Voice telephony, the next peer-to-peer application?

Georg Schwarz

Table of Contents

A technical invention gives birth to a new industry

When in 1876 Alexander Graham Bell was granted a patent for the electrical telephone he was not the first one, as we know today, to have come up with such a device. Beyond doubt however, it was the historical achievement of that enterprising American to clearly recognize the vast economic potential of that invention, which soon formed the basis of a fast-growing, multi-billion dollar industry.

Even though today's phones only quite remotely resemble the first models from Bell's days, their basic technological principles of operation have remained largely unchanged to date, and so has the core business model of telephone companies. Today just like a hundred years ago telephone customers for a monthly fee get connected to the operator's local switch and thus to the telephone network. They are assigned a telephone number through which they can be reached by other subscribers. When a call is placed it is the operator's task to establish the callee's line corresponding to the dialed number, and to set up a connection between the two parties. In the early days, this was done over dedicated pairs of wires which were interconnected over the operator's switchboards. Today, classical voice telephony typically uses TDM (time division multiplexing) technology in the carriers' backbones. In any case dedicated resources (wire pairs or fixed time slots) are temporarily assigned exclusively to a connection for the duration of the call. Not surprising, pricing models have come in use which charge callers by the length of a call, with rates going up by distance to account for the longer "copper lines" being occupied. The deployment of modern electronic switching technologies and large bandwidth fiber networks as well as, probably most importantly, the fierce competition in liberalized telecommunications markets has helped to bring down the price per minute of a long distance or international phone call often to only fractions of what they used to cost in former decades. Nevertheless, the principle of charging by the minute remains unchallenged, and due to the increased volume of calls still accounts for a considerable part of the telecom industry's revenues.

The telco business model

The basis of the telco business model is the exclusive access to and control over the customer a telephone company obtains when it connects a subscriber to its network. All incoming and outgoing connections to and from that customer will necessarily be established through the local switch and can thus be recorded and potentially charged for by the telco. When a subscriber dials a number it is the telephone company's switching equipment that performs the task of setting up the connection to the called destination. Reachability of subscribers on other operators' networks, for example in other countries, is ensured by telephone companies forging interconnection

agreements between each other and taking care of handing over calls. Here as well the charging model among operators follows the lines of "paying by the minute". For the end customer, the subscriber, the local telco remains the only contractual partner and service provider; it is the telco's role to offer voice services, set the price tag and send the bill.

From the customer's perspective this rigid model has been softened in part by liberalization of long distance markets. Subscribers in many countries for their telephone calls can select other providers than their local telco and make calls at these companies' tariffs. This new freedom of choice however is not brought about by a change in technology but by incumbents being forced through legislation to hand over calls to competitors at conditions and prices defined by regulatory authorities.

Particularly in Germany liberalization has led to an unparalleled decline of the price per minute of a long distance call. In spite of the increased overall call volume these changes have resulted in a considerable reduction of revenues from the voice minute business for ex-monopolist Deutsche Telekom. Markets in many other countries have seen or will see a similar, albeit often less dramatical development. The basic telco business model however remains untouched. The new entrants, enabled by regulation, simply take away some of the market share from the incumbents, but they operate at the same technology and the same business principles. Most importantly, it is the operator of the "last mile" that continues to ensure reachability of the subscriber under his or her telephone number and that determines, within the framework set by regulation, which services customers can use.

That "ownership of the customer" is a direct result of the technical voice telephony network architecture; and it is that fact which enables operators to control or at least participate in revenues from non-voice transport related "value added" services such as directory assistance or other "premium services". The operator of the last mile draws its dominating position within the telephony value chain from its technical monopoly on the local call setup and termination service. This holds true for all telephone networks, analog (POTS) or digital (ISDN), fixed or mobile (GSM, etc.).

The success of the Internet

With the advent of the Internet in the 1980s and, driven by the success of the World Wide Web, increasingly in the 1990s for the first time in history a world-wide communications infrastructure emerged whose philosophy and whose operational and technological concepts do not follow the approach of a "carrier-centric network". Traditional communications networks such as the world-wide telephone network (including mobile networks) or the telex network (or to a certain extend even radio and television broadcast networks) have been set up by one or several operators with the aim of connecting subscribers and offering them certain communications services. Consequently, these networks are designed along the lines of the carrier-centric network architecture model. They are typically architectured to offer a specific service only, e.g. voice telephony. End user devices (e.g. telephones) have a network interface and protocol with rather limited capabilities. They must be capable of selecting the service offered by the network (i.e. dial a destination number or pick up an incoming call) as well as some means of using that service. The service creation is always done by the provider, i.e. the network.

In contrast, the Internet is based on a completely different approach. Conceived as a robust means of interconnecting computer systems over large distances and enabling the remote use of these computing resources, it followed the approach of the "simple" network. With TCP/IP, the network protocol the Internet is built on, the "intelligence", e.g. error correction, the resolution of network addresses, all "higher level" services, in short almost the entire service creation occurs on the hosts connected to the network. One basic design goal of TCP/IP and a fundamental assumption of the Internet is that every host connected to the net can communicate with any other host. Likewise, there is no intrinsic limitation or preference from the network architecture which services a host connected to the Internet can use, or who will provide these services. Once a host has Internet connectivity, it is (ignoring for the moment artificial, deliberate restrictions through firewalls etc.) up to the operator of the host, not of the network infrastructure, which other hosts (servers) to connect to for whatever services and in turn which services to offer to whatever other hosts on the net. In this respect, the Internet's design as a simple, universal network is fundamentally opposed to the "operator-centric", special-purpose network approach, where service creation occurs within the network and is provided and controlled by the network operator. It is undoubtedly this property, which enables new Internet services to be deployed by practically anyone without the network having to be upgraded first, that has greatly contributed to the tremendous success the Net has seen over the past decades.

From a traditional telecommunications network provider's point of view, the Internet is a commercial accident, if not a nightmare, since it does away with the operator's monopoly or even preferred position on service creation. Successful online businesses such as Ebay, for example, need not share their revenues with local ISPs or telcos as it would have been the case in a carrier-centric network architecture (for example Deutsche Telekom's BTX service from the 1980s and early 1990s or similar "video text" systems). With its roots as a research network there has never been such a thing as a "business case for the Internet"; otherwise the Net would most likely not look the way it does, if it existed at all.

The impact of IP on voice telephony

The success of the Internet in the 1990s led to the establishment on IP as the universal, ubiquitous communications protocol. With the availability of suitable devices, e.g. workstations or personal computers with audio support, and sufficient bandwidth between Internet sites such as universities, people pursued the idea of using these communication lines for voice (and later video) conversations. Soon a number of applications and protocols were developed, most of them more or less incompatible with each other. When the popularity of the Internet rose as more and more people got "online" and the typical dialup bandwidth increased due to better and cheaper modem technology "Internet telephony" found its way into the commercial software and service provider market. The motivation was and still is to use the Internet as a means of toll bypassing, especially in countries with high long distance and international telephone rates due to limited competition in that market.

In the mid 1990s Internet telephony was sometimes portrayed as the imminent death to telephone operators. That scenario did not materialize for a number of reasons: Thanks to liberalization and the abundance of available fiber capacities prices for long distance and even international telephone calls have dropped sharply in many markets reducing the financial incentive of toll bypassing. Secondly, until only very recently, almost all end users have been using dialup connections for Internet access, preventing most of them from having a permanent online connection that would have given then reliable reachability for incoming Internet telephony calls. Also, while narrow-band Internet access in principle is sufficient for good quality voice communications over IP it requires some prioritization mechanism which are not available in today's Internet. Thus Internet telephony was quickly dismissed as unreliable and qualitatively inferior. Last, the lack of standardized Internet telephony protocols that would have ensured interoperability between various applications as well as the technical problems of determining one's intended conversation partner's availability and location on the Internet greatly limited their usefulness. Consequently, so far, for end customers such applications are a niche product at best.

Nevertheless, Voice over IP (VoIP), i.e. the technology of transmitting voice conversations over an IP network (not necessarily the Internet), quietly matured over the years. While originally the notion of saving call charges prevailed, later the focus of VoIP deployment shifted towards reducing operational cost by replacing proprietary and thus expensive to maintain classical telephone systems and, since data application and networks had become an essential part of enterprises, by doing away with the need for separate on-site telephony network infrastructure and support staff. Standardization efforts for VoIP led to the development of the H.323 family of protocols by the ITU and somewhat later SIP (Session Initiation Protocol) by the IETF (Internet Engineering Task Force). While both protocols are similar in their basic approach towards VoIP by separating call setup between terminals (e.g. IP telephones) from the actual exchange of media (e.g. an audio stream of a conversation) most people will probably agree that although H.323 due to its earlier start today still enjoys the bigger market penetration, SIP, thanks to its simple yet elegant, "Internet-like" design will ultimately prevail.

The emergence of these open VoIP standards has created the opportunity for new entrants to the previously highly proprietary telephony equipment market (and likewise poses a severe long-term threat for vendors of traditional telephony systems). As with other IP-based technologies, companies have realized the potential of Open Source products (and in particular Linux) as a basis for their development activities. Some of them, for example Cisco Systems with their VOCAL project, have recognized the Open Source model as a means of efficient and quality product development.

IP telephony equipment not only holds the promise of reducing investment and maintenance costs but moreover, thanks to its open nature, greatly facilitates the development of new services particularly in the field of Computer Telephony Integration (CTI), for example the seamless integration of a call center into an enterprise database and directory environment. In the domain of voice carriers, the development of robust VoIP equipment together with the tremendous growth of data traffic has led to the concept of the Next Generation Network, or NGN for short. The idea of the NGN consists of replacing the traditional TDM-based voice backbone of a carrier with a

VoIP-based backbone, ideally using a single IP network for both voice and data services. In the latter scenario of a "converged" network voice telephony becomes just another IP service, reversing the situation of the early days of data communications when data connections were realized on top of voice networks.

Today, even incumbent carriers are seriously considering or are already starting to deploy Next Generation Networks, often encouraged by vendors of IP equipment who have realized their opportunity of extending their scope of business. Proponents of NGN technology point out that it not only helps the operators to reduce cost but enables the deployment of new services. For consumers however, the transition to NGN first of all does not bring about any significant change as long as customer network access and the end user equipment remain traditional telephony technology. In other words, the mere transition to NGN does not touch upon the traditional telco business model since it does not affect ownership of the customer but rather provides operators with a new, potentially more cost effective network technology. In this respect, the introduction of NGN is somewhat comparable with the replacement of analog switching technology by digital carrier voice networks. It did enable a number of new services, but did not challenge the telephony value chain.

A real change however is brought about the increasing availability of broadband Internet access, typically via DSL or cable, for both enterprise and residential customers. Access is typically charged by traffic volume (or customers just pay a flat monthly fee), no longer by connection time. This allows users to have permanent or quasi-permanent connectivity to the Internet instead of "dialing up to the Net". Always-on access already did and will continue to change the way people make use of the Internet.

Broadband Internet access eliminates many of the obstacles Internet telephony has been struggling with in the past. Since it provides "always-on" functionality it has the potential of delivering constant availability for incoming communications requests just like a traditional telephone line. Broadband access, together with fast, cheap personal computers boosts the use of large-bandwidth applications, so ISPs are impelled to provide these customers with adequate IP backbone capacity. As a result, the available bandwidth is typically significantly larger than that required for an audio stream. New popular real-time applications such as online gaming create additional demand for high-quality IP connectivity. So even with the absence of any strict quality of service mechanisms for IP broadband Internet access typically allows for IP-based voice telephony of good or, with increasing available bandwidth, even excellent quality.

This development provides companies with the opportunity of offering telephony services over IP to broadband Internet subscribers, bypassing the local telco’s monopoly on the last mile. While it might seem natural for the local ISP to come up with such a "value added" service offering, and in fact many already do, there is no such strict technical requirement. Just as with any other IP services, in principle anyone with adequate Internet connectivity can offer them.

Companies like Vonage in the US have built their business around that idea. They are providing their customers with a SIP-based VoIP terminal adapter to be connected the residential IP (cable or DSL) router. Up to two conventional telephones can be plugged into that adapter, making it a small IP-enabled home PBX. Using any broadband Internet access (which Vonage does not provide) users can make outgoing calls via Vonage’s servers. Vonage provides a gateway to the PSTN (i.e. the plain old telephone network) enabling their customers to use their telephones just as if they were connected to a conventional phone line. Moreover, Vonage provides conventional telephone numbers for their subscribers, enabling them to receive incoming phone calls which are terminated at their local SIP adapter. Ideally, users do not even realize that they are using Voice over IP for their telephone calls. Vonage thus has assumed exactly the role of the traditional telco, just making use of a different access technology. Since the company is in control of their subscribers’ incoming and outgoing calls the traditional telecom business model still applies, albeit with a new provider.

Other players will soon pursue a far more radical approach. With the ever- increasing popularity of Internet-based messaging and online gaming it is only a matter of time until high quality voice functionalities are integrated into these applications or even new IP-enabled voice communications consumer devices will appear on the market that have the potential of substituting conventional voice telephony. Microsoft, for example, has started to offer a voice communications kit for their Xbox gaming machine. It enables subscribers of Microsoft’s XBox Live online gaming service to make voice conversations with one another over that platform. Besides, that company is busily integrating telephony functionality into their Windows CE operating system.

Do consumers need a telephony service provider?

With growing proliferation of always-on broadband Internet access so-called Integrated Access Devices (IADs) will increasingly appear on the market which allow people to use their Internet access for voice telephony much like they do today with the conventional telephone network. The interesting question will be whether these devices will come bundled with specific service offerings and whether customers will be forced to use specific platforms or providers for their Internet-based telephone calls. Under the latter scenario the dominating position within the telecommunications value chain would just pass from the network operators to user device (or software) vendors.

The acceptance of any public communications network highly depends on the number of communications partners users can reach over it. In the case of the Internet the open, standardized nature of its communications protocols and its liberal policies ensured success over closed, proprietary data communications platforms. With the development of SIP (Session Initiation Protocol) in the past years a powerful IETF standard for Internet-based messaging has been made available for IP-enabled communications devices and applications. In some respect, SIP extends the philosophy of electronic mail (i.e. Internet-based non-real-time messaging) to all sort of real-time messaging, including but by far not limited to voice telephony. Strictly speaking, SIP as a messaging protocol deals with the setup and control of communication sessions only, not with the actual media exchange (for the latter job, a stack of other suitable protocols has been standardized). This approach helps to maintain a simple, yet powerful protocol and ensures its far-ranging applicability. SIP's addressing scheme is based on the proven, scalable, existing world-wide DNS (domain name service) system. SIP addresses, which assume the role of the telephone numbers from traditional telephony, thus bear a high degree of similarity to email addresses, and since they are independent from the actual communications media (voice, video, text, etc.) they can naturally provide the basis for true unified messaging (UM).

SIP solves the problem of mobility and locating users that do not necessarily have a fixed IP address by allowing them to temporarily register with a SIP server (which can be located anywhere on the Net), much like connecting to a POP3 or IMAP server to retrieve mail. The SIP server will then forward any incoming connection requests for that user to the present location on the Internet. Since only the session setup and control is handled by the SIP server while the actual media stream is typically passed directly between the two parties' computers or IP phones SIP servers require relatively little resources. SIP services, just like web or email services can thus be offered by anyone with Internet access and control over a DNS domain. Due to the flexibility of the SIP protocol users can easily combine SIP services (e.g. voice mail) from any provider. Whether providers will succeed in bundling services and assuring customer loyalty becomes primarily an issue of marketing, not of technology. Bare voice transport as simply yet another comparatively low-bandwidth service in an open IP environment will be a commodity for which consumers will no longer be willing to pay by the minute.

In short, SIP has the potential of turning the Internet into a universal, open platform for real-time messaging. All it takes for broadband Internet users to participate is a suitable SIP client (i.e. a software application running on the user's PC or a messaging device such as an IP phone) and a SIP address and corresponding SIP server to register with. The latter service can be expected to be (and in fact already is) available for free or a very small monthly fee from many providers; or users with their own DNS domain might even run their own SIP server, ensuring their reachability. SIP makes voice telephony or other types of real-type communications a peer-to-peer application, essentially returning to the original aim of the Internet: enabling direct communications between any of the systems connected to the network.

This implies the question whether there is still a role for a telephony service provider. Even if the use of SIP addresses will gain significant popularity for voice communications over the Internet it will probably be a long way until they become an accepted mainstream form of communications. It is hard to imagine that Internet telephony will replace the traditional phone network altogether in any foreseeable future, nor that people and businesses will completely cease using traditional telephone numbers for their reachability. Therefore, even those users who completely switch to Internet telephony will continue to subscribe to some form of gateway to and from the PSTN. Such gateways however can in principle be located anywhere in the Net and run by anyone, not necessarily just the local ISP or telco. The speed at which non-traditional addresses, i.e. SIP addresses or some equivalent addressing scheme, become accepted as a replacement for telephone number for voice communications is one important factor that determines how fast telcos will loose their influence.

Aside from the social acceptance of new telephone "numbers" there is another issue that makes it doubtful whether a true global peer-to-peer scenario will find widespread use outside the group of Internet hobbyists. Since the Internet so far does not provide any global user authentication mechanisms peer-to-peer telephony is destined to become an easy victim to all sorts of malicious or junk calls, especially since IP-based phone calls could easily be automated and placed at negligible cost. Although people have been putting up with that very problem when it comes to email the effect would probably be far more damaging on real-time communications. Moreover, unlike electronic mail, which was a completely new service when the general public became aware of the Internet in the 1990s, Internet voice telephony must stand the comparison with traditional telephony in terms of quality, security and reliability if it is to replace it. Some sort of caller authentication mechanism therefore represents an important prerequisite. Such a requirement however does not only exist for telephony but for many other applications as well. This makes it easily conceivable that such authentication services for Internet applications and users will successfully be introduced in the future. Whoever provides and has control over these services will assume part of the role that telephone companies still hold in the traditional telephony world.

Can Open Source make a difference?

Since with IP as the network protocol the service creation occurs on the terminal device (e.g. the IAD, the PC or the IP-enabled telephone) it is these devices, no longer the network, that determine what services and providers the users see and can subscribe to. Domination over these devices thus implies control over the telecommunications services market.

Open Source can give enterprises and consumers the option to make use of the opportunities an IP-based communications infrastructure offers. In a closed source environment software vendors will be tempted to promote or even enforce the use of a limited set of services from specific providers only according to their economic and strategic interests, limiting consumers' choice and competition. On the service provider side, the Open Source model and the existing code base will help enterprises to rapidly develop new services and to tailor and combine services according to their clients' needs. Open Source can thus be expected to extend its success from traditional non-voice Internet services, many of which, such as DNS or electronic mail, rely almost entirely on Open Source implementations, to IP telephony services. Security and reliability requirements, which are particularly high for telephony services, also favor the Open Source development model for such applications.

What about mobile telephony?

So far, considerations have focused primarily on fixed line telephony services. With mobile phone services, the situation is still somewhat different, although ultimately a similar development can be envisioned. With the proliferation of 2.5G and particularly 3G (i.e. GPRS and UMTS) mobile services, network operators aim at selling data services to customers in order to increase revenues. Although these network architectures are still well within the traditional "carrier-centric" model, allowing operators to retain ownership over the customer, they pave the way for acceptance of IP-enabled mobile devices. Operators must make sure that new data communications services, which one the one hand they hope will enable them to sell new content-based services to subscribers, will on the other hand not lead to customers making use of that data connectivity to substitute for traditional services such as SMS or ultimately, with the help of VoIP-capable mobile devices, voice telephony, i.e. services which measured by data volume today are extremely highly-priced and thus profitable to operators. Such a cannibalization scenario requires the availability of the respective functionality in end user devices. Here again Open Source development might empower consumers to make use of their devices in the best of their own interests.

An even bigger threat for mobile network operators comes from the emergence of WLAN (Wireless LAN) hotspots, not only to data services that the owners of UMTS licenses have put their hopes on for consolidation of their business cases, but beyond that to the traditional voice telephony business, i.e. their present cash cow. Although it is still unclear today what the wireless data market will look like and who will be the dominating players in a few years (traditional network operators, local hotspot operators, or vendors of mobile devices or respective software) there should be little doubt that the availability of comparatively inexpensive always-on broadband network connectivity and IP as an open network protocol will lead to a similar development as in the fixed line world. Recent standardization efforts for Fixed Wireless Broadband (IEEE 802.16a) and particularly Mobile Wireless IP Access (IEEE 802.20) might pave the way for further deployment scenarios.

Outlook

The perspective of the increased proliferation of inexpensive broadband "always-on" Internet access leading to a qualitatively new kind of access to consumers and thus undermining traditional business models not only holds for telephone companies. In principle, any established communications and information infrastructure, whether it is for example radio or cable television, is potentially affected by the competition of IP as an open, universal network protocol. Internet radio has already developed into a serious player in the media industry. Compared to conventional radio it offers the advantage to the entertainment industry of not having to apply for scarce frequency resources (and often licenses) and at the same time enables potentially global reach. It is only a matter of typically available bandwidth for digital television to become a mainstream Internet service. In the field of Video on Demand distribution via Internet is already a reality, turning broadband Internet access into a direct competition to plans of cable network operators of selling to consumers personalized premium entertainment as a lucrative value-added product for their networks. Just as telephone companies, cable network operators are about to loose their technical monopoly for such "higher-value" services and are endangered of seeing themselves confined to the role of yet another high-speed Internet access provider, i.e. of having to live from the revenues of a "bit transport provider".

In contrast to radio or television operators, the added value telephone companies offer to their customers is not some content but merely the call setup (i.e. location of the callee) and the voice transmission. The fact that on the one hand this kind of service, measured by its data volume, is extremely highly-priced and on the other hand could be provided by intelligent end-user devices with the help of the existing DNS infrastructure, makes the telco business model particularly vulnerable - or, depending on the perspective, attractive - to substitution by IP-based technology and the consequent recreation of the value chain.

Today’s telcos will increasingly find themselves trapped by the success of high-speed Internet access. Even if they recognize that threat there ultimately seems to be little what they could do to prevent it. With a mix of technical as well as legal and regulatory obstruction and clever marketing they are arguably in a good initial position of delaying that process. Over time however, pressure from other influential industries such as entertainment or consumer device vendors as well from public and politics, who realize that failing to adopt new communications schemes will put their countries behind other societies, will prevail. That leaves for telcos the choice to either transform themselves into an entertainment and service provider amongst a competition of many others, or to confine themselves to the role of the operator of the last mile, maybe providing basic IP connectivity in addition.

It will be interesting to see whether other players will be able to take over the telecom companies’ dominating role in telecommunications service markets or whether consumers will be capable of making use of the freedom of choice the Internet architecture model offers them. Just as with computing services so far, Open Source has the potential to enable both consumers and businesses to efficiently solve their communications needs, including voice telephony, the way they want.

Acknowledgements

The author would like to thank his colleagues from Detecon’s "Strategic Telecommunications Services and Architectures" group for numerous valuable discussions on that subject over the past months.