

An Experimental Study of the Skype Peer-to-Peer VoIP System

Paper report

Georgios Mathioudakis

Skype is one of the most popular and valuable applications on the web nowadays. It manages to provide services of VoIP (voice over IP), instant messaging and file sharing in an easy-to-use package that works even behind Network Address Translators (NAT) and firewalls. This paper provides a measurement study of the Skype P2P VoIP network and one of its main goals is to understand how P2P VoIP traffic in Skype differs from traffic in P2P file-sharing networks and from traffic in traditional voice-communication networks.

Skype network is closely related to KaZaA network. Both companies were founded by the same individuals, there is an overlap of technical staff, and that much of the technology in Skype was originally developed for KaZaA. Network packet level analysis of KaZaA and of Skype supports this claim by uncovering striking similarities in their connection setup. Both Skype and KaZaA use Supernode-based peer-to-peer networks that organize participants into two layers: *supernodes*, and *ordinary nodes*. Typically, supernodes maintain an overlay network among themselves, while ordinary nodes pick one (or a small number of) supernodes to associate with. Supernodes also function as ordinary nodes and are elected from amongst them based on some criteria, such as node's spare bandwidth level and public reachability.

Ordinary nodes in Skype send control traffic including availability information, instant messages, and requests for VoIP and file transfer sessions over the supernode peer-to-peer network. If the VoIP or file transfer request is accepted, the Skype clients establish a direct connection between each other. If one client is behind a NAT, Skype uses connection reversal whereby the node behind the NAT initiates the TCP/UDP media session regardless of which end requested the VoIP or file transfer session. If both clients are behind NATs, Skype uses STUN-like NAT traversal to establish the direct connection. In the event that the direct connection fails, Skype falls back to a TURN-like approach where the media session is relayed by a publicly reachable supernode. This latter approach is invoked when NAT traversal fails, or a firewall blocks some Skype packets.

In order to understand the Skype network, three experiments were performed in parallel. In the first experiment a supernode's activity was observed for 135 days and 13GB of data sent and received were captured. The second experiment used the Skype client's supernode-cache in order to discover as more as possible supernode addresses. The last one gathered snapshots of which supernodes were online at a given time. Using the list of supernodes from the second experiment application-level pings were sent, to check randomly selected supernodes status.

Churn in P2P networks is the continuous process of nodes joining and leaving the system. In Skype there is very little churn in the supernode network. The measurements showed that there are large diurnal variations with peak usage during normal working hours and significantly reduced usage (40% - 50%) at night. In addition, there are weekly variations with 20% fewer users online on weekends than on weekdays. As has been observed widely for interactive applications like telnet, web, and email, node arrivals in Skype are concentrated towards the morning, while departures are concentrated towards the evening. The median supernode session time is 5.5 hours.

Skype uses spare network and computing resources of hundreds of thousands of supernodes, and little additional infrastructure to handle calls, as compared to traditional telephone companies and wireless carriers who rely on expensive, dedicated, circuit-switched infrastructure. However, Skype supernodes use very little bandwidth most of the time. Measurements showed that a supernode is engaged in relaying data 9.6% of the time. For relayed data, the supernode uses 60 kbps in the median.

As observed Skype calls seem to last longer than calls in traditional telephone networks and files transferred seem to be smaller than in file sharing networks. Particularly, the median Skype call lasted 2m 50s, while the average was 12m 53s and the median file transfer size was 346KB. One reason for the long lasting calls is that Skype-to-Skype VoIP is totally free while phone calls are charged. Altogether, it seems that Skype users appear to behave differently from file-sharing users as well as traditional telephone users.