Research Statement

Junchen Jiang

Today's Internet has become an "eyeball economy" dominated by applications such as video streaming (e.g., Netflix counts for 37% of Internet traffic) and VoIP (e.g., Skype users spend over 2 billion minutes talking to each other every day). With most application providers relying on user engagement to generate revenues, maintaining *high user-perceived QoE* (Quality of Experience) has become crucial to ensure high user engagement. For instance, recent research shows even one short buffering interruption leads to 39% less time spent watching videos and causes significant revenue losses for ad-based video sites. Despite increasing expectations for high QoE, existing approaches have limitations to achieve the QoE needed by today's applications. They either require costly re-architecting of the network *core*, or use suboptimal *endpoint*-based protocols to react to the dynamic Internet performance based on limited knowledge of the network.

My dissertation explores a new approach, inspired by the recent success of data-driven approaches in many fields of computing. I demonstrate that *data-driven* techniques can improve Internet QoE by utilizing a centralized real-time view of performance *across millions of endpoints* (clients), rather than per-endpoint information. My work focuses on two fundamental challenges that are unique to applying data-driven approaches in networking: the need for *expressive* models to capture complex factors affecting QoE, and the need for *scalable* platforms to make real-time decisions with fresh data from geo-distributed clients. I address these challenges in practice by integrating several domain-specific insights in networked applications with machine learning algorithms and systems, and achieve better QoE than using off-the-shelf machine learning solutions. The unifying theme underlying these domain-specific insights is that there exist some *structure* (e.g., performance bottlenecks) that identifies network sessions with similar QoE-determining factors, and that such structure tends to be *persistent* on timescales of tens of minutes. During my Ph.D., I developed end-to-end systems that substantially improved QoE for video streaming and VoIP. Such quality improvements can lead to substantial increase in user engagement, and eventually more revenue for application providers. Two of my projects, CFA [1] and VIA [3], have been used in industry by Conviva and Skype, companies that specialize in QoE optimization for video streaming and VoIP, respectively.

Thesis Work

My thesis work shows QoE can benefit from applying data-driven approaches in networking – QoE can be improved by learning optimal client configurations in real time from a global view of performance across millions of endpoints. This section presents three systems, each addressing a key algorithmic or architectural challenge in realizing this idea.

CFA: Improving QoE via Expressive Prediction Models. Delivering high QoE is crucial to the success of today's subscription-/ad-based business models for Internet video. There is substantial room for improving video QoE by dynamically selecting the optimal CDN (Content Delivery Network) and bitrate based on a real-time global view of network conditions [8, 5]. The key to realizing this promise is a *prediction oracle* that can accurately predict the quality of a video client if it uses a certain CDN and bitrate. The challenge is that this prediction system must be (a) expressive enough to capture complex relations between video quality and observed session features, and (b) capable of updating quality predictions in near real time. We used several off-the-shelf machine learning techniques, such as random forests and SVM, but found they did not produce expected QoE improvements, because the long-term historical data is too coarse grained for these algorithms to capture the dynamics of video quality, while the short-term historical data is not sufficient for the algorithms to learn complex relations between video quality and observed session features. My solution, CFA [1], leverages the domain-specific insight of *persistent critical features*; each video session has a small set of critical features that ultimately determines its video quality, and these critical features change much more slowly than video quality [5]. This insight enables us to learn complex prediction models from long-term historical data (thus expressing complex relations between video quality and session features), and update the models by short-term historical data in near real time (thus capturing quality fluctuation as well). A real-world pilot deployment shows that CFA leads to non-trivial improvements in video quality, e.g., 32% less buffering time than industry-standard algorithms. Our conversation with domain experts confirmed that these improvements are significant for content providers and can potentially translate into substantial benefits in revenues. An end-to-end implementation of CFA has been deployed and used by Conviva, a company that offers video quality optimization services to many premium content providers.

The insight of persistent critical features turns out to be more general; e.g., I have also applied the same insight to accurate prediction of TCP throughput [7], which leads to 11% higher video bitrate than state-of-the-art adaptive bitrate players (e.g., Netflix players) with no extra buffering.

Pytheas: Improving QoE via Exploration and Exploitation at Scale. While formulating data-driven QoE optimization as a prediction problem has shown promising QoE improvement (e.g., CFA), it is necessarily incomplete, as it suffers from many known biases such as incomplete visibility, and cannot respond to sudden changes such as flash crowds. Drawing a parallel from machine learning (e.g., ad recommendation), I argued that data-driven QoE optimization should instead be cast as a *real-time exploration and exploitation* process. Measurement collection (exploration) could be informed by decision making (exploitation) to explore the decisions with less data, thus addressing the shortcomings of the prediction-based formulation. This new formulation is complementary to CFA, since CFA's prediction model can be reused to capture similarities among clients. While many exploration-and-exploitation algorithms are available, enabling these algorithms in networking introduces an architectural challenge: we need a *scalable control platform* to update decisions with fresh global data from clients, despite data coming from a variety of geo-distributed front-end clusters, each with only a partial view of the clients. A baseline approach is to run control logic in a single back-end cluster with global data from front-end clusters, but this approach leads to non-trivial staleness of the global data and suboptimal decisions. I take an alternative approach; the control logic is run by geo-distributed front-end clusters which have fresh data, rather than the back-end cluster. The intuition is that the clients that exhibit similar QoE behavior will have similar network-level features (e.g., same IP prefix), and thus their fresh data will likely be collected by the same front-end cluster. Inspired by this insight, Pytheas [2] uses the concept of *group-based exploration and exploitation*, which decomposes the global exploration and exploitation process into subprocesses, each controlling a group of similar clients and running in the front-end cluster that has these clients' fresh data. Using an end-to-end implementation in CloudLab, I show that compared to a state-of-the-art prediction-based system, Pytheas improves the average video QoE by up to 31% and the 90th percentile QoE by 78%. Pytheas is an open source project that enables application providers to deploy the proposed architecture at scale within their existing infrastructure.

VIA: Improving QoE in the Face of Large Decision Spaces. In the first large-scale study on Internet telephony quality, I found that a substantial fraction of Skype calls suffer from poor network performance, and that there is much room for improving Skype quality by routing each call through the optimal relay clusters in Microsoft's cloud. However, identifying a close-to-optimal relay in practice is very challenging, due to the sheer number of possible relay paths (in hundreds) and their dynamic performance (which could change on timescales of minutes). Neither prediction-based methods (e.g., CFA) nor those based on exploration and exploitation (e.g., Pytheas) suffice to handle such a large decision space. The key insight to address this challenge is that, for each pair of caller ISP and callee ISP, there is a *small and stable subset* of relays that almost always contains the best relay. This insight has two implications: (1) because this subset of relays is stable, it can be learned from history; and (2) because this subset has only a few relays (less than five), it can be explored efficiently even with limited data. Inspired by these intuitions, I developed VIA [3], a relay selection system that achieved close-to-optimal quality using the concept of *guided exploration*. The idea is to learn a small set of promising relays for each ISP pair based on long-term (e.g., daily) historical data, and explore these relays using most calls in real time. Trace-driven analysis and a small-scale deployment shows that VIA cuts the incidence of poor network conditions for calls by 45% (and for some countries and ISPs by over 80%) compared to today's Skype quality. VIA has been used in Microsoft internal deployment with real Skype users, and is in the process of being fully deployed. VIA was also used by Microsoft to identify the countries and ISPs where Skype quality can benefit the most if relaying services are deployed.

Future Research

My vision for the future is that there will be tremendous opportunities for data-driven techniques in networking driven by both "use pulls" (high QoE/performance requirements and the increasing complexity of networked systems) and "technology pushes" (availability of big data in networking and new data analytics capability). My future research will explore the confluence of data-driven paradigms and networking along three dimensions.

Rethinking Classic Problems in Networking. The data-driven paradigm encourages rethinking many classic problems in networking as well as enabling new services.

- *Rethinking congestion adaptation:* As the Internet grows more complex, it has become increasingly difficult for transport protocols (e.g., TCP) and applications (e.g., Facebook Live) to adapt in real time to changes in network conditions and resource availability. I believe data-driven techniques offer a promising new approach. Consider TCP; I envision a new service that aggregates real-time performance of similar TCP connections and *predicts* the largest window size for which a TCP connection will not experience congestion losses. This service can potentially lead to better TCP performance than many state-of-the-art techniques, as it needs little active probing and can train the logic in near real time, rather than offline. An early promise of this approach was shown in my prior work of CS2P [7] which can accurately predict HTTP throughput using information of similar HTTP sessions, and CS2P could potentially be extended to inferring optimal window sizes from the HTTP throughput prediction.
- *Real-time Internet traffic map:* Many Internet services can benefit from a traffic map service that can predict performance (e.g., available bandwidth) of any network path, but prior efforts towards such a service suffer from limited

coverage, high probing overhead, or significant data staleness. Fortunately, passive measurements of applications like video streaming offer a new enabler for such a service, as they provide real-time measurements of network state from millions of vantage points without incurring any probing overhead! My prior work [3, 6] shows it is possible to predict coarse-grained performance metrics (e.g., RTT between ISPs) by mapping video quality to simple network models. Extending these maps to predicting fine-grained (e.g., link-level) metrics provide an interesting next step.

Custom data analytics stack for networking. While data-driven techniques can be applied to many problems, there are several *common* challenges across these applications. First, while it is tempting to use general-purpose techniques to analyze networking data, one must take into account *network-specific knowledge* to avoid undesirable outcomes; e.g., to optimize overall performance, ESPN may use small ISPs to use suboptimal CDNs and let Comcast users use optimal CDNs, causing a reverse network neutrality violation in which content providers discriminate against ISPs! Second, one also needs to leverage *application-specific resilience*; e.g., to probe both optimal and suboptimal servers without affecting quality, a video player can use the optimal server to fetch most content, but use suboptimal servers when the buffer is full. While data-driven approaches will have more benefits if these network-/application-specific opportunities can be utilized in a systematic way, unfortunately, there are few formal approaches to integrating them with existing "black-box" algorithms. To best utilize various network-/application-specific opportunities, I envision a *custom data analytics stack for networking* built on top of existing analytics stacks (e.g., Spark), which offers abstractions to express network/application concepts (e.g., sessions) and requirements (e.g., data staleness).

Experience-Oriented Network Architecture. So far my work has been primarily within the scope of a single service provider, looking forward I argue that there is greater room for improving application QoE by bringing *more* service providers into the loop. These providers can be ISPs, CDNs, and cloud services, whose revenue is also driven, albeit indirectly, by high QoE. Unfortunately, in today's federated Internet structure, many problems arise because each service provider has little visibility into either the QoE of its sessions or the decisions made by others. For instance, ESPN clients use HTTP-based bitrate-adaptation video players, which can select both CDN (e.g., Akamai and Limelight) and bitrate, and in many cases these players perform poorly because ESPN cannot know whether the bottleneck link is at edge ISPs (e.g., Comcast) or CDN servers (e.g., Akamai). Now, if Comcast shares with ESPN additional information that attributes bottlenecks to Comcast rather than Akamai, the players can react to ISP congestion by lowering the bitrate, rather than blindly switching between Akamai and Limelight. Building on this intuition and my early work [4], I argue that the fundamental problem lies in the limited interfaces between service providers, and that a principled approach to re-architecting these interfaces with explicit QoE optimization in mind is needed. I envision an *Experience-Oriented Network Architecture* (*EONA*), where service providers share QoE information and key configuration changes in order to optimize user-perceived QoE in a coordinated way, while still keeping the federated structure. EONA can be realized without changing the data/control plane protocols; instead, it adds new interfaces between the "controllers" of service providers (e.g., an SDN controller, a cloud orchestrator, or a global video control plane [8]). Despite EONA's potential, there are many open-ended questions: How to incentivize EONA to attract service providers? What is the minimal set of information that has to be shared? How to preserve privacy when sharing data between providers?

The past years have witnessed the coming of age of data-driven paradigms in various aspects of computing (partly) empowered by advances in system research (cloud computing, MapReduce, etc). Now, I believe the benefits can flow the opposite direction; I am excited to explore the possibilities of improving distributed systems by data-driven paradigms.

Ask not what networking can do for data; Ask what data can do for networking!

References

[1] **J. Jiang**, V. Sekar, H. Milner, D. Shepherd, I. Stoica, H. Zhang. "CFA: A Practical Prediction System for Video Quality Optimization" *NSDI 2016*

[2] **J. Jiang**, S. Sun, V. Sekar, H. Zhang. "Pytheas: Enabling Data-Driven Quality of Experience Optimization Using Group-Based Exploration-Exploitation" *NSDI 2017*

[3] **J. Jiang**, R. Das, G. Anathanarayanan, P. A. Chou, V. Padmanabhan, V. Sekar, E. Dominique, M. Goliszewski, D. Kukoleca, R. Vafin, H. Zhang. "Via: Improving Internet Telephony Call Quality Using Predictive Relay Selection" *SIGCOMM 2016*

[4] **J. Jiang**, V. Sekar, X. Liu, I. Stoica, H. Zhang. "EONA: Experience-Oriented Network Architecture" *HotNets 2014*

[5] **J. Jiang**, V. Sekar, I. Stoica, H. Zhang. "Shedding Light on the Structure of Internet Video Quality Problems in the Wild" *CoNEXT 2013*

[6] Y. Sun, **J. Jiang**, V. Sekar, H Zhang, F. Lin, N. Wang. "Using Video-Based Measurements to Generate a Real-Time Network Traffic Map" *HotNets 2014*

[7] Y. Sun, X. Yin, **J. Jiang**, V. Sekar, F. Lin, N. Wang, T. Liu, B. Sinopoli. "CS2P: Improving Video Bitrate Selection and Adaptation with Data-Driven Throughput Prediction" *SIGCOMM 2016*

[8] X. Liu, F. Dobrian, H. Milner, **J. Jiang**, V. Sekar, I. Stoica, H. Zhang. "A Case for a Coordinated Internet-Scale Video Control Plane" *SIGCOMM 2012*