IEEE ComSoc Distinguished Lecture Tour – Europe, June 8-24, 2015

Ying-Dar Lin, IEEE Fellow
National Chiao Tung University, TAIWAN

June 2015

DLT Planning and Itinerary

I packed this DLT (Distinguished Lecture Tour) with five talks and piggybacked it to IEEE ICC'15 in London. These five talks cover four topics, including software defined networking (SDN), traffic forensics, international academic services and research, and research roadmap driven by a test lab. Right after ICC where I participated meetings of some editorial boards, program committees, and technical committees, my first talk was on SDN at University of Surrey near London, hosted by Prof. Zhili Sun who co-chaired with me ICC'15 Next Generation Networking Symposium. Then I flew to Barcelona to give my second talk also on SDN at Polytechnic University of Catalonia (UPC), hosted by Prof. Josep Solé-Pareta, with other participants led by Prof. Jose Marzo from University of Girona. Jose Marzo awarded me the best paper award back in SPECTS'12. After that, I took Spanish high speed rail, AVE, to Zaragoza to give my third talk on international academic services and research at University of Zaragoza, hosted by Prof. Jose Saldana and Prof./Director Juan Ignacio Garces. I have known Jose also in SPECTS'12. Then I was on another AVE to Madrid to give my fourth talk on traffic forensics at University Carlos III of Madrid (UC3M), hosted by Prof. Victor P. Gil Jimenez. I flew to Amsterdam and drove to Louvain to give my final fifth talk on research roadmap driven by Network Benchmarking Lab (NBL) at University of Catholic Louvain (UCL), hosted by Prof. Marco Canini who I've known through editorial works. It was a one-day workshop packed with eight other smaller talks following my keynote.

Four Topics in Five Lectures

The lecture on SDN was an overview. I argued why, where, and when for SDN. Then I illustrated how SDN works in sections of research, standardization, development, and testing. The lecture on traffic forensics summarized a series of 20 research papers of mine done in the past 6 years, ranging from traffic capture on the campus beta site, replay from captured traces, classification leveraging commercial devices, to detection and analysis of intrusions and malware with techniques of signature matching, statistical behavior analysis, and hybrid approach. In the lecture on research roadmap driven by NBL, I reviewed the model of development-driven research and its advantages over literature-driven research. Finally, the lecture on international academic services and research was a non-technical career talk on the motivation and logistics of serving on the editorial boards, program committees, technical committees, special issues, etc., and delivering research with academic or industrial impacts. As it was the end of the semester in Europe, the attendees were mostly faculty members, post-doc researchers, and Ph.D. students.

The number of attendees per lecture was about 20-40. The number of questions asked was 3-6. In summary, I received the positive comments like “very informative talk”, “very logical arguments”, “it clarifies my confusion about XXX”, etc. I myself also enjoyed and appreciated better the rich and diverse cultures of southern and northern Europe.

In-Depth Discussions During and After Lectures

The lectures triggered good questions from the audience. I list major questions and my answers below. For questions similar to the ones listed in my previous DLT reports, they are not repeated here.

1. [SDN] Will SDN practice more cloud computing or more fog computing?

More fog computing than cloud computing. Unlike traditional cloud computing where the number of data centers is quite limited to a few big ones, SDN-NFV would need a hierarchical data centers, with the low-tier data centers transformed from the “central offices” of operators. The reason is that redirected data plane, i.e., virtualized network functions, should not go too far.

2. [SDN] Does the society define API between SDN controllers?

Yes, it’s called east-west-bound API in ONF, as compared to south-bound API between switches and controllers and north-bound API between apps and controllers. It is to build the scalable architecture of a controller hierarchy.

3. [Traffic Forensics] To classify traffic into a PCAP library, why do we need separate devices from multiple vendors?

Each vendor has their own false positives (FPs) and false negatives (FNs). Thus, we need devices from more vendors to jointly reduce FPs and FNs.

4. [Research Roadmap Driven by NBL] What if the redirected data-plane traffic overload NFV, including the virtualized packet classifiers and the virtualized network functions (VNFs)?

A classifier or a VNF could be run on multiple VMs for better scalability. But that requires the management plane to increase or decrease the VM allocation, the control plane to do load balancing, and the data plane to steer traffic to the correct VM.

5. [International Academic Services and Research] How do I know whether I am qualified to serve on a TPC or an editorial board?

If you have published on a conference or a journal consecutively, say more than three times, you could volunteer yourself to its TPC chair or editor-in-chief. Though your request might be judged with other criteria and declined, it is always worth a trial.


Left: Univ of Surrey – lecture attendees

Right: Univ of Surrey – Zhili Sun, me, statue of Alan Turing


Left: University Politechnica de Catalunya (UPC) – lecture attendees

Right: UPC – after-talk lunch


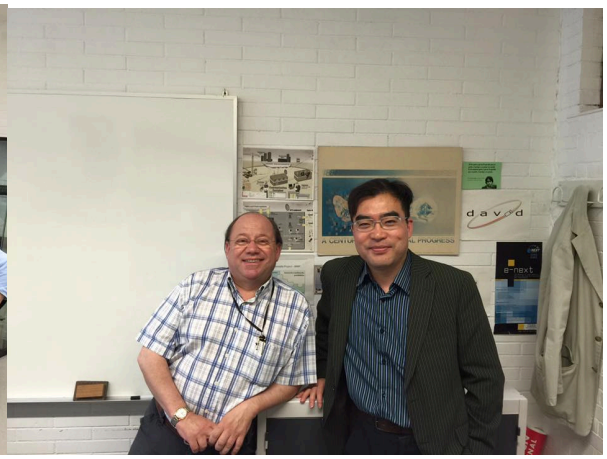

Left: UPC – after-talk panel discussion

Right: UPC – Josep Solé-Pareta and me


Left: Lecture at Univ of Zaragoza

Right: Univ of Zaragoza – Juan Ignacio Garces (Director), me, Jose Saldana

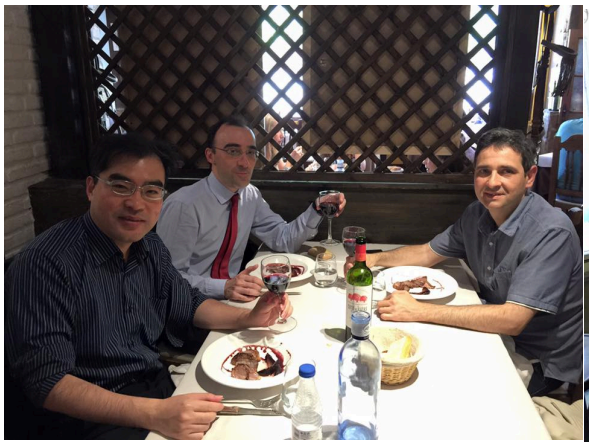


Left: Univ of Zaragoza – me, Jose Saldana, Julian Fernandez-Navajas

Right: Univ of Zaragoza – after-talk dinner


Left: UC3M – lecture attendees

Right: UC3M – Andres Marin Lopez, me, Victor Gil Jimenez, David Larrabeiti

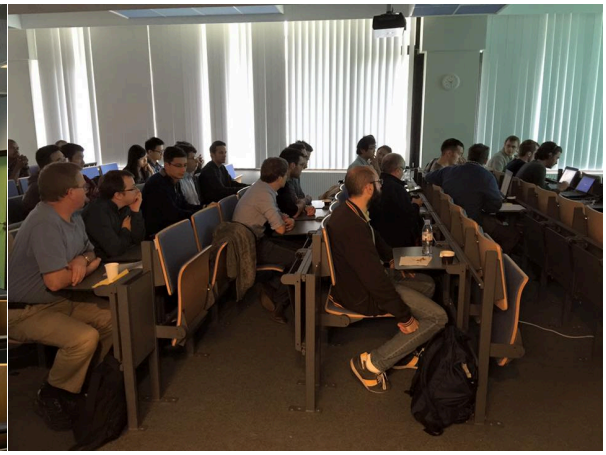

Left: Université Catholique de Louvain (UCL) – lecture attendees
Right: Lecture in UCL

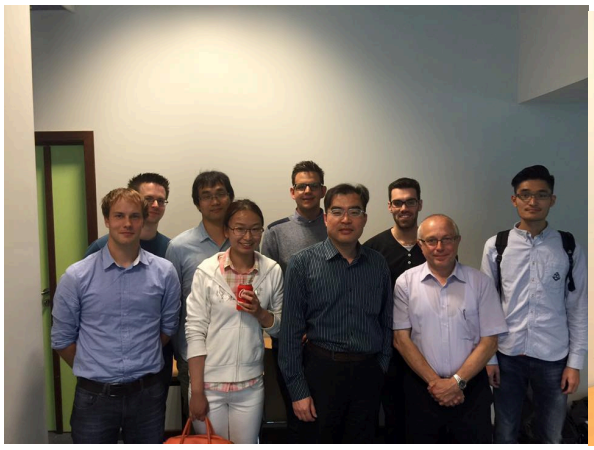


Left: UCL - Luc Claesen and his group
Right: UCL - Luc Claesen and me

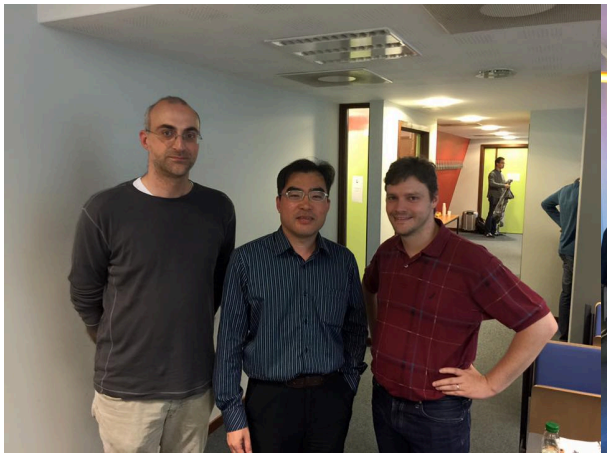


Left: UCL - Ramin Sadre, me, Marco Canini
Right: UCL – lunch break

Appendix:

Talk Title: **Software Defined Networking: Why, When, Where, and How**

Abstract:

The first wave of cloud computing was to centralize and virtualize servers into the clouds, with a phenomenal result. The emerging second wave, named Software Defined Networking (SDN), is to centralize and virtualize networking, especially its control, into the clouds. SDN deployment started from data centers and now expands to the model of "networking as a service" (NaaS) offered by the operators to enterprise and residential subscribers. By centralizing the control-plane software of routers and switches to the controller, and its applications, and controlling the data-plane of these devices remotely, SDN reduces the capital expenditure (CAPEX) and operational expenditure (OPEX) because the devices become simpler and hence cheaper and number of administrators could be reduced. SDN also enables fast service orchestration because the data plane is highly programmable from the remote control plane at controllers and applications. However, as we detach control plane from where data plane resides, new protocols shall be introduced between control plane and data plane, as the southbound API between controllers and devices and the northbound API between controllers and applications. As we further extend the control plane from controllers to applications such as Service Chaining (SC) and data plane from devices to Network Function Virtualization (NFV), newer mechanisms and APIs need to be added to these APIs. We argue why, when, and where SDN would prevail, and then illustrate how to make it happen. We shall introduce the key technology components, including OpenFlow, SC, NFV, and Network Service Header (NSH) and then review the issues on standardization, development, deployment, and research. At the end, the development and deployment experiences of a campus SDN solution for Wi-Fi/switch control and management are shared.

Talk Title: **Traffic Forensics: Capture, Replay, Classification, Detection, and Analysis**

Abstract:

If computer forensics is to identify, preserve, recover, and analyze who did what on a computer, network forensics is to do the same on a network. Compared to network forensics, which has wider forensics targets on devices (e.g., switches, routers, access points, firewalls, gateways) and packets between them, traffic forensics focuses on packets alone. When these devices are black boxes and do not have storage to record what happened, which are often true, traffic forensics then approximates network forensics. In this talk, we present a series of technologies and tools we developed to capture, replay, classify, detect, and analyze traffic. From the architectures of a beta site embedded into an operational campus network with live traffic, to replay captured traffic with stateless or stateful replayers in wired or wireless environments, we

build the basic infrastructure and tools to play with real traffic. A case study is reported to see how effective the accumulated packet traces are in triggering bugs in products under development. Then we present another class of techniques leveraging the domain knowledge of existing products to classify traffic into various applications or malicious intrusions and malware. A classified PCAP library, associated techniques, and their evaluation are illustrated. With these integrated, a case study is reported to redefine security criteria with functionality, robustness, performance, and stability testing, in order to complement existing criteria such as Common Criteria, ICSA, and NSS. As sources of intrusions are often malware carried in application payloads, collect, analyze, and detect malware are the essential ways to build the defense lines. Thus, we present the mechanisms to collect and analyze active and passive malware through honeypot and P2P, respectively. At the end, we present detection mechanisms for traditional malware, Android malware, and Advanced Persistent Threat (APT).

Talk Title: **Research Roadmap Driven by Network Benchmarking Lab (NBL): Deep Packet Inspection, Traffic Forensics, WLAN/LTE, Embedded Benchmarking, Software Defined Networking, and Beyond**

Abstract:

Most researchers look for topics from the literature. But our research has been driven mostly by development which in turn has been driven by industrial projects or lab works. We first compare three different sources of research topics. We then derive two research tracks driven by product development and product testing, named as the blue track and the green track, respectively. Each track is further divided into development plane and research plane. The blue track on product development has fostered a startup company (L7 Networks Inc.) and a textbook (Computer Networks: An Open Source Approach, McGraw-Hill 2011) at the development plane and also a research roadmap on QoS and deep packet inspection (DPI) at the research plane. On the other hand, the green track on product testing has triggered a 3rd-party test bed, Network Benchmarking Lab (NBL, www.nbl.org.tw), at the development plane and a research roadmap on traffic forensics, WLAN/LTE, embedded benchmarking, and software defined networking at the research plane. Throughout this talk, we illustrate how development and research could be highly interleaved. At the end, we give lessons accumulated over the past decade. The audience could see how research could be conducted in a different way.

Talk Title: **Sharing Experiences on International Academic Services and Research**

Abstract:

Just like doing campus academic services on top on teaching and research obligations, a researcher could volunteer to international academic services after years of research, which gains one visibility and opportunities to co-work in all dimensions with other researchers. These services and co-work experiences could in-turn inspire and elevate one’s future research. But

they were seldom discussed publicly in the society. In this talk, based on my 22 years of research and 7 years of international academic services, I'd share my personal humble viewpoints on why, what, when, and how of international academic services and research. The 1st part of the talk reviews incentives and logistics behind serving journal editorial boards, special issues, conference program committees, society technical committees, and other positions. Whether waiting to be invited or to volunteer oneself is compared. The 2nd part first compares research (big R or small r) and development (big D or small d), in academic (with Rd) and industry (with rD), to inspect their motivation. Then I compare the "criteria of survival" and "impacts of lifetime", where the former and the latter are like basketball games and football games, respectively. In the long run, research should be campaigned as a football game instead of a basketball game. Next I share some logistics on (1) evolving independent work model to co-work model, (2) managing research processes from proposals to publications, (3) how to graduate your students on time, and (4) how to campaign for IEEE Fellow. At the end, I list lessons and skills I've learned so far and those yet to be learned by me.

Autobiography:

YING-DAR LIN is a Distinguished Professor of Computer Science at National Chiao Tung University (NCTU) in Taiwan. He received his Ph.D. in Computer Science from UCLA in 1993. He served as the CEO of Telecom Technology Center during 2010-2011 and a visiting scholar at Cisco Systems in San Jose during 2007–2008. Since 2002, he has been the founder and director of Network Benchmarking Lab (NBL, www.nbl.org.tw), which reviews network products with real traffic. NBL recently became an approved test lab of the Open Networking Foundation (ONF). He also cofounded L7 Networks Inc. in 2002, which was later acquired by D-Link Corp. His research interests include design, analysis, implementation, and benchmarking of network protocols and algorithms, quality of services, network security, deep packet inspection, wireless communications, embedded hardware/software co-design, and recently software defined networking. His work on "multi-hop cellular" was the first along this line, and has been cited over 650 times and standardized into IEEE 802.11s, IEEE 802.15.5, WiMAX IEEE 802.16j, and 3GPP LTE-Advanced. He is an IEEE Fellow (class of 2013), an IEEE Distinguished Lecturer (2014&2015), and a Research Associate of ONF. He is currently on the Editorial Boards of *IEEE Transactions on Computers, IEEE Computer, IEEE Network, IEEE Communications Magazine - Network Testing Series, IEEE Wireless Communications, IEEE Communications Surveys and Tutorials, IEEE Communications Letters, Computer Communications, Computer Networks, Journal of Network and Computer Applications, and IEICE Transactions on Information and Systems.* He has guest edited several Special Issues in IEEE journals and magazines, and co-chaired symposia at IEEE Globecom'13 and IEEE ICC'15. He published a textbook, *Computer Networks: An Open Source Approach* (www.mhhe.com/lin), with Ren-Hung Hwang and Fred Baker (McGraw-Hill, 2011). It is the first text that interleaves open source implementation

examples with protocol design descriptions to bridge the gap between design and implementation.