My research interests are in the area of **systems and networking.** I am currently a **research fellow** at **Microsoft Research India** where I work with **Dr. Venkat Padmanabhan** and **Dr. Ranjita Bhagwan**. During my undergraduate studies, I have undertaken research projects at Tsinghua University and Technical University of Munich. These research experiences have bolstered my interests in this field. I am applying to the Ph.D. programme at the Massachusetts Institute of Technology to expose myself to your best people, both students and researchers, and to learn to build a career in research. I share many common research interests with the professors at MIT.

1. Research Experience

1.1 Inter-datacenter traffic optimization:

There exists a lot of prior research on inter-datacenter traffic engineering with the aim of improving network utilization by exploiting demand deadlines and availability targets. However, most of these techniques don't actually result in capacity savings; the inter-datacenter networks are over-provisioned, which fuels the appetite of first-party applications (applications owned by the cloud provider). For instance, background demands are not smoothed over time which results in a higher peak, causing the network to forecast higher usage for the future, leading to increased network provisioning.

I worked on developing a formal framework, *Approv*, that provisions the network optimally, without affecting SLA requirements, using a data-driven link failure model. Approv takes as input different types of demands: interactive and background demands (with deadline information) along with the availability target for each demand. Such demand differentiation in conjunction with the link failure model results in optimizations such as provisioning less network redundancy for a background demand if the time to repair the link is lower than the demand deadline. We evaluate *Approv* using data from a large cloud service, Microsoft WAN topology including its historical failure data, and show **capacity reduction of ~30% and ~40%** in NAM and APAC regions respectively. This work, in which I am the lead author, has been accepted at **USENIX NSDI '22**. Ideas based on this work are also currently in the process of being deployed by the Azure Networking team.

1.2 Network Function Verification:

Network Function(NF) verification allows verification of software NFs for properties such as reachability, absence of loops, etc. These NFs are often stateful and this state is modified either when a new packet arrives or when a timer expires. The prior work on NF verification focused on the former, i.e, packet arrival event logic. We found, from real-world NF implementations such as the Linux kernel Netfilter subsystem, that the latter, i.e. time-driven logic, is quite common. For instance, consider a simple firewall in which a SYN packet of a TCP flow punches a hole in the firewall. This allows subsequent packets to be routed through, and if no further packets arrive by a threshold time, i.e. by the time a timer expires, the connection is closed.

I worked with **Prof. Wenfei Wu** at Tsinghua University, to extend their NF modeling language to include the timer-driven logic, built the language parser and a symbolic execution tool, and verified the correctness of 14 different types of NFs. We also proposed some optimizations to improve the scalability of our symbolic execution tool. This work is published at **IEEE MASCOTS 2020**.

2. Journey from physics to systems:

I graduated from IIT Roorkee in 2020 majoring in Engineering Physics. My interest in systems and networking arose during my second year when I started contributing to the **Linux Kernel** Netfilter subsystem (both as a **Google Summer of Code** student and as an **Outreachy** intern). I was also an **ns-3 Summer of Code** student under the guidance of **Dr. Tom Henderson**. From these experiences, I recognized how active research in networking leads to the development of new features in open-source projects and how contributing features to open-source projects allow for the democratization of research ideas while also making them more easily reproducible. This encouraged me to pursue research in this field. My interests, motivation, and persistence to learn on my own, by taking math and

CS electives at my university and online courses, are some of my main strengths which I believe will prove crucial during my Ph.D.

3. Why MIT?

Prof. Hari Balakrishnan and **Prof. Mohammad Alizadeh** have contributed seminal work in congestion control such as Congestion Manager, DCTCP, Conga, HPCC, COPA, ABC and more. Their recent work on **CCAC** - formal verification of the CCAs has the potential to improve not just congestion control research but also ML for networking and network verification research. Formally verifying ML algorithms for CCA and ABR algorithms such as **Remy** and **Pensieve** is a difficult problem due to the lack of ML model interpretability. Existing work on Reinforcement Learning (RL) for CCA, **Aurora**, has promise in performing better than the hand-crafted state-of-the-art CCAs such as BBR and COPA. It would be interesting to compare Aurora's performance with non-ML CCAs like BBR and COPA in the adversarial network traces generated by CCAC for BBR and COPA. Further, these generated network traces can also be introduced to the RL model during it's training phase itself.

CCAC can be extended to allow for multiple flows competing on a bottleneck, and also in-network support, to verify **fairness** among competing CCAs. We could further extend CCAC for combinations of **CCAs and AQMs**; this would require rethinking the generalized token-bucket filter approach for modeling network boxes. With support for multiple flows and AQMs, manually writing the first order logic for network settings would be very tedious; there exists a lot of work in formal verification which could allow learning this first order logic from the implemented code. If we implement these network models in ns-3 for example, from my previous experiences of working in ns-3 and network verification, I believe we could learn first order logic from the ns-3 models. I'm very interested in working with **Prof. Balakrishnan** and **Prof. Alizadeh** on such problems!

In his recent SigComm talk, Prof. Balakrishnan spoke about how **applications** have inspired his research over the years. From my experience of collaborating with the Azure Networking and Azure Traffic Manager teams, I realize that network management is still often a manual and tedious process and I see how this affects the end user applications and network managers. This has inspired me to also pursue network management research and I'm very interested in working with Prof. Balakrishnan and Prof. Alizadeh in this domain. Recent advances in programmable switch monitoring, such as their work on **Marple,** allow scalable in-network monitoring. The *collection servers* defined in Marple can be extended as a network management plane which monitors the global network traffic. This can then send feedback to the Traffic Engineering engine, to rate limit, or to the control plane, to install new routes.

Although deployment of programmable switches has started in private networks, like Alibaba and Baidu, we are still ways away from a full wide-scale deployment. I'm interested in working on the problem of **in-network monitoring for hybrid networks,** networks with programmable and legacy switches, thereby further incentivizing incremental deployment of programmable switches. From my experience of working with the Microsoft WAN, I know that the network topology and path information is readily available for private datacenter and inter-datacenter networks; I envisage an application analogous to Google Maps for such hybrid private networks which predicts queuing delay, link utilization, partial and full failures by learning from sparse real-time switch measurement data and the underlying network topology. As a next step, I believe, we can fix these network issues using Traffic Engineering and control plane updates.

I also really like **Prof. Hari Balakrishnan**'s recent work on RFocus. I believe my physics background has provided me with a unique perspective and excitement for such interdisciplinary problems. Although I have mentioned a few professors and research ideas above, I am broadly interested in pursuing systems and networking research with these and other researchers at MIT such as **Prof. Manya Ghobadi** . I believe my strong background in research and my software development experience, and alignment of research interests with the professors at MIT make me a very good fit for your graduate program.