Institution: University College London

Unit of Assessment: 11 – Computer Science and Informatics

Title of case study: SIP/SDP as an enabler of real-time internet communication

1. Summary of the impact

Pioneering research at UCL Department of Computer Science (CS) into multimedia communications over the Internet led directly to the development of central techniques used in voice-over-IP (VoIP), videoconferencing, and instant messaging. Millions of people worldwide today use applications that incorporate these techniques. In particular, UCL CS created the Session Initiation Protocol (SIP) and the Session Description Protocol (SDP), two Internet standards that comprise the primary way multimedia calls are established on the Internet. They are at the core of products made by Microsoft, Apple, Cisco, Siemens, and Polycom, among many others, and are used in most 3G mobile telephone networks. Implementing the technology reduces costs for businesses, with Oracle, for example, realising $18 million in savings since 2010.

2. Underpinning research

In the early 1990s, the conventional wisdom was that real-time communications such as telephony and videoconferencing required circuit-switched or virtual-circuit networks. Many fundamental research questions were raised about whether – and how – one could instead conduct such communications over packet-switched networks such as the Internet. These included questions, for example, about how to ensure call quality, efficiently distribute audio and video to multiple participants, and locate users and establish calls between them. Specific challenges within user location and call establishment included minimising call setup latency, preserving users' privacy during decentralised user location, and providing an extensible mechanism to allow the negotiation of diverse types of sessions. This case study focuses on underpinning research at UCL that solved these problems in user location and call establishment for Internet multimedia calls, culminating in the design of the Session Initiation Protocol (SIP) and Session Description Protocol (SDP).

This underpinning research took place under the EU MICE project (starting in 1992) and EU MERCI project (starting in 1995), as part of which researchers at UCL conducted extensive research into and subsequent pilot studies of Internet multimedia conferencing. The work incorporated three intertwined research strands: algorithm and protocol design; building prototypes; and active participation in the Internet Engineering Task Force (IETF), the technical standards body that codifies protocols used for communication over the Internet. Contributions to the IETF help increase the chance that research impacts practice. At the IETF, intensive discussions with equipment vendors, network operators, and experienced protocol designers yield insights beyond those typically gleaned from experiments in a university laboratory, as well as highlight factors that impact on a design's deployability.

SIP/SDP provide low-latency call setup by adopting an optimistic codec and parameter negotiation mechanism. The protocols, which subsequently found such diverse uses as instant messaging and file transfer, provide an extensible solution for the types of session being negotiated. Finally, SIP/SDP allow extensive use of proxy servers to decentralise the processing of calls and enable a framework for privacy-preserving user-location, no matter which of several devices a user currently employs. This enabled telecommunications providers to embed SIP within legacy networks using proxies as gateways, and eased the transition from circuit-switched to packet-switched telephony.

As this work on SIP/SDP matured, Mark Handley introduced the new protocols at the IETF for consideration as possible Internet standards. SDP's first draft specification was submitted in November 1995 [3], and SIP's in February 1996 [2]. Today, these protocols are extremely widely implemented and used.

Mark Handley wrote the SDP specification, which originated in a generalisation of and set of extensions to an earlier protocol implemented but never specified by Van Jacobson at LBNL. Mark Handley was the main author of the SIP version 1.0 specification [2] submitted to the IETF in February 1996 (co-authored with Eve Schooler of the California Institute of Technology and initially known as the Session Invitation Protocol). Henning Schulzrinne (today of Columbia University and the US Federal Communications Commission, then based at the Gesellschaft für Mathematik und Datenverarbeitung [GMD]) wrote a competing proposal called SCIP, also in Feb 1996. Handley published a workshop paper in October 1996 that further described the ongoing work on SIP/SDP [1]. In late 1996, Handley and Schulzrinne merged SIP v1.0 and SCIP to form SIP v2.0. Although this version was later enhanced by additional specifications clarifying SIP's use in particular circumstances, SIP v2.0 is essentially the protocol used today. Most of its protocol semantics came from SIP v1.0, as did its use of SDP to describe sessions, but the syntax and ability to control sessions after establishment came from SCIP. SIP v2.0 resulted from roughly equal contributions made in Handley's work at UCL and Schulzrinne's work at GMD. The SIP/SDP standards have continued to be updated [4]; the most recent revision of SDP was re-published as RFC 4566 in July 2006 [5].

From 1991 to October 1996, Mark Handley was a research fellow at UCL. He returned to UCL in 2003 as Professor of Networked Systems. A unique corroboration of Handley's central role in the development of multimedia communication protocols for the Internet was the award to him of the 2012 IEEE Internet Award, a career-long achievement award given to the handful of technology leaders deemed by the IEEE to have most impacted the design and deployment of the Internet. The citation for Handley's award reads: "For contributions to Internet multicast, telephony, congestion control, and the shaping of open Internet standards and open-source systems in all these areas."

3. References to the research

References 1, 2 and 3 best demonstrate the quality of the research.

An early research paper describing multimedia research at UCL CS, including SIP:

[1] Mark Handley, Applying Real-Time Multimedia Conferencing Techniques to the Web; World Wide Web Consortium Workshop on Real Time Multimedia and the Web, October 1996
http://www.w3.org/pub/WWW/AudioVideo/9610_Workshop/paper32/paper32.ps
http://www.w3.org/AudioVideo/RTMW96.html

Initial versions of the SIP and SDP specifications, which form the basis of all subsequent versions of these protocols:

[2] Mark Handley, Eve Schooler, Session Invitation Protocol, Internet Draft draft-ietf-mmusic-sip-00, 22 Feb 1996, www.cs.columbia.edu/sip/drafts/mmusic/draft-ietf-mmusic-sip-00.pdf

[3] Mark Handley, Van Jacobson, SDP: Session Description Protocol, Internet Draft draft-ietf-mmusic-sdp-01, 22 Nov 1995, http://tools.ietf.org/html/draft-ietf-mmusic-sdp-01

Like most standards, these protocols have continued to evolve over time. The most recent versions are:

[4] J. Rosenberg, H. Schulzrinne, G. Camarillo, A. Johnston, J. Peterson, R. Sparks, M. Handley, E. Schooler, SIP: Session Initiation Protocol; RFC 3261, June 2002;
http://tools.ietf.org/html/rfc3261; Citation count (Google Scholar, mid-April 2013): 6219

[5] M. Handley, V. Jacobson, C. Perkins; SDP: Session Description Protocol; RFC 4566, July 2006
http://tools.ietf.org/html/rfc4566; Citation count (Google Scholar, mid-April 2013): 2667

4. Details of the impact

Today, SIP is a key enabler of essential functionality within voice-over-IP (VoIP) and videoconferencing systems, as the standard interoperable method by which VoIP and videoconferencing systems signal and establish calls. Since 2008 it has been of benefit to the many millions of users who place multimedia calls, as well as the many vendors – including Microsoft, Apple, Cisco, Siemens, and Polycom – who market the software and hardware that underpins these calls.

The market for these products, all of which interoperate by virtue of SIP/SDP, is substantial. A report by Visiongain, a telecommunications market research firm, estimates global VoIP revenues at $65 billion in 2012 [d]. IBISWorld, another telecommunications market research firm, estimates that the size of the US VoIP market has grown by 15.3% annually in the five years to July 2013 [e]. As Baset et al. state in their February 2012 journal article, “It [SIP] has been an unquestionable success in becoming the lingua-franca for voice calling between vendor products and between domains.” [a, p. 101] Although SIP/SDP are by no means the entirety of the technologies encompassed by VoIP and videoconferencing, SIP is the dominant open standard for the call signalling functionality in VoIP and videoconferencing. Baset et al. note: “There are hundreds of SIP product vendors, thousands of deployed SIP networks, and millions of SIP hosts in deployment today.” [a, p. 93]

SIP/SDP’s standardisation in the Internet Engineering Task Force (IETF) and publication as Requests for Comments (RFCs) played a central role in their dissemination to the broad community of vendors of software and hardware for Internet-based communication. Many of these vendors actively participate in the IETF standards body, and nearly all monitor developments there in order to keep abreast of and influence new Internet protocol designs. RFCs are standards documents that specify the protocols used in Internet communication. They undergo a rigorous process of review by Internet researchers and technologists. SIP/SDP were proposed sufficiently early in the development of Internet-based multimedia communication that, when vendors first began designing products in this area, they were available as known, standardised solutions for the core problems of call routing and establishment, and for describing a session's content.

As documents, RFCs are rigorous specifications of protocols, but somewhat abstract. Often vendors are more likely to adopt a protocol when there is a prototype implementation of it that they can examine and experiment with, since reading and running code reveals many nuances that aren't immediately apparent in an English-language specification. Accordingly the prototype SIP/SDP implementation, distributed as open-source software in Handley’s Session Directory, Revisited (SDR) session directory tool and widely used in the 1990s by many thousands of Internet multimedia researchers, contributed to SIP/SDP's adoption. Since its adoption, as Baset et al. note:

> SIP has been implemented in hundreds (if not thousands) of different products. It has been implemented in phones, ranging from IP hardphones to soft clients to telephony adapters. It has been implemented in PSTN gateways, from single port analog gateways to massive, carrier-grade SS7 gateways. It is part of many enterprise PBX products, from small-scale small-medium-enterprise solutions to large multinational IP PBXs. It is a feature on nearly every carrier softswitch. Most firewalls have a SIP module. Indeed, SIP has given rise to entirely new product categories. Session Border Controllers (SBC) were born of industry needs around inter-domain SIP deployment. [a, p. 99]

These impacts have been on-going through the entire REF impact period and have, as Baset et al. suggest, facilitated the development and improvement of a great many products used by a great many people around the world. There are, however, two very large such constituencies who have benefited both particularly directly and particularly significantly from SIP/SDP. One, as noted by Baset et al., is the many millions of individuals and enterprises who today place multimedia calls that are routed, signalled, and described using SIP/SDP. Some of those calls are placed over the Internet; others are placed over mobile telephone networks that have adopted Internet protocols for some of their functionality. This class of calls includes VoIP telephony and videoconferencing.

Wired telephone handsets that use VoIP are now widely used in business and home settings alike. BT, Vodafone and all major telecoms companies provide SIP trunking capabilities, allowing businesses to directly connect SIP-based Private Branch Exchange (PBX) systems to the global phone network in a cost-effective manner, circumventing the need for conventional dedicated telephone lines, and providing access to advanced features. Thus, for example, by connecting to its worldwide audio conferencing service provider via a SIP trunking service, Oracle Global IT “realised a cumulative $18 million in cost savings since 2010, while experiencing a 32% increase in minutes of use to over 60 million minutes per month.” [f] Moreover, as Baset et al. further state, “SIP has been deployed by dozens, if not hundreds, of service providers, and runs within countless enterprises. Billions of minutes of traffic are carried on SIP" [a, p. 99]. Commercial benefits comparable with those within Oracle Global IT have accrued to many of those enterprises.

Many 3G and newer mobile telephone networks use SIP for call signalling of enhanced features such as user location and video calls; as such, SIP also underpins the routing and signalling of calls for the world's vast population of 3G mobile telephone users. Videoconferencing, another application in which SIP/SDP provides call signalling, is also very widely used, not only in high-end telepresence systems and VoIP hardware phones, but also in software available for laptops or smartphones.

The second main constituency benefitting from SIP/SDP since 2008 consists of the many vendors who sell software and hardware for VoIP and videoconferencing that route, signal, and describe calls with SIP/SDP. Many of these vendors discuss SIP-related issues in an industry forum devoted entirely to SIP [g]. Today, Cisco, Siemens, and Polycom manufacture VoIP handsets that implement SIP/SDP; Cisco and Siemens also make and sell VoIP infrastructure such as SIP proxy servers and SIP registration servers, and Cisco and Polycom now make high-definition telepresence systems using SIP [c, h]. To date, Cisco has shipped over 50 million SIP-enabled IP phones. A statement from Cisco's Chief Technology Officer for the company's Cisco Collaboration division noted: "SIP is the core call signaling protocol across our portfolio today and it has helped us interoperate between our products and with other vendors. Cisco collaboration is a nearly 4 billion dollar business of Cisco representing over 8% of revenue." [h]

Myriad mobile telephone base station vendors also implement SIP for 3G (and later) call signalling. Furthermore, Apple's popular FaceTime videoconferencing software, which is now included with every iPhone, iPad, and MacBook, signals calls using SIP, a fact noted by Steve Jobs in his introduction of the iPhone 4 during a keynote speech at Apple's Worldwide Developers' Conference (WWDC) in 2010 [b].

5. Sources to corroborate the impact

[a] Confirmation of SIP's central role in interoperability between vendor products can be found in Baset, S. A., Gurbani, V., Johnston, A., Kaplan, H., Rosen, B., and Rosenberg, J., The Session Initiation Protocol (SIP): An Evolutionary Study, in Journal of Communication, 7(2), Academy Publisher, February 2012, pp. 89-105. http://doi.org/p52

[b] For Apple's use of SIP in FaceTime: http://appleinsider.com/articles/10/06/08/inside_iphone_4_facetime_video_calling.html

[c] All Cisco's telepresence end-systems and multipoint control units listed in Cisco's brochure (http://bit.ly/16NKIVV) use SIP call control; confirmed, as an example, for Cisco TelePresence 3200: http://www.cisco.com/en/US/prod/collateral/ps7060/ps8329/ps8330/ps9573/data_sheet_c78-457905.html

[d] The Visiongain report cited in Section 4 describes the size and nature of the VoIP marketplace. It comments on the market in 2012 and forecasts market value and subscriber numbers for VoIP for 2012-2017. This report is available at: http://www.visiongain.com/Report/854/The-Voice-Over-Internet-Protocol-%28VoIP%29-Market-2012-2017-Prospects-for-Skype-and-Other-Players

[e] The IBISWorld report cited in Section 4 describes growth in the US VoIP market in the five years to July 2013. It is available at: http://www.ibisworld.com/industry/default.aspx?indid=1269

[f] "Oracle Improves Communications and Reduces Costs with SIP Trunking", page 1, http://www.oracle.com/us/industries/communications/enterprise-border-controller-wp-2010610.pdf

[g] SIP has its own industry forum and network operator's conference: http://www.sipforum.org/content/view/22/199/. The SIP forum has 31 full commercial members, including Cisco, Alcatel/Lucent, Ericsson, Microsoft, Nokia, Polycom, Samsung and Siemens.

[h] Statement from Chief Technology Officer, Cisco Collaboration division, confirms Cisco's use of SIP, number of SIP-enabled phones shipped, revenue linked to SIP, and the relationship between SIP and UCL's research. Available on request.