Skype was one of the more important of Strato's specific issues and a favourite topic of those who wish to attack it. Garfinkel, [2] reviewed the situation with Skype security as follows. Skype claims that conversations, instant messaging and file transfer are all AES-encoded (AES = Advanced Encryption Standard) with 256 bit encryption and that the required symmetric keys are exchanged securely by public key cryptography using a 1024 bit RSA algorithm. Packet analysis confirms that Skype traffic is not readily decipherable. Given the proprietary nature of Skype protocols and algorithms, independent confirmation is not possible, but Garfinkel notes that the PSTN and most competing VoIP systems are not encrypted at all.

As mentioned above, some Skype calls, particularly multi-participant bridged calls, transit super node computers. It is believed that hardly any of these super nodes are owned by Skype - they are mostly machines with public IP addresses owned by ordinary Skype users, frequently university machines as already noted. As such, there may be a security risk for Skype users whose machines end up as super

nodes with transit traffic being carried on their machines. There is also the issue of their machine resource being used for supporting other people's calls. Super nodes also seem to be involved in the distributed directory search to set up calls, and when searching for other users logged onto the Skype 'cloud'.

Skype makes no attempt to hide the identity of connected users. It appears that a Skype search can find any recently logged-on user. This may raise privacy concerns. In particular, organisations may not wish the names and registration details of their employees to be visible to any Skype user in the world. However, it is a Skype licence condition of use that Skype is not used for commercial purposes.

Garfinkel's conclusion is that Skype's security status is pretty good for a consumer product, but that there are a number of potential security vulnerabilities and issues, largely due to Skype's unwillingness to be open about how it addresses them. However, it is certainly more secure than most competing VoIP services.

It should be noted that Skype is neither more nor less open about its product than other software companies. Microsoft does not make its source code public, nor does it generally publish its algorithms. There are few independent guarantees that Microsoft security protocols necessarily do exactly what they claim. Most Microsoft users have little idea of what the large number of active Operating System processes on their machines are doing. Users just trust the brand. Now that Skype has been acquired by eBay, one assumes that similar brand guarantees might apply.

It should be relatively easy for a sufficiently concerned organisation to test Skype's claims. Its client is under 10 Megabytes and could be disassembled. Known plaintext can be sent via Skype's instant messaging subsystem and the cipher text analysed to assess the power of the encryption. Ditto with voice using a tone generator.

Given the opportunities for competitors to Skype, it is not wholly surprising that Skype see an advantage in retaining commercial secrets. In fact there is ample room for competitor P2P products which addresses many of these security issues, and which focus more on commercial markets.

Here's the bottom line. Millions of people are actively using Skype at any particular moment. There are a huge number of vested interests who wish to discredit Skype on security grounds. Yet not a single case of a Skype security violation has ever been published. And then there is a final twist. Skype was developed by an entrepreneur and a group of Estonian programmers. It was quite likely that for a while, Skype

messaging and calling were genuinely difficult to impossible for the world's intelligence agencies to crack. I doubt that anyone would say the same today, following its acquisition.

Skype is not using official standards

This was another of Strato's major points, and he felt on stronger moral ground here. Standards are, after all, ethically positive. They prevent lock-in to particular manufacturers, enlarge the market and by increasing competition, lower prices.

So are users locked into Skype? Well, Skype certainly encourages its users to evangelise their friends, even when this turns a SkypeOut call (= revenues to Skype) into a free on-net call. These guys are plainly serious about network effects. But it's easy to remove the Skype client and use one of the many other free VoIP soft clients out there. Switching costs and lock-in are minimal.

Would standards enlarge the market in this case? The standards which matter here are for interconnect, not the interior signalling standard. This is not about SIP. At some stage it will be necessary for a Skype user to be able to talk to a Vonage user, for example. However, SkypeOut proves that interconnect is possible, so it is down to Skype to make it happen. Standards and interconnect are not the same thing at all. And lower prices? Well, as Strato himself said, 'how do you compete with free?'

It is understandably mortifying for people who have given years of their lives to developing scalable, robust and sophisticated protocols to find that the most successful product in the world has used something quite different, and proprietary and secret to boot. But unfortunately, that happens, and it's important not to obsess about it.

It's the reasons for the standards, not the standards themselves, which are essential. In the case of Skype, the issue is not so much the use of a novel P2P protocol as the fact that it's kept secret so that its properties cannot be verified. That in the end may restrict its applicability in higher-value applications.

Skype undermines the carriers' business model

Skype's cost base is a login server cluster, perhaps a few PC bootstrap super nodes, and the salaries and expenses of its few executives and programmers. The operational infrastructure for the millions of Skype users consists of their already existing computers and broadband Internet connections. Skype's fixed costs are low, and variable costs essentially zero. This is how Skype was able to scale to hundreds of millions

of downloads and millions of concurrent users so rapidly. Skype makes its money from its value-added services: SkypeOut, SkypeIn and Voicemail.

SkypeOut allows the user to call PSTN numbers from the Skype client on the PC. The call is carried from the Skype client across the Internet for the long-haul part of the call, and then breaks out locally to the PSTN for the final part. Skype has to pay the local telephone operator to carry that final leg, and that charge is pushed back to the Skype user. To use the service, you have to set up a pre-pay account, and Skype retains a proportion of this revenue stream. Per-minute charges are very low.

SkypeIn allows the user to purchase PSTN numbers in a variety of global locations for around \$30 per year (up to ten numbers can be bought). Incoming calls to these numbers are then transferred to the Internet and conveyed to the user's Skype client. This service provides a second revenue stream to Skype.

Skype's third chargeable service, at time of writing, is voicemail, although this is currently bundled at no extra cost with SkypeIn.

From an economics point of view, Skype efficiently utilises an available resource (user computing power and broadband connectivity) and harnesses it to satisfy a demand at very low cost. This is wholly to be applauded, and is what is meant to happen in competitive markets.

Skype's lack of substantive owned infrastructure means that almost all non-trivial functions must be carried out in the client. If the client is doing infrastructure work for other people (e.g. acting as a directory, transit node or conference bridge) this may be perceived by the user as effort expended without recompense on behalf of freeloading others. This sets a limit as to how much generic functionality the client can do in pure peer-to-peer model.

The pure P2P model therefore works less well where information needs to be managed in a way which is decoupled from any particular user's machine. An example is buddy lists, which cache contact information for family, friends and contacts. Originally, this information was stored on the local machine but this was irritating if you were running Skype on multiple devices - you had to re-enter the details on each new machine. Now buddy lists are stored on a central Skype server, which adds to Skype's costs.

Another problem for the P2P model is where substantial application functionality is needed. In more conventional architectures, this includes servers for announcements and value-added services. If Skype

wishes to introduce these services, it will probably have to put in place special platforms to host them. There are limits to the number of spare university machines!

It is extremely likely that large-scale carrier services based on IMS, and light-weight multimedia-over-IP services based on P2P will both co-exist. There is ample mileage in both approaches. Skype has proved that P2P works and can give excellent service quality, and there are reasons to expect the P2P ecosystem to diversify, perhaps with Skype-like variants for business, provably secure applications, contact centres, ecommerce and so on. A contrary argument, however, argues for the negative impact of network effects in the absence of interconnect. Why download and install someone else's 'secure-pseudo-Skype' when everyone you know is using Skype? Skype's obsession with expanding its user base implies that it could give master classes on the positive network externalities linked with first-mover advantage.

The apparent drawback with a pure P2P architecture is that fixed system functions (directories, conference bridges) are wholly dependent on enough users, and the right kind of users, being logged onto the network. Once this distributed infrastructure is in place, however, incremental end-user facilities are brought by the users themselves. As Chairman Mao once observed, each mouth comes with a pair of hands. A possible hybrid solution for an enterprise or operator is to run a small server farm - a collection of 'super nodes of last resort' - which can also support specialist services. This can be kept up all the time, and ensures that even the second user logging in will get service. Skype may do this itself, but to-date it isn't telling.

Strato, from his standards pedestal, could only see how unfair it was that those cowboy entrepreneurs, refugees really from KazaA, had sidelined his years of hard work on H.323, SIP and now IMS. But we *have* been here many times before. Sometimes the opportunities which have to be grasped are not those which we predicted with such pedestrian foresight all those years ago. Time to move on.

Is Skype really that easy to use?

Strato might have been on surer ground if he had emphasised usability. I have Skype running on my computers, as do some of my colleagues. The usage patterns are quite diverse. A colleague whose family is in France, but who works in the UK during the week, uses Skype to keep in touch. These are scheduled, laptop-to-PC calls. It seems to work well. Some of my colleagues call me using Skype but it often doesn't work. If I am away from my desk, I usually don't notice a call to the laptop. There are two cases:

- The laptop internal speakers are online, or an external speaker system is plugged in. In this case, I *may* hear the ringing depending on the volume settings, but answering the call is a problem as the headphone/speaker will not be plugged in. Disentangling wires, looking for the right line-out socket and struggling to put the headset on before the caller rings off is not a pretty sight.
- The headphone/speaker is already plugged in - this is rare. However, if I am away from the desk, I will not hear ringing tone as the headphone volume is too quiet.

I am not the only one with this kind of problem. Calling my colleagues using Skype I encounter similar issues. Is it surprising then that we fall back on the plain old telephony system, with its loud ringing and easy-to-use handsets? Domestic fixed telephony call-costs are low, so unless one is very price-sensitive, there is little incentive to incur the additional hassles of using Skype. And Skype does not substitute for mobile phones except for the smallest niches (where Skype has been installed as a client on a mobile network).

However, things will certainly improve. Even a cursory network search typing the words 'skype handset' into a search engine will turn up an increasing number of USB/WiFi handsets. A WiFi Skype handset can run a full client and does not need a PC or laptop. So it seems that if people want to emulate their current fixed phone service with Skype, they will be able to do so. Of course, in marketing, we know very well that just because an existing service can be emulated with a new technology, this doesn't necessarily mean that everyone will defect just like that. Those enticing free Skype-to-Skype calls are mixed with chargeable Skypeout and SkypeIn calls and the handsets themselves are far from cheap. Although voice quality today is good, who knows whether it will stay good - that is outside Skype's control.

Incidentally, not all of Skype's uses are what they seem. The Skype icon features on the system tray, bottom right of a normal Windows XP screen. When the Skype client is in contact with the Skype network, it shows a white tick-mark on a green background. If the client can't find the Skype network, it shows a cross on a grey background. Since Skype is highly adept at making contact, it serves as an excellent indicator of network connectivity. I find myself checking that icon quite a bit!

Conclusions

In economics terms, the Skype service demonstrates an efficient utilisation of the relevant factors of production. If the carriers don't like it, they are simply exhibiting the distortions inherent in their current pricing and business models, which positively invite an arbitrage response. In the short-term, of course, it is easier to use administrative action to block innovation than to adapt, but the imperatives of economics

cannot be suppressed for ever. In the end, in a competitive market, services which reside on users' already-bought end-systems and exploit vanilla connectivity in the network will be priced pretty much at the marginal cost of said connectivity. In this case, close to zero.

The idea that carriers will thereby go bankrupt, or have their business seriously disrupted, is ludicrous. Their market structure is oligopolistic, they retain control of network assets - where barriers to entry are severe, and all the major fixed and mobile players have significant market power. The only issue is how they will rebalance their tariffs nearer to their real cost structures, and the extent to which they intend to rely upon extracting rents for vanilla services such as network access vs. their abilities to develop new revenue-generating services where they have a competitive advantage (integrated network-hosted services, to be specific). The global herd of fixed and mobile oligopolies will have to lumber to a new business model without too much overt collusion, but lumber they will, and they will get there in the end.

In conclusion, at the moment, Skype still looks niche, but service innovation in the P2P space can be fast, and if they can put together something new with real mass appeal, then they could seriously frighten the IMS-bound carriers once again. But don't forget, the carriers control the infrastructure.

References

[1]. Baset, S. A. and Schulzrinne. H. G., An Analysis of the Skype Peer-to-Peer Internet Telephony Protocol, Technical Paper, Columbia University, 2006.

<http://www.eecs.harvard.edu/~mema/courses/cs264/papers/skype-infocom2006.pdf>

[2]. Garfinkel, S. L., VoIP and Skype Security, Technical Paper, March 2005.

www.tacticaltech.org/files/Skype_Security.pdf