JUNZHI GONG

Webpage: junzhig.me
(+1) 617-852-7336 ⋄ jug257@g.harvard.edu

EDUCATION

Harvard University, Cambridge, United States Sept. 2018 -
Ph.D. student in Computer Science *Advisor: Minlan Yu*

Peking University, Beijing, China Sept. 2014 - Jul. 2018
B.S. in Computer Science and Technology *Advisor: Tong Yang*
Major GPA: 3.86/4.00 **Cumulative GPA: 3.79/4.00**

RESEARCH INTERESTS

I am currently interested in doing research regarding computer networks & systems, including mobile network systems, network diagnosis, NFV systems, programmable switches, SmartNICs, *e.t.c.* Before attending Harvard, I was interested in network measurement and analysis, data stream processing, and database systems in Peking University.

AWARDS & SCHOLARSHIPS

- National Scholarship, Peking University 2014 - 2015
- Merit Student Honor, Peking University 2014 - 2015
- May 4th Scholarship, Peking University 2015 - 2016
- Learning Excellence Award, Peking University 2015 - 2016
- Honorable Mention in the Mathematical Contest in Modeling 2016
- The Okamatsu Scholarship, Peking University 2016 - 2017
- Learning Excellence Award, Peking University 2016 - 2017

SELECTED RESEARCH PROJECTS

Enabling Resilience in Virtualized RANs
Mentor: Anuj Kalia *Microsoft Azure*

· Virtualized RANs, which run RAN processing on commodity servers instead of specialized hardware, are gaining adoptions in cellular networks. The real-time deadline and black-box nature prevent existing resilience techniques, like virtual machine migration or state replication, from being successfully used in vRANs. This project focuses on proposing an effective resilience solution for vRANs. The key idea is to repurpose existing cellular network mechanisms, including handovers and cell reselections, to provide software resilience for vRANs. Our evaluation shows that our solution achieves minimum disruptions to cellular connectivity during resilience events, while incurring low overhead.

Scalable distributed massive MIMO baseband processing
Mentor: Anuj Kalia, Minlan Yu *Microsoft Research, Harvard University*

· Massive MIMO (multiple-in multiple-out) is a key wireless technique to get higher bandwidth in modern mobile networks such as 5G. The large amount of computation required for massive MIMO baseband processing poses a challenge to the ongoing softwarization of radio access networks (RAN), in which mobile network operators are replacing specialized baseband processing chips with commodity servers. Existing software-based systems for massive MIMO fail to scale to increasingly larger MIMO dimensions with an ever-increasing number of antennas and users. In this project, we present a new scalable distributed system called Hydra, designed to parallelize massive MIMO baseband processing while minimizing the overhead of distributing computation over multiple machines. Hydra's high scalability comes from reducing inter-server and inter-core communication at different stages of baseband processing. To do so, among other techniques, we take advantage of hardware features in modern commodity radios in novel ways. Our evaluation shows that Hydra can support over four times larger MIMO configurations than prior state-of-the-art systems, handling for the first time, 150×32 massive MIMO with three servers.

Detailed NFV Performance Diagnosis
Advisor: Minlan Yu *Harvard University*

· By moving network appliances from hardware to software, Network Function Virtualization (NFV) allows flexible resource sharing among network functions and achieves scalability with low cost. However, due to resource contention, network functions sometimes suffer from performance problems that are hard to diagnose. In particular, when many flows traverse a complex DAG of NFs, it is hard to locate the root causes of a flow experiencing low throughput or high tail latency. Simply maintaining resource counters at individual NFs is not sufficient because resource contentions can propagate across NFs and over time. In this project, we propose a performance diagnosis tool for network functions that leverages queue information between NFs to identify the root causes (i.e., flows, resources, NFs). Our evaluation on realistic NF chains and traffic shows that we can capture many performance problems with high accuracy and low overhead.

Identify Elephant Flows to Increase Accuracy for Network Measurement
Advisor: Tong Yang *Peking University*

· In most network measurement tasks, elephant flows are more important than mice flows. In order to achieve high accuracy for network measurement, an elegant solution is to identify elephant flows in real time, and accurately record their sizes. We aim to design novel algorithms to intelligently separate and guard elephant flows from mice flows.
- Designed novel data structures and algorithms which can separate and guard elephant flows from mice flows.
- Applied the data structure to typical network measurement tasks.

Achieve Better Performance for Sketches in Network Measurement
Advisor: Tong Yang *Peking University*

· Sketch is a compact data structure which can estimate flow sizes in network traffic. However, in real traffic, sizes of most flows are small and only a few flows have a large size, which leads to significant memory waste for sketches. In order to increase memory efficiency for sketches, we aim to propose better designs for sketches.
- Proposed better designs for sketches, which achieve higher memory efficiency.
- Made mathematical analysis for sketches using our new designs.
- Applied sketches using our new designs to typical measurement tasks.

INTERNSHIP

Google Cloud, Sunnyvale, United States Jun. 2023 - Sept. 2023
Software engineer intern *Mentor: Yuliang Li, Devdeep Ray*

Microsoft Research, Redmond, United States Jun. 2022 - Sept. 2022
Research intern *Mentor: Anuj Kalia*

Microsoft Research, Redmond, United States Jun. 2020 - Sept. 2020
Research intern *Mentor: Anuj Kalia*

Carnegie Mellon University, Pittsburgh, USA Jun. 2017 - Sept. 2017
Student Reseach Intern in CMU-PKU Internship Program *Advisor: Peter Steenkiste*

PUBLICATIONS

- **Enabling Resilience in Virtualized RANs with Atlas**
  Jiarong Xing*, Junzhi Gong*, Xenofon Foukas, Anuj Kalia, Daehyeok Kim, Manikanta Kotaru
  MobiCom 2023
  **: equal contributions*
- **Scalable Distributed Massive MIMO Baseband Processing**
  Junzhi Gong, Anuj Kalia, Minlan Yu
  USENIX NSDI 2023
- **Rearchitecting the TCP Stack for I/O-Offloaded Content Delivery**
  Taehyun Kim, Deondre Martin Ng, Junzhi Gong, Youngjin Kwon, Minlan Yu, KyoungSoo Park
  USENIX NSDI 2023

- **Microscope: Queue-based Performance Diagnosis for Network Functions**
  Junzhi Gong, Yuliang Li, Bilal Anwer, Aman Shaikh, Minlan Yu
  ACM SIGCOMM 2020
- **q-MAX: A Unified Scheme for Improving Network Measurement Throughput**
  Ran Ben Basat, Gil Einziger, Junzhi Gong, Jalil Moraney, Danny Raz
  IMC 2019
- **Elastic sketch: Adaptive and fast network-wide measurements**
  Tong Yang, Jie Jiang, Peng Liu, Qun Huang, Junzhi Gong, Yang Zhou, Rui Miao, Xiaoming Li, Steve Uhlig
  ACM SIGCOMM 2018
- **HeavyGuardian: Separate and Guard Hot Items in Data Streams**
  Tong Yang*, Junzhi Gong*, Haowei Zhang, Lei Zou, Lei Shi, Xiaoming Li
  ACM SIGKDD 2018
  **: equal contributions*
- **HeavyKeeper: An Accurate Algorithm for Finding Top-k Elephant Flows**
  Junzhi Gong, Tong Yang, Haowei Zhang, Hao Li, Steve Uhlig, Shigang Chen, Lorna Uden, Xiaoming Li
  2018 USENIX Annual Technical Conference (ATC '18)
- **ABC: a Practicable Sketch Framework for Non-uniform Multisets**
  Junzhi Gong, Tong Yang, Yang Zhou, Dongsheng Yang, Shigang Chen, Bin Cui, Xiaoming Li
  IEEE International Conference on Big Data 2017 (IEEE BigData), Second workshop on Real-time and stream processing in Big Data
- **SSS: An Accurate and Fast Algorithm for Finding Top-k Hot Items in Data Streams**
  Junzhi Gong, Deyu Tian, Dongsheng Yang, Tong Yang, Tuo Dai, Bin Cui, Xiaoming Li
  IEEE International Conference on Big Data and Smart Computing 2018 (IEEE Bigcomp)
- **Difference Bloom Filter: a Probabilistic Structure for Multi-set Membership Query**
  Dongsheng Yang, Deyu Tian, Junzhi Gong, Siang Gao, Tong Yang, Xiaoming Li
  IEEE International Conference on Communications 2017 (ICC)

SELECTED TALKS

Scalable Distributed Massive MIMO Baseband Processing
NSDI conference talk — April 2023

Microscope: Detailed NFV Performance Diagnosis
SIGCOMM conference talk — August 2020

Microscope: Detailed NFV Performance Diagnosis
MNR group meeting in MSR — July 2020

HeavyGuardian: Separate and Guard Hot Items in Data Streams
SIGKDD conference talk — Aug 2018

HeavyKeeper: An Accurate Algorithm for Finding Tok-k Elephant Flows
ATC conference talk — July 2018

SKILLS

Programming Languages	C/C++, Java, Python, Bash
Tools	Latex, Git, Vim, SVN