Allison Turner
UCSD Department of Computer Science & Engineering, CS PhD Program

I want to become an internet measurement and architecture researcher because it marries the sheer joy of solving computer systems puzzles with my steadfast belief in equitable access to resources.

Though I have never lived in a world without computers or the internet, they are still in infancy, growing faster and bolder every year. I've traded playing Webkinz on a CRT monitor for flipping between apps on my smartphone. We have an incredible opportunity to shape this growth into something radically accessible, and to take heed of early cautionary tales of poor stewardship.

Net neutrality has fascinated me since I first became aware of its contention; the debate over access to resources and information provoked my sense of justice, and the nebulous inner workings of the internet sparked my technical curiosity, but I had no idea where to begin.

I've never let not knowing where to begin stop me. Instead, I sought a mentor, and found one in Dr. Christine Bassem when I went to her office hours to ask about her work on crowdsensing. My curiosities about internet fairness surfaced, and Dr. Bassem offered to mentor my independent studies!

We started from square one: what did internet research look like? I buried myself in digital libraries, discovering measurement and data analysis techniques. The studies I read on adversarial network measurement inspired my hopes for the future of free access. Eventually I found myself limited by lack of fundamental knowledge of networking and internet architecture, and decided I had to strengthen my foundations. Thus, in fall 2019 I registered for a graduate-level networks research foundation class at MIT. 6.829 was an incredibly difficult but rewarding class, where I ran congestion simulations, read research papers both classic and current, listened to guest lectures by industry researchers, and built a proof-of-concept peer-assisted livestreaming system.

6.829 gave me the missing tools I needed to start figuring out how to design my own study. I began trying to articulate my questions into experiments, and to tinker with measurement tools encountered in class or in literature review. It was then that Dr. Bassem encouraged me to submit to undergraduate conferences and think about graduate school. At first, I was confused at the idea that I could qualify for conferences or become a PhD student. I didn't know anybody who had; it seemed like a career made for mysterious experts. Besides, my grades weren't the best. First jobs, new responsibilities, harder classes, and my then-undiagnosed ADHD meant I was struggling with deadlines and exam scores for the first time in my life. Despite my doubts, professors kept telling me they saw I understood much more than my grades suggested, and that my curiosity, intuition, and passion were some of the best tools a researcher could have. So, here I am, aiming for a future career as a computer scientist that I never would have thought was within reach.

My first-ever conference was the National Collegiate Research Conference at Harvard University in January 2020. I wrote an abstract for the first time for the application, and presented a poster outlining the different categories of possible network inefficiencies or unfairness I had discovered in my literature search. NCRC was interdisciplinary, so I got lots of practice both explaining computer systems in accessible terms, and getting insightful feedback on experiment ideas. Discussing my work gave me an unforeseen level of clarity on my purpose in researching internet fairness, and talking to other undergraduate researchers about their work filled me with senses of wonder. If becoming a computer scientist meant building a career filled with this wonder, then there was nothing I wanted to do more.

I returned to my research reinvigorated, this time to define my research goal: to measure how ISP and AS routing decisions affect everyday consumers' quality of experience via crowdsourced network measurement Unfortunately this is when the COVID-19 pandemic descended. The National Conference on Undergraduate Research, where I had been accepted as a plenary speaker, and the Consortium for Computing Sciences in Colleges New England Conference, where I had been accepted as a poster presenter, were cancelled. Though these cancellations were sensible public health decisions, I regretted the loss of experiences akin to NCRC.

In summer 2020 I pressed on virtually as a member of Dr. Bassem's lab, returning to my experiment design with my mentee and collaborator, Ariel Traver. As I perused internet topology studies, it became apparent that we would likely not be able to get a full view of a measurer's local topology with our own tests. However, I knew that public data sets already existed, and that their data would likely use better techniques for alias resolution and topology inference - so why reinvent the wheel? Thus I formulated stage 1: run traces to common domains from a measurement vantage point, and match traceroute points to nodes in the Center for Applied Internet Data Analysis' Internet Topology Data Kit (CAIDA ITDK). I then turned to the next stage: knowing how different packets are routed. I chose to use HTTP requests, varying the target domain, to apply the strategies of Paris Traceroute for path correctness, to "disguise" tracing packets, and to choose non-measurement-associated ports. Time-to-live value filtering and ICMP drops were unavoidable obstacles, but I could mimic web browser traffic to attempt circumvention of other filters. Stage 3 would focus on client-side QoE metrics. By the end of summer, I had produced a proof-of-concept for stage 1 and created a parser to convert the ITDK into a database schema, and Traver, under my guidance, created our stage 2 test authoring tool with the Python packet crafting library Scapy. I tried out submitting to more serious research events with peer reviewed proposal processes, such as the USENIX Workshop on Free and Open Communications on the Internet and Tapia Conference's ACM Student Research Competition. Though FOCI rejected my proposal, reviewers validated my progress and offered valuable critique. I gained many of the same benefits from feedback at Tapia, but contrastingly, I won 1st place at Tapia's ACM SRC!

Now, in fall 2020, I'm writing an IRB proposal for the crowdsourced participation phase, setting up the data repository, and making the measurement code into cohesive, portable software. My goal is to crowdsource geographically distributed participants to measure from a diverse set of vantage points, each with different autonomous system peering or providing relationships, approaches to privacy, and resource availability. I have also virtually attended USENIX Security and the ACM Internet Measurement Conference to immerse myself in current research.

During my undergraduate career, I also explored the future of computing from another angle: elementary computer science education. I joined Code Ninjas Wellesley, a small CS education franchise, in November 2018. I guide kids ages 7 to 14 through a self-paced curriculum of game and animation projects in Scratch and JavaScript. At least, that's what my basic job description is. After identifying conceptual gaps in our program, I have exceeded my teaching role to meet curricular needs by creating original lessons, workshops, and videos, performing curriculum audits and revisions, and am currently developing content for an in-house learning management system. This work includes a workshop introduction to binary, boolean algebra, and logic gates using Snap Circuit toys and Scratch simulators; a multi-module Scratch checkers game, teaching clone management and skills for managing larger projects; a revamped JavaScript curriculum to integrate supplementary modules.

My foray into elementary CS education also led me into HCI research. I joined Dr. Orit Shaer's TangiBac project, where I worked on a prototype of CRISPEE, a toy used to teach kindergarteners about genes. Using results of previous testing, my teammates and I designed and fabricated a user interface with a more intuitive workflow, sturdier and cheaper internal components, and new, extendable educational features. CRISPEE's secondary goal was to be an open source fabrication project for high schoolers, and so after I helped administer user testing at the Boston Children's Museum, my teammates and I began building a website guide to replicating CRISPEEv3.

I am applying to UC San Diego because of its faculty's groundbreaking work in internet measurement for privacy, system integrity, and efficiency. UCSD's close relationship with the Center for Applied Internet Data Analysis is also a point of interest. CAIDA's ITDK plays a significant role in my crowdsourced routing and QoE experiment design, and I studied several publications from CAIDA members such as Dr. Kimberly Claffy to learn about "adversarial" topology and network measurement. UCSD and CAIDA have long stood out to me as pioneers in the development of network auditing and inference methods. My own goals as an aspiring computer scientist to keep network entities accountable and build better systems align with their missions. I would be eager to contribute my skills to the work of Dr. Claffy, Dr. Stefan Savage, Dr. Geoffrey Voelker, or the other standard-bearers of the field at UCSD.

As I look to a possible future in graduate school, I happily anticipate many aspects of doctoral scholarship: helping to teach undergraduates, learning everything I can about computer systems, contributing code, writing, and ideas to ongoing research, and forging ahead with my own experiments. My goal is to use my skills to help push the shape of the internet towards decentralization, democracy, efficiency, and resiliency, and become a better computer scientist in doing so. I am broadly aiming towards the fields of internet measurement and privacy, but other areas of interest include voting security, censorship detection and evasion, and network protocol or architecture design.