

My research lies at the intersection of usable **security/privacy** and **natural language processing/machine learning**. I am building systems which automatically analyze app permissions and privacy policies to detect abuse of collection and sharing of personal data in Android apps. I previously worked on **mobile systems** and **networks** for 5 years. I developed novel quality-of-service-guaranteed energy-efficient technologies for the **Chrome web browser** and **real-time streaming**.

**EDUCATION**

---

<b>University of Michigan</b> , Ann Arbor, MI	<b>2017 – present</b>
• PhD candidate in Computer Science · Advisor: Professor Kang Shin	
<b>Korea Advanced Institute of Science and Technology (KAIST)</b> , South Korea	<b>2011 – 2013</b>
• MSc in Computer Science · GPA 4.0/4.3 (96.7/100) · Outstanding Master's Thesis Award	
<b>Hanoi University of Science and Technology</b> , Vietnam	<b>2005 – 2010</b>
• BSc in Computer Science · GPA 8.5/10 (top 1% of CS Dept.)	

**INDUSTRIAL EXPERIENCE**

---

<b>Facebook</b> , Menlo Park, CA	<b>Research Intern</b>	<b>May – Aug 2019</b>
• <i>Personalized Federated Learning</i> : Designed the 1 <sup>st</sup> privacy-preserving user representation learning for Federated Learning. My techniques improve 8-51% AUC/accuracy than non-personalized models and provide similar performance with centralized approaches while preserving user privacy. Facebook R&D continued to use my implementation (3000+ lines of code) after my internship. [PyTorch, PyText, Python]		
• This work was submitted to ICLR 2020 and published on arXiv.		
<b>Google</b> , Mountain View, CA	<b>Software Engineering Intern</b>	<b>May – Aug 2017</b>
• <i>Performance metrics for the Chrome web browser</i> : Added congestion tracking metrics to Chrome to quantify page load times and the responsiveness of foreground tabs when the browser loads multiple tabs simultaneously. The metrics are deployed in production. Chrome engineers use them to direct optimization efforts to improve the user experience when loading many tabs. I made 16 code changes (added 3000+ and deleted 500+ lines of code) to the Chrome codebase. [C++]		
• <i>Power consumption estimation from CPU usage</i> : Investigated the correlation between Chrome's CPU usage and power consumption to estimate the power profiles of devices in the wild. [Python, JavaScript]		
<b>Samsung Electronics</b> , Suwon, South Korea	<b>Research Intern</b>	<b>May – Sep 2015</b>
• <i>Transparent cross-device resource sharing</i> : Used real-time streaming protocols (RTSP/RTCP/RTP) to improve camera preview frame rate by 6X and reduce photo capture time across two smartphones by 4X, compared to the state of the art (Rio, MobiSys 2014); developed Unified Resource Management Framework with 11,000+ lines of C/C++ code that manages resources (camera, sensors, and apps) across multiple heterogeneous-platform mobile devices. [C, C++, CMake, Tizen OS]		
<b>Microsoft Research</b> , Beijing, China	<b>Research Intern</b>	<b>Jan – Jun 2014</b>
• <i>Energy-efficient mobile web browsing</i> : Reduced the whole-system energy consumption of the Google Chrome mobile web browser by 24.4% with no perceivable impact on page load time; analyzed inefficiencies, developed, and implemented energy-saving techniques for the Chrome and Firefox web browsers on Android. Compared with state-of-the-art solutions, my techniques leverage browser internals (process/thread structure and resource fetching/processing pipelines) and the heterogeneous multi-core big.LITTLE CPU architecture rather than based on characteristics of web pages. [C++, Python, .NET]		
• This work was published in a top-tier conference (ACM MobiCom) and highlighted on an ACM SIGMOBILE magazine and online news ( <a href="#">Microsoft Research Blog</a> , <a href="#">Phys.org</a> , and <a href="#">On MSFT</a> ); open source on GitHub: <a href="#">energy_efficient_web_page_loading</a> , <a href="#">browser_profiler</a> .		

**ACADEMIC EXPERIENCE**

---

<b>University of Michigan</b> , Ann Arbor, MI	<b>Research Assistant</b>	<b>Sep 2017 – present</b>
• <i>Assisting Automated Comprehension of Privacy Policies</i> : Created the <i>first</i> ML-based fully automated extraction of personal information types from privacy policy documents, using ELMo-BiLSTM sentence classification and BiLSTM-CRF entity recognition; extracted fine-grained word-level information automatically with high accuracy in previously unseen documents using deep learning (better than state-of-the-art solutions which extract coarser-grained paragraphs, need manual refinement, or use rule-based methods). [AllenNLP, PyTorch, TensorFlow, Python]		

- *Energy-efficient multi-link real-time streaming*: Created the first LTE-enabled prototype implementation that aggregates bandwidth over multiple asymmetric mobile wireless interfaces (WiFi & LTE) for energy-efficient real-time delivery; provided quality-of-service (QoS) for high-bitrate video streaming while reducing energy consumption by leveraging Radio Resource Control (RCC) states of LTE. Compared to state-of-the-art solutions, my method does not require any support on existing servers and changes to the existing Internet infrastructure. [C, Java, Android]
- This work was published in a top-tier conference (IEEE RTSS); open source on GitHub: [greenbag](#).

## PUBLICATIONS

---

1. *Federated User Representation Learning*.  
**Duc Bui**, Kshitiz Malik, Jack Goetz, Honglei Liu, Seungwhan Moon, Anuj Kumar and Kang G. Shin.  
arXiv preprint arXiv:1909.12535, 2019.
2. *Active Federated Learning*.  
Jack Goetz, Kshitiz Malik, **Duc Bui**, Seungwhan Moon, Honglei Liu and Anuj Kumar.  
Workshop on Federated Learning for Data Privacy and Confidentiality (in Conjunction with **NeurIPS**) 2019.
3. *Cross-Platform Support for Rapid Development of Mobile Acoustic Sensing Applications*.  
Yu-Chih Tung, **Duc Bui**, and Kang G. Shin.  
ACM International Conference on Mobile Systems, Applications and Services (**MobiSys**) 2018. (**27%** acceptance rate)
4. *Mobile Plus: Mobile platform for Transparent Sharing of Functionalities Across Devices*.  
Sangeun Oh, Hyuck Yoo, Dae R. Jeong, **Duc Bui**, and Insik Shin.  
ACM International Conference on Mobile Systems, Applications and Services (**MobiSys**) 2017. (**18%** acceptance rate)
5. *Demo: Mobile Plus: Mobile platform for Transparent Sharing of Functionalities Across Devices*.  
Sangeun Oh, Hyuck Yoo, Daelyong Jeong, Sooyoung Park, **Duc Bui**, Sungsoo Moon, and Insik Shin.  
ACM International Conference on Mobile Systems, Applications and Services (**MobiSys**) 2016.
6. *Rethinking Energy-Performance Trade-Off in Mobile Web Page Loading*.  
**Duc Bui**, Yunxin Liu, Hyosu Kim, Insik Shin, and Feng Zhao.  
ACM SIGMOBILE GetMobile Magazine (**Research highlights**), April 2016.
7. *Rethinking Energy-Performance Trade-Off in Mobile Web Page Loading*.  
**Duc Bui**, Yunxin Liu, Hyosu Kim, Insik Shin, and Feng Zhao.  
ACM International Conference on Mobile Computing and Networking (**MobiCom**) 2015. (**18%** acceptance rate)
8. *GreenBag: Energy-efficient Bandwidth Aggregation for Real-time Streaming in Heterogeneous Mobile Wireless Networks*.  
**Duc Bui**, Kilho Lee, Sangeun Oh, Hyojeong Shin, Insik Shin, Honguk Woo, and Daehyun Ban.  
IEEE Real-Time Systems Symposium (**RTSS**) 2013. (**22%** acceptance rate)
9. *A Case Study of the Application of Dynamic Symbolic Execution to Real-World Binary Programs*.  
**Duc Bui**, Yunho Kim, and Moonzoo Kim.  
Korea Conference on Software Engineering (**KCSE**) 2012.
10. *A method of verifying web service composition*.  
Thang Huynh, Quynh Pham, and **Duc Bui**.  
ACM International Symposium on Information and Communication Technology (**SoICT**) 2010.

## TALKS

---

- *VinAI research lab seminar series*, October 2019: Federated User Representation Learning.
- *MobiCom 2015 oral presentation*, September 2015: Rethinking Energy-Performance Trade-Off in Mobile Web Page Loading.
- *RTSS 2013 oral presentation*, December 2013: GreenBag: Energy-efficient Bandwidth Aggregation for Real-time Streaming in Heterogeneous Mobile Wireless Networks.
- *KCSE 2012 oral presentation*, February 2012: A Case Study of the Application of Dynamic Symbolic Execution to Real-World Binary Programs.
- *Workshop on Software Analysis for Error-free Computing in Hong Kong University of Science and Technology (HKUST)*, January 2012: A Case Study of the Application of Dynamic Symbolic Execution to Real-World Binary Programs.

## PROFESSIONAL SERVICES

---

- *External reviewer for conferences/journals:* IEEE Real-Time Systems Symposium (RTSS), ACM International Conference on Embedded Software (EMSOFT), IEEE International Conference on Embedded and Real-Time Computing Systems and Applications (RTCSA), ACM Wireless of the Students, by the Students, and for the Students Workshop (S3).

## HONORS AND AWARDS

---

- 2016 **Naver PhD Fellowship** by Naver Corp. (the largest search engine in South Korea) for excellent PhD students.
- 2015 **Qualcomm Innovation Award** finalist, being in the top ten out of 37 final-round projects.
- 2015 **Microsoft Research Asia Fellowship Nomination Award**, being one of 90 students selected from top Asia universities.
- 2014 **Microsoft Research Asia Excellent Award** in the Stars of Tomorrow Internship Program.
- 2014 **Outstanding Master's Thesis Award** by Computer Science Department, KAIST.
- 2010 **Korean Government Scholarship** for the master's program at KAIST.
- 2009 **Vietnam Ministry of Information and Communications Scholarship** for outstanding students.

## REFERENCES

---

(Available upon request)