

STATEMENT ON RESEARCH

Ye Yu

Research of computing technologies is facing emerging challenges in generating, transmitting, storing, and processing large-scale data. As a computer software systems researcher, I bring efficient algorithms and data structures to bear on practical problems by building real-world systems that tackle challenging issues of the **distributed networking systems and big data areas**. My research philosophy is to bring together theory and practice, with a special focus on algorithmic design/analysis and system implementation. The unique characteristic and strength of my research is the comprehensive study based on the expertise of a variety of fields including *cloud computing, management and processing of large-scale data (especially large-scale genomics data), high-data-rate network architectures and processing algorithms, and fault tolerance of distributed systems*.

I have published more than 15 papers in refereed journals and conference proceedings, most of which are in highly selective venues including *Bioinformatics* by Oxford Academic (Impact Factor 7.3), *IEEE/ACM Transactions on Networking* (ToN), *IEEE Transactions on Parallel and Distributed Systems* (TPDS), *IEEE ICNP*, *ACM CoNEXT*, *ACM SIGMETRICS*, and *ACM/IEEE IoTDI*. My research fields are closely related to the “Strategy for American Innovation” release by the White House in 2010 and 2015, for the highlighted topics including nationwide high-speed networks, big data, and secure cyberspace.

I participated in writing three grant proposals of my advisors: “CRII: NeTS: A Coordinate System based Network for Data Plane Scalability and High Throughput”, “CAREER: Othello Hashing and Its Applications to Scalable and Dynamic Network Forwarding and Functions”, and “BigData: IA: Collaborative Research: Finding Needles in a Haystack: A Cloud-based Query System to Empower Search in Large Scale Sequencing Data”. I contributed most of the preliminary data and helped to plan the future work for these proposals. **The first proposal was successfully funded by National Science Foundation**, and the other two are still pending.

1. EXISTING WORK

1.1. Othello Hashing: Ultra-fast and Memory-efficient Key-value Lookups for Large-scale Data. Systems that process big data (e.g., high-traffic networks and large-scale storage) prefer data structures and algorithms with small memory and fast processing speed. Our vision of computer system research by looking forward ten years is that efficient and fast algorithms will continue to play an essential role in system design, despite the improvement of hardware. For example, even though fast memory would become cheaper and larger, algorithms and portable primitives that demonstrates higher speed and less memory cost than existing solutions still provide tremendous benefits in resource saving and management flexibility. **The most significant achievement in my Ph.D. study is an ultra-fast and memory-efficient key-value lookup method called Othello Hashing, which can fit the requirements of the core algorithms of many large-scale systems and big data applications.** Othello achieves 5x to 20x speed in lookups of network addresses, data IDs, or DNA sequences, using significantly smaller memory (7% to 40%) compared to recent solutions (such as Cuckoo Hashing used in Google TensorFlow). My journey of exploring the applications of Othello Hashing starts from the high-traffic computer networking for cloud application and expanded to the big data and bioinformatics area. Using Othello hashing, combined with domain expertise in cloud, computer networks, big data, and bioinformatics, I developed the following systems and applications that resolve several major challenges in the area.

[Cloud/Networking Theme] Forwarding Information Base. A Forwarding Information Base is a data structure, typically a table, used to determine the proper forwarding actions for packets at the data plane of a forwarding device (e.g., switch or router). The increasing size of FIBs, due to the ever-growing number of connected network devices or entities, causes many problems including large memory cost (hence more expensive switches/routers), slow table lookups, and coarse-grained flow management. Meanwhile, the increasing line speed also requires a fast and small FIBs. The polymorphic property enables the separation of Othello’s query and control functionalities, which is a perfect match to the programmable networks such as Software Defined Networks. I built a fast and scalable FIB named Concise. Extensive evaluation results on three different platforms show that Concise outperforms other FIB designs. Compared to existing FIB designs for name switching, Concise supports much faster name lookup using significantly smaller memory (2x ~ 4x faster with only 10% ~ 30% memory size compared to state-of-art solutions).

[Cloud/Networking Theme] Cloud Load Balancer. In a cloud, the load balancer is a device that acts as a reverse proxy and distributes network or application traffic across a number of servers. Load balancers are widely used in cloud computing and mobile edge computing to increase the capacity and reliability of applications. A layer-4 load balancer servers the following two functions. 1) A lookup to find

the designated server if a packet belongs to an ongoing session. 2) If the packet is new or session-free, the load balancer selects one of the available servers according to the capacity of the servers. I built a software load balancer with Othello Hashing techniques named SDLB, which enables the load balancer to perform the two functions simultaneously using one query on the table. Evaluation results show that the system is faster by 4x to 10x and uses much less (< 50%) memory, than the current widely-used load balancer designs.

[Big Data/Bioinformatics Theme] Taxonomic Classification of Metagenomic Sequences. Metagenomic read classification is a critical step in the identification and quantification of microbial species sampled by high-throughput sequencing. Although many algorithms have been developed to date, they suffer significant memory and computational costs. Due to the growing popularity of metagenomic data in both basic science and clinical applications, as well as the increasing volume of data being generated, efficient and accurate algorithms are in high demand. We built a system to support efficient classification of taxonomic sequences using its k -mer¹ signatures. Othello shows 20x speed improvement and less than 5% memory cost compared to the state-of-art solutions.

[Big Data/Bioinformatics Theme] Transcript Query on RNA-seq datasets. Advances in the study of functional genomics produced a vast supply of RNA-seq datasets. However, how to quickly query and extract information from sequencing resources remains a challenging problem and has been the bottleneck for the broader dissemination of sequencing efforts. The challenge resides in both the sheer volume of the data and its nature of unstructured representation. We carefully designed the algorithms so that it enables sequence search on a vastly compressed data domain. Using the Othello Hashing techniques, we built the first indexing structure capable of answering the *sequence coverage* query by extracting its k -mer coverage across large sequencing datasets. Such information can be used to identify mutations, distinguish transcript isoforms and predict a sequence’s expression level. For sequence containment query, which is also supported by traditional approaches, it achieves a two orders of magnitude improvement over the state-of-art while retaining the same accuracy.

1.2. Data Center and Cloud Architecture. The increasing need for cloud and big data applications requires data center networks to be scalable and bandwidth-rich. Network topology and the corresponding routing protocol are determinant factors of application performance in a data center network. Recent work has been investigating architectures and protocols of large data centers to achieve three important design goals, namely (1) high throughput, (2) routing and forwarding scalability, and (3) flexibility for incremental growth. However, existing data center network architectures focus on one or two of these three properties and pay little attention to the others. We designed a novel flexible data center network architecture, SpaceShuffle (S2) , which applies greedy routing on multiple ring spaces to achieve high-throughput, scalability, and flexibility. Extensive experimental studies show that S2 provides high bisectional bandwidth and throughput, near-optimal routing path lengths, extremely small forwarding state, fairness among concurrent data flows, and resiliency to network failures. In addition, existing data centers do not provide specific fault-tolerance mechanisms to recover the network from failures and to protect network performance from downgrading. To address this problem, we extend the S2 design and achieve fault tolerance in a self-fixing manner with the FTDC data center design. Upon failures, the servers automatically explore valid paths to deliver packets to the destination by exchanging control messages among servers. We are collaborating with a computer architect Prof. Jishen Zhao at UC San Diego on using S2 to build a memory interconnection for in-memory big data processing such as Spark. A paper is in submission to ISCA 2018.

1.3. SDN Management. Software-defined networking (SDN) is an emerging paradigm that separates the network’s control logic from the underlying routers and switches to a logically centralized network controller. Network traffic monitoring supports fundamental network management tasks. However, monitoring tasks introduce a non-trivial overhead to network devices such as switches. We proposed a Distributed and Collaborative Monitoring system, named DCM. DCM is a memory-efficient solution that performs per-flow monitoring in SDN. It allows switches to fulfill the monitoring tasks and balance measurement load collaboratively. DCM uses novel two-stage Bloom filters to represent monitoring rules using small memory space. It utilizes the centralized SDN control to install, update, and reconstruct the two-stage Bloom filters in the switch data plane. Experiments on real network topologies show that DCM achieves highest measurement accuracy among existing solutions given the same memory budget of switches.

1.4. Industry Collaboration. I have participated in multiple collaborative projects with the industry. In particular, at the Ericsson Research, I developed a rollback-recovery mechanism for the SDN system to enhance its robustness. I also developed a new framework for load-balancing and bandwidth allocation

¹ k -mer refers to a string of length k that represents k consecutive nucleotide bases in sequencing results. Suppose the genome is a book, then k -mer is the corresponding term for “word” in this book.

for Facebook’s internal inter-data center backbone SDN. These collaborations equipped me with first-hand experience and knowledge about large-scale cloud systems in the industrial production environment, which further strengthens my ability in system research, design and implementation. I also treasure the connection with the industry and expect further collaboration opportunities.

2. FUTURE RESEARCH PLAN

My long-term goals are to explore directions in algorithms and their applications in networking systems, big data systems, and cloud systems for applications including bioinformatics data processing and Internet of Things. During the journey of exploring the Othello Hashing algorithm, I have located several important but challenging problems to be solved. Currently, I have already embarked on ambitious multi-disciplinary research agendas along these directions, where many promising research results are expected to come out soon.

Storage and Query of large-scale sequencing data. Existing sequencing data available in both public and private domains are growing exponentially at a rate far beyond the growth of computation power. How to organize and maintain a searchable library of such ever-growing large-scale set of sequencing data remains to be challenging. With Othello, we reduced the storage overhead and achieved sequence *coverage* query. Can we go even further by bringing more data with biological information meaning to the Othello Hashing techniques and achieve *expression* query? Instead of *k*-mer signatures, is it possible to store and query the transcriptome as a whole efficiently?

Fast and scalable programmable data planes. While we conquered the scalable problems for FIB explosion by Othello Hashing, there are still essential but challenging issues remain opening for programmable data planes. How does the data plane deal with multihoming where multiple entities in a network may share the same name? How shall we design the communication protocol and update mechanism that better supports the concurrent update and the concurrency requirements among multiple data plane devices?

Bioinformatics data analytics in the cloud. The unprecedented resolution enabled by high-throughput sequencing (HTS), coupled with its decreasing cost, has fueled an explosion of primary data. With the recent development of cloud computing services such as Amazon EC2, Azure, and Google Cloud, more and more research institutions have been moving the sequencing data or the processing procedure to the cloud. However, although researchers are moving to cloud computing services, the current bioinformatics tools are not specialized for the cloud, and hence the cloud infrastructure can only be utilized as a “big computer with many CPUs”. In my point of view, there is still a major obstruction of the improvement of the performance and availability of bioinformatics applications. Such obstruction sources from the lack of system-level design that takes advantage of cloud computing. More importantly, less effort has been made into the investigation of how the bioinformatics researches may benefit from the distributed nature of cloud computing infrastructures. For example, how does online sequence query in a distributed cloud system benefit the collaboration among researchers that are geologically separated apart? Meanwhile, the *privicay and anonymity* concerns that hinder processing sequencing data from being stored or processed in the cloud remains to be addressed. I will further explore the abundant merits of algorithms and tools in the networking area, and try to fully utilize the cloud for bioinformatics applications by designing new tools that feature domain information of both cloud computing and bioinformatics.

In addition, I am planning to continue my research in Cloud and Data Center Networks. The objectives of include but not limited to: i) new applications that adopt fast and scalable processing of network packets in various types of enterprise networks, edge computing, data centers and cloud systems. ii) new mechanisms and algorithms for file systems that improve the reliability, availability, and mobility for the encoding, storage, and query of data. iii) new tools and mechanisms to enhance the robustness and reliability of networks such as wide-area Software Defined Networks.

We are entering a great time to conduct multi-disciplinary research. This is especially true for networking, big data, and bioinformatics. The digital revolution has drastically changed many major aspects of our life, and the new progresses in computer science will shape the world even more. I believe future for networking research is bright, as we are motivated to achieve more scalable, flexible and reliable network infrastructures. Meanwhile, as the sequencing techniques and bioinformatics algorithmic approaches are ready, with the recent drastic improvement in the computing infrastructure, I expect a breakthrough of genomics applications with new algorithms and system designs. These revolutionary breakthroughs may change the lifestyle of the human race again in the foreseeable future. I am excited to see myself as a multi-disciplinary researcher that may make my contribution in addressing the emerging challenges.

E-mail address: ye.yu@uky.edu