RESEARCH STATEMENT

Ragavendran Gopalakrishnan

I am interested in the **design of service systems** and the accompanying operational policies and economic mechanisms that involve interactions among multiple strategic entities whose objectives are not necessarily aligned. *An emerging focus of my research is to better understand key trade-offs between competing economic objectives and how managing these trade-offs constrains the underlying operational decisions and vice versa.* My research emphasizes fundamental contributions to methodology, increasingly motivated by applications to *urban transportation* and the issues impacting its major stakeholders. I combine tools from multiple areas such as optimization, probability, queueing theory, decision theory, game theory and mechanism design for obtaining analytical results. In addition, I employ numerical methods and, more recently, data-driven simulation for validation in real-world environments.

In the following, I discuss three major research streams that I have been broadly pursuing. Within each section, I briefly introduce and motivate the theme, highlight recent work, and discuss future directions.

1. Sustainable Market Models for Modern Urban Mobility

Urban mobility has been undergoing a global transformation into a *user-centric*, service-oriented playing field, from a traditionally *vehicle-centric*, infrastructure-oriented framework. By 2030, roughly 60% of the world's population will live in cities, where most users are just a few taps away from accessing a diverse set of mobility options. While users seek the best price-convenience trade-off, public transit agencies want increased ridership, private mobility providers want their revenue targets met, and the city wants what is best for society. The recent spike in interactions between these stakeholders (both 'friendly' cooperation and 'unfriendly' regulation) is evidence of a growing interdependence between their interests—they cannot ignore each other while making crucial strategic decisions. This presents a huge opportunity for new, multidisciplinary research to make a significant impact in shaping the next generation of urban mobility, by *guiding policy and decision-making towards a fair balance of competing interests, while promoting a sustainable future.*

Ongoing Work: In collaboration with Chamsi Hssaine, Siddhartha Banerjee, and Samitha Samaranayake, I seek to address the challenges in designing a welfare-maximizing urban mobility market mechanism that simultaneously encourages honest participation from utility-maximizing commuters, nonprofit public transit agencies, and profit-sensitive on-demand mobility services. Technology advances that enable Mobility-as-a-Service (MaaS) and the sharing economy mold the underlying problem design spaces along new dimensions. For example, single-payment multi-modal travel (involving more than one mobility service provider) with seamless transfers not only introduces new real-time pricing and capacity/revenue-sharing problems, but also connects *economic* objectives, e.g., social welfare and revenue, to multiple time-sensitive *operational* objectives, e.g., feasibility and 'quality' of multi-modal trips. The market's ability to facilitate efficient multi-modal solutions by injecting complementarity into otherwise substitute services, can be exploited to possibly overcome limitations of traditional mechanism design. In order to test our analytical insights in a real network, I am concurrently leading the design and development of a data-driven, city-scale, multi-modal mobility simulator that can support various market models and economic control policies, and have initiated discussions with TCAT (Ithaca's public transit agency) to identify opportunities for collaboration.

Future Agenda: Natural extensions include investigating the environmental footprint of mobility markets and impact on personal vehicle ownership. Given the promise of multi-modal travel, an important problem is to identify meaningful characteristics of scheduled and on-demand services (in relation to the underlying network structure and demand pattern) that most benefit from multi-modal trips. For example, multi-modal trips are unlikely to matter much in cities where public transit is either too bad (because they are infeasible) or too good (because they are unnecessary). Such results would help identify existing cities for potential pilots, and provide valuable inputs to long-term urban planning and smart-city projects.

2. Designing Service Systems with Strategic Servers

The vast literature on service systems has traditionally modeled servers as having fixed (possibly heterogenous) service rates, given which, system policies such as routing, matching, staffing, etc. are proposed and analyzed. In reality, when the servers are *people*, the rate a server chooses to work can be impacted by these policies. For example, the Fastest Server First routing policy may lead to servers slowing down to avoid becoming too busy. Similarly, always assigning extra staff to a busiest division of a service system can lead to servers slowing down to ensure that their division is assigned the extra staff. Indeed, strategic server behavior has been observed in the context of journal peer-reviewing (referees enjoy flexibility over scheduling their time towards working on their assignments and when to submit their reports), call centers (agents have some flexibility over how much time they spend on any given call), and taxicabs (drivers may have flexibility in their route choice and driving speed), to name a few. Therefore, system policies that are provably optimal under classical models where servers are nonstrategic, can end up being far from optimal as a result of undesirable strategic incentives created by these policies. Consequently, it is crucial for service systems to be designed in a manner that provides the proper incentives for such strategic servers. While there has been considerable research on designing performance-based rewards/penalties that augment the system policies, the *systemic incentives* created by the policies themselves are much less understood.

Recent Work: In joint work with Sherwin Doroudi, Amy R. Ward, and Adam Wierman [3], published in *Operations Research*, we introduce a model for strategic servers where they enjoy idle time, but incur a cost of effort/fatigue. Then, we study routing and staffing policies in a strategic queueing system that result in Nash equilibrium service rates, and minimize the total equilibrium system cost (a linear combination of staffing and waiting costs). First, we show a surprising policy-space-collapse: all "idle-time-order-based" routing policies (e.g., Longest/Shortest Idle Server First) are equivalent. Then, we characterize asymptotically optimal staffing policies, which we find must staff strictly more than the counterpart "square-root staffing policy" in the nonstrategic setting. In subsequent independent work [2], presented at the *2018 MSOM Conference*, I jointly optimize the routing policy (over a large class of "rate-based" policies, e.g., Fastest/Slowest Server First) and the system configuration (pooled with a single queue vs. dedicated with parallel queues). Interestingly, while servers prefer working faster in a dedicated configuration, it may not be enough to overcome its systemic inefficiency for a superior performance. Moreover, system-optimal policies also lead to the least utilities for servers, resulting in a 'moral dilemma' for a manager who also values employee satisfaction.

Future Agenda: It would be interesting to revisit several strong/fundamental insights in the traditional queueing literature, when servers are strategic, e.g., service systems with strategic customers, optimal system configuration with heterogeneous servers, analysis of the popular "Join the Shortest Queue" policy.

3. Sequential Individual Rationality and Fairness

It is well-known that *variability* in waiting and/or service times in a service system negatively affects the perceived Quality-of-Service (QoS). This problem is further exacerbated by the perception of unfairness in systems that do not guarantee First-In-First-Out (FIFO) service. More generally, when system policies (e.g., routing, matching, staffing, pricing, etc.) allow for future events to cause adverse "shocks" to current customers' *expected* utilities, they experience "frustrations". For example, a burst of high-priority arrivals may cause low-priority customers to wait more. Similarly, ridesharing detours may overshoot promised service time guarantees. Experiments suggest that firms benefit from proactively compensating customers who experience such shocks, without waiting for frustrated complaints [1]. These observations motivate introducing utilitarian concepts of QoS and fairness in shared service systems, that regulate fluctuations in customers' expected utilities *during* their time in the system. Accordingly, we call them *Sequential Individual Rationality* (SIR) and *Sequential Fairness* (SF). For example, a strong notion of SIR could require that a customer's expected utility be nondecreasing. Our goal is to study appropriate notions of SIR/SF in various settings, characterize SIR/SF-compliant system policies, and design optimal *QoS-aware* policies that may be SIR/SF-defiant, but account for the resulting penalties to 'smartly' manage violations.

Recent Work: In joint work with Theja Tulabandhula and Koyel Mukherjee [5], presented at the *2017 MSOM Service Operations SIG Meeting*, we study the consequences of imposing strong notions of SIR/SF in dynamic ridesharing systems. First, we consider a commercial setting, where the matching and pricing policies are driven by a desire to increase profit. We show, perhaps surprisingly, that the firm can potentially obtain a higher profit by adopting 'risky' matching and pricing policies that could be SIR-defiant, despite incurring penalties for such violations. Next, we consider a community carpooling setting where fair group formation and cost sharing are the most important concerns. Here, requiring SIR-compliant matching and cost sharing policies induces natural limits on the incremental detours permissible, leading to sublinear bounds on the fractional delays. In addition, our characterization of sequentially fair cost sharing policies implies a strong requirement that *passengers must compensate each other for the detours that they cause*. This work guided the implementation of the core cost sharing method (patent pending) of Xerox/Conduent's ridesharing subsystem, which was then successfully piloted as *Go Bengaluru*, a multi-modal trip planner with peer-to-peer ridesharing for the city of Bangalore. Together with my sustainability-focused research on electric vehicles, it earned me a *Scientific Excellence Award*. I am working with Tulabandhula to test these myopic policies on system-wide operations through simulations based on New York City's taxi data.

Future Agenda: Extensions include exploring the environmental (i.e., vehicle-miles) and algorithmic (i.e., computational hardness of vehicle routing problems) consequences of SIR/SF in ridesharing systems, and possible connections to characterizations of stable cost sharing schemes from my doctoral work [4]. Suitable notions of SIR/SF are also relevant in other models such as priority queues, where SIR-compliant responses to shocks (e.g., bursty high-priority arrivals) can be *systemic* (e.g., adding more staff) rather than *monetary*.

In summary, I am motivated by a strong desire to execute fundamental methodological research concerning the foundations of strategic service operations, with an equally deep focus on applications to the transportation domain, where my interdisciplinary technical background in academia, coupled with industrial experience in developing MaaS solutions, is well-suited to tackle emerging complex and multi-faceted problems. In executing this research agenda, I envision establishing active collaborations with fellow researchers in my area as well as from companion areas such as behavioral operations management, machine learning and data science, and strategic partnerships with interested government and business stakeholders.

REFERENCES

[1] M. Cohen, Fiszer M.-D., and B. J. Kim. Frustration-based promotions: Field experiments in ridesharing. Working Paper, 2018. Available at SSRN: `https://ssrn.com/abstract=3129717`.

[2] R. Gopalakrishnan. On the architecture of service systems when servers are strategic. Presented at the MSOM Conference, 2018. Working Paper.

[3] R. Gopalakrishnan, S. Doroudi, A. R. Ward, and A. Wierman. Routing and staffing when servers are strategic. *Oper. Res.*, 64(4):1033–1050, 2016.

[4] R. Gopalakrishnan, J. R. Marden, and A. Wierman. Potential games are necessary to ensure pure Nash equilibria in cost sharing games. *Math. Oper. Res.*, 39(4):1252–1296, 2014.

[5] R. Gopalakrishnan, T. Tulabandhula, and K. Mukherjee. The costs and benefits of sharing: Sequential individual rationality and fairness. Presented at the MSOM Service SIG Meeting, 2017. Working Paper.