RESEARCH STATEMENT

Abhishek Sinha

1 Introduction

In the twenty-first century of the all-pervasive Internet, when the annual global IP traffic is expected to exceed the threshold of a Zettabyte (10^{21} bytes), the need for efficient networking policies, protocols, and algorithms cannot be overemphasized. The ongoing global initiative *Internet of Things* (IoT) proposes interconnecting more than 50 billion objects by the year 2020, thus creating an urgent need for highly scalable, robust, and efficient network control algorithms. My research is primarily based on a principled and rigorous approach, seeking to understand and solve diverse problems arising in the broad area of Networking. In particular, I use tools from Queueing Theory, Control Theory, Optimization Theory, Information Theory, and Discrete Mathematics to address fundamental problems such as throughput-optimal routing and scheduling, optimal load balancing, and delay-optimal resource allocation for emerging wireless paradigms such as 5G. During my industrial stints at Qualcomm Research, Bell Labs, and Microsoft Research, I have delightfully experienced how concrete theoretical ideas can lead to elegant and practical solutions to pressing industrial problems. Some of the central results that I obtained in this pursuit are described below.

1.1 Throughput-Optimal Policy for Generalized Network Flow Problems

One of the most fundamental problems in Computer Networking is to *efficiently* distribute messages from the source(s) to the destination(s) over a communication network. This problem has been studied extensively over the period of last five decades, resulting in a number of routing and scheduling algorithms. In particular, in their seminal paper [1], Tassiulas et al. introduced the *Backpressure* algorithm, which is a throughput-optimal dynamic routing and scheduling policy for *unicast* traffic - where every packet has a *single* destination. Although Backpressure is a solution to the unicast problem, a significant fraction of traffic carried by the Internet and military applications today are of *multicast* and *broadcast* in nature - where copies of a packet are to be distributed among *multiple* destinations. Moreover, there is a growing trend of *anycast* traffic, especially in CDN applications - where the destination node for a packet has to be optimally chosen from a given set of possible destination nodes. Unfortunately, little progress has been made towards a universal algorithm which simultaneously addresses all of the above problems, collectively known as the *Generalized Network Flow* problem. In fact, to the best of our knowledge, there is no known efficient throughput-optimal policy even for the Broadcast and the Multicast problem. The main technical difficulty stems from the fact that, unlike unicast, there is no per-node flow conservation due to packet duplications. Moreover, even for the unicast case, the Backpressure policy is not entirely satisfactory as it results in large delays due to an indefinite circulation of packets inside the network [2].
In a recent series of papers [3], [4], [5], [6], we have designed an efficient throughput-optimal dynamic control policy for the generalized network flow problem. These papers introduce a surprisingly general algorithmic paradigm which can handle a simultaneous mix of unicast, broadcast, multicast and anycast flows. Our generalized policy yields an alternative delay-efficient solution to the unicast problem, improving upon the Backpressure policy and its enhanced variants. The proposed policy is particularly suitable for implementation in the emerging Software Defined Networking (SDN) framework. The key algorithmic insight in our framework is to relax the physical *precedence* constraints and consider a suitably defined *virtual* queueing system. The control actions are taken based on the virtual queue lengths, as opposed to the physical queue lengths, common in the rest of the literature. In this work, we make the following major contributions:

- We design a generic algorithmic framework which gives a unified solution to *all* flow problems, including unicast, broadcast, multicast and anycast.
- For the unicast problem, our framework yields a dynamic algorithm which overcomes many limitations of the state-of-the-art throughput-optimal policies, and drastically reduces end-to-end latency.
- Our analysis effectively combines Lyapunov drift theory with adversarial queueing theory, which is of independent theoretical interest.

1.2 Throughput-Optimal Broadcast Policy

Before delving into the Generalized Network Flow problem described above, I focused on the particular case of the pure *Broadcasting* problem. It turns out that by exploiting the special structure of this problem, we can design throughput-optimal broadcasting policies which are very different from the one derived from the previous generalized policy. Moreover, this new policy also offers some distinct computational advantages. We make the following contributions in this pursuit:

- In our papers [7], [8], we showed that it is possible to achieve the broadcast capacity in a directed acyclic wireless network while delivering the packets in the order of their arrival (*in-order*). This result leads to a low-complexity, distributed, and efficient broadcast policy for wireless DAGs. Interestingly, this policy does not require keeping track of expensive network-wide structures, such as spanning trees. We also proved a surprising converse result that, in general, the broadcast capacity cannot be achieved by any *in-order* policy if the network is not a DAG.

- During my visit to the Yale University, I showed in [9] that the above policy can be extended to time-varying wireless networks. Moreover, the broadcast capacity is also shown to be efficiently computable under some standard interference models. This paper received the **Best Paper award** in the conference ACM MobiHoc 2016, held in Paderborn, Germany.
- In a follow-up work [10], [11], I considered the problem of broadcasting in arbitrary networks (not necessarily DAGs). I designed a new dynamic policy which can be viewed as "Backpressure on Sets". Moreover, I showed that the limitation imposed by the previous converse result may be overcome by using *multi-class in-order* policies.
- All of the above work on broadcasting assumes wireless networks with *point-to-point* links (*e.g.,* 5G mmWave). In our recent paper [12], we tackled the optimal broadcasting problem for networks with wireless broadcast advantage (i.e., networks with *point-to-multipoint* links). This paper makes use of the general methodology introduced in [4].

1.3 Optimal Multi-User Scheduling

I studied the multi-user scheduling problem in wireless networks from the dual perspective of minimizing the average delay and maximizing the aggregate throughput.

- Minimizing latency is chosen to be one of the most important design criteria in the emerging 5G standard (*e.g.,* for URLLC traffic). To explore the topic of low-latency scheduling in this context, we studied an optimal scheduling problem over a multi-user wireless erasure channel. The objective was to minimize the *Age of Information*, which measures the freshness of information available to the users. In this pursuit, we derived a converse and a 2-optimal scheduling policy [13], [14]. The paper [13] received the **Best Paper Award** in the conference IEEE INFOCOM 2018, held in Honolulu, HI, USA. At Qualcomm Research, we are currently studying the efficacy of the proposed policy in the context of minimizing scheduling delay due to hybrid beamforming constraints in 5G mmWave networks.
- The majority of the backhaul Internet traffic today is transported via Passive Optical Networks (PON). For a large number of wireless end-users, with time-varying channels and stochastic demands, it is imperative to design low-complexity, yet efficient scheduling algorithms for accessing the limited-capacity backhaul PON. During my internship at Bell Labs, I designed an optimal pseudo-polynomial time Dynamic Program, and a polynomial time LP rounding-based 2-approximation algorithm for this scheduling problem. A US patent application has been filed for this work [15].

1.4 Load Balancing in Anycast-based CDNs

Content Distribution Networks or CDNs are the cornerstones of modern cloud-based services. Popular large-scale applications, such as e-mails, social networking, and video streaming make extensive use of CDNs for efficient operation. CDNs operate by geo-replicating the content across multiple servers for faster access across the globe. To load-balance among multiple servers, *anycast* addressing is used to route the users to their "closest" server. However, due to its inherently greedy routing policy, service hotspots often arise, overloading some servers and disrupting the overall system performance. To solve this interesting problem, while I was interning at Microsoft, I started out by developing a realistic yet analytically tractable dynamical model of a commercial CDN system - `Microsoft Azure`®. Utilizing tools from convex optimization and nonlinear system theory, we designed a *Kelly-type* dual decomposition algorithm for large-scale CDNs. An interesting feature of our proposed policy is that it effectively exploits the underlying anycast architecture to coordinate actions among multiple nodes in a decentralized fashion. A preliminary version of the results appeared in [16], and a comprehensive account of the work appeared in [17].

1.5 Optimal Network Deployment Policy

During my Master's studies at the Indian Institute of Science, I worked on the problem of *impromptu* deployment of ad hoc networks in uncertain environments, e.g., battlefields or places hit by a natural disaster. The objective of the problem was to design an optimal deployment policy for a limited number of sensor nodes such that, (1) coverage and connectivity requirements are met, and (2) end-to-end latency is minimized. We formulated the problem as a countable state total cost Markov decision process and obtained the following results [18]:

- We showed that a simple One-Step-Look ahead (OSLA) threshold policy is optimal.

- Exploiting the regenerative structure of the problem, we developed an algorithm which is substantially faster than the usual Value Iteration, and yields the *exact* optimal deployment policy in *finite* number of steps. This result is surprising, as the general total cost MDP on a countable state space is known to be notoriously hard to solve.

In addition to the above, along with my colleagues at IISc, I worked on the problem of optimal deployment of full-duplex relays along a line with power-constrained nodes. In our papers [19], [20], we designed optimal node deployment and power allocation policy using Information-theoretically achievable rate characterizations.

1.6 Theory of Evolutionary Optimization Algorithms

In my undergraduate thesis, I worked on bio-inspired meta-heuristic algorithms, which are surprisingly effective in practice in solving hard optimization problems. However, most of these algorithms (e.g., *Genetic Algorithm* (GA), *Particle Swarm Optimization* (PSO), *Biogeography based Optimization* (BBO) etc.) are heuristic in nature with little or no performance guarantees. In our paper [21], we showed that the migration dynamics of BBO is globally asymptotically stable. Our proof uses techniques from non-negative matrix theory.

2 A Vision for Future Research

It is widely accepted in the networking community that even after tremendous theoretical and practical advances over the last fifty years, the subject of optimal control and optimal resource allocations in large-scale networked systems is still in infancy. Keeping this in view, I would like to address the following problems immediately:

2.1 Efficient Policies for 5G Networks

The emerging paradigm of 5G technology brings the excitement of high-throughput connectivity at a massive scale, along with a new set of technical challenges, especially in the mmWave (6-100 GHz) regime. As discussed earlier, a key performance indicator for the 5G networks is user-perceived latency. Although various throughput-optimal schemes are known in different networking settings, the necessary practical issue of delay-optimality is mostly open due to its technical difficulty. My recent work on generalized network flow problem takes an entirely new approach to designing and analyzing throughput-optimal dynamic policies with low delay. Moreover, we obtained some interesting preliminary results in [13] for the problem of maximizing freshness of updates. More recently, at Qualcomm Research, I am working on designing practical scheduling algorithms for 5G mmWave networks taking into account the hybrid Beamforming constraints. Along with my colleagues, I have recently filed two US patents on beam management and Uplink Control [22], [23]. In future, I plan to explore this exciting area at the intersection of technical standards, industrial practices, and academic research. Creating IP and contributing to the standards will also be on my research agenda.

2.2 Information and Control in Cyber Physical Systems

The connection between information and control in the context of networking problems is relatively unexplored. In practice, a Cyber-Physical System needs to take control actions based on noisy, delayed, and incomplete information on the system state. A fundamental problem in this context is how to quantitatively interpret the “value of information” in connection with the Quality of Service (QoS) of a large-scale interconnected and interdependent system. Our recent work on Age of Information takes a first step in answering this question. One of my future research plans is to incorporate recent advances from diverse disciplines, such as Information theory, Network Control, and Decision Theory to explore the key questions in this area.

2.3 Learning Problems in Networks

Due to the complexity of large-scale networks, such as the Internet, designing fast, adaptive, and robust control policies for such systems have always remained a challenge. The last decade has witnessed crucial breakthroughs in the field of Machine Learning and Data Science. However, most of these techniques have not made their way into the mainstream networking research yet due to their complexity. In future, I plan to work in the emerging interface between networks and machine learning, and design effective control techniques inspired by the advances in Machine Learning.

References

[1] L. Tassiulas and A. Ephremides, “Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks,” *Automatic Control, IEEE Transactions on*, vol. 37, no. 12, pp. 1936–1948, 1992.

[2] L. Ying, S. Shakkottai, A. Reddy, and S. Liu, "On combining shortest-path and back-pressure routing over multihop wireless networks," *IEEE/ACM Transactions on Networking (TON)*, vol. 19, no. 3, 2011.

[3] **A. Sinha** and E. Modiano, "Optimal control for generalized network-flow problems," in *IEEE Conference on Computer Communications (INFOCOM), 2017*, Available online https://arxiv.org/abs/1611.08641

[4] **A.Sinha** and E. Modiano, "Optimal control for generalized network flow problems," minor revision submitted to *IEEE/ACM Transactions on Networking (TON).*

[5] **A. Sinha** and E. Modiano, "Network utility maximization with heterogeneous traffic flows," in *Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks (WiOpt), 2018 16th International Symposium on* (to appear), Shanghai, China.

[6] J. Zhang, **A.Sinha**, J. Llorca, A. Tulino, and E. Modiano, "Optimal control of distributed computing networks with mixed-cast traffic flows," to appear in *INFOCOM 2018, Honolulu, HI.*

[7] **A. Sinha**, G. Paschos, C.-P. Li, and E. Modiano, "Throughput-optimal broadcast on directed acyclic graphs," in *IEEE Conference on Computer Communications (INFOCOM), 2015.* IEEE, 2015, pp. 1248–1256.

[8] **A. Sinha**, G. Paschos, C. P. Li, and E. Modiano, "Throughput-optimal multihop broadcast on directed acyclic wireless networks," *IEEE/ACM Transactions on Networking*, vol. PP, no. 99, pp. 1–15, 2016.

[9] **A. Sinha**, L. Tassiulas, and E. Modiano, "Throughput-optimal broadcast in wireless networks with dynamic topology," in *Proceedings of the 17th ACM International Symposium on Mobile Ad Hoc Networking and Computing*, ser. MobiHoc '16, 2016, pp. 21–30, (**Best Paper Award**).

[10] **A Sinha**, G. Paschos, and E. Modiano, "Throughput-optimal multi-hop broadcast algorithms," in *Proceedings of the 17th ACM International Symposium on Mobile Ad Hoc Networking and Computing*, ser. MobiHoc '16, 2016.

[11] **A. Sinha**, G. Paschos, and E. Modiano, "Throughput-optimal multi-hop broadcast algorithms," *IEEE/ACM Transactions on Networking*, vol. 25, no. 5, pp. 3088–3101, Oct 2017.

[12] **A Sinha** and E. Modiano, "Throughput-optimal broadcast in wireless networks with point-to-multipoint transmissions," in *Proceedings of the 18th ACM International Symposium on Mobile Ad Hoc Networking and Computing*, ser. MobiHoc '17, 2017.

[13] I. Kadota, **A.Sinha**, and E. Modiano, "Optimizing age of information in wireless networks with throughput constraints," to appear in *INFOCOM 2018, Honolulu, HI*, (**Best Paper Award**).

[14] I. Kadota, **A.Sinha**, E. Uysal, R. Singh, and E. Modiano, "Scheduling policies for minimizing age of information in broadcast wireless networks," submitted to *IEEE/ACM Transactions on Networking (TON)*, Available online http://bit.ly/2miZy4l

[15] M. Andrews, P. Ananth, and **A. Sinha**, "Integrated scheduler for scheduling with x-haul capacity constraints," Patent, invention submission 81991 at Nokia Bell Labs. US Patent application serial no: 15/630,367.

[16] **A. Sinha**, P. Mani, J. Liu, A. Flavel, and D. A. Maltz, "Distributed load management in anycast-based CDNs," in *2015 53rd Annual Allerton Conference on Communication, Control, and Computing (Allerton).*

[17] **A. Sinha**, P. Mani, J. Liu, A. Flavel, and D. Maltz, "Distributed load management algorithms in anycast-based cdns," *Computer Networks*, vol. 115, pp. 1–15, 2017.

[18] **A. Sinha**, A. Chattopadhyay, K. Naveen, P. Mondal, M. Coupechoux, and A. Kumar, "Optimal sequential wireless relay placement on a random lattice path," *Ad Hoc Networks*, vol. 21, pp. 1–17, 2014.

[19] A. Chattopadhyay, **A. Sinha**, M. Coupechoux, and A. Kumar, "Deploy-as-you-go wireless relay placement: An optimal sequential decision approach using the multi-relay channel model," *IEEE Transactions on Mobile Computing*, vol. PP, no. 99, pp. 1–1, 2016.

[20] A. Chattopadhyay, **A. Sinha**, M. Coupechoux, and A. Kumar, "Optimal capacity relay node placement in a multi-hop network on a line," in *Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks (WiOpt), 2012 10th International Symposium on.* IEEE, 2012, pp. 452–459.

[21] **A. Sinha**, S. Das, and B. K. Panigrahi, "A linear state-space analysis of the migration model in an island biogeography system," *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans*, vol. 41, no. 2, pp. 331–337, 2011.

[22] M. Wilson, S. Akkarakaran, T. Luo, Y. Zhou, X. Wang, W. Nam, and **A. Sinha**, "TCI Indication Procedure for UL Control," March 2018, Qualcomm IDF #182761IDF, US Patent Filed.

[23] X. Wang, T. Luo, S. Akkarakaran, M. Wilson, Y. Zhou, W. Nam, **A. Sinha**, and J. Montojo, "PUCCH Reliability Enhancements," March 2018, Qualcomm IDF #182660IDF, US Patent Filed.