Research Statement - Igor Kadota

My research has two intertwined goals: (i) develop theory and models that capture the fundamental limits and trade-offs of next-generation communication networks; and (ii) contribute to the design and implementation of the networks that will bring to reality emerging applications such as the Internet-of-Things (IoT), Smart-City Intersections, and Shared Augmented Reality.

Common to each of these emerging applications are their *potential to benefit society* and their *need for an underlying communication network* that can satisfy stringent performance requirements in terms of data rates, latency, scalability, and/or resiliency. For example, consider the Smart-City Intersection in Fig. 1 with multiple cameras sending high-resolution video streams to a Remote Server that detects objects (e.g., people, bicycles, and cars) and then sends intelligent real-time feedback to autonomous vehicles and traffic lights. This application has the potential to avoid accidents, improve traffic efficiency, and provide valuable information for policy makers. However, to enable such an application, the underlying communication network should be able to provide high data rates (on the order of multiple gigabits per second) to accommodate the transmission of multiple video streams and ultra-low latency (on the order of milliseconds) to support the transmission of real-time feedback to autonomous vehicles. Notice that *outdated feedback to autonomous vehicles can lead to safety risks*. In addition to high data rates and ultra-low latency, the Smart-City Intersection may also require a communication network that can scale to hundreds of connected devices and that is resilient to communication disruptions (e.g., rain-induced signal attenuation [1]). *These stringent performance requirements - common to many emerging applications - are unachievable by traditional networks, including current 5G deployments.*

Advanced wireless technology such as millimeter-wave (mmWave) and full-duplex will enable beyond-5G networks to support applications that require ultra-high data rate and ultra-low latency. In addition, Software-Defined Radios (SDRs) and Dynamic Spectrum Access (DSA) techniques will allow these applications to coexist and efficiently share the wireless spectrum. However, as wireless systems evolve and become more capable, they can also become more complex. For example, consider the full-duplex system illustrated on the right in Fig. 2. To enable full-duplex communication, the innovative microchip developed in-house by our collaborators [2] needs to be adaptively configured. To that end, I developed an FPGA-embedded algorithm that dynamically adapts the chip parameters aiming to minimize the time-varying self-interference. The experimental results in the ongoing FlexICoN project, described later in detail, show how *the principled design of algorithms combining theory-based and data-driven approaches can be critical for the operation of chips (and other advanced wireless systems) as they become more powerful and complex.*

My interdisciplinary background and research approach (illustrated in Fig. 2) **combining theory, machine learning (ML), and advanced wireless systems** places me in a unique position to contribute to the development of next-generation communication networks. To develop network control algorithms that satisfy application-specific requirements in terms of data rates, latency, and/or information freshness, I use *classical tools* such as Renewal Theory [3–9], Stochastic Coupling [8–10], Dynamic Programming [8–12], Lyapunov Optimization [4–9,13], Multi-Armed Bandits [6–10], and Model Predictive Control [1], as well as *ML-based approaches* such as Long Short-Term Memory (LSTM) models [1], Stochastic Bandit algorithms [14] (e.g., Upper Confidence Bound and Thompson Sampling), and Reinforcement Learning (in two ongoing projects). Then, to validate my solutions, I employ *advanced wireless systems* from the NSF PAWR COSMOS testbed, including

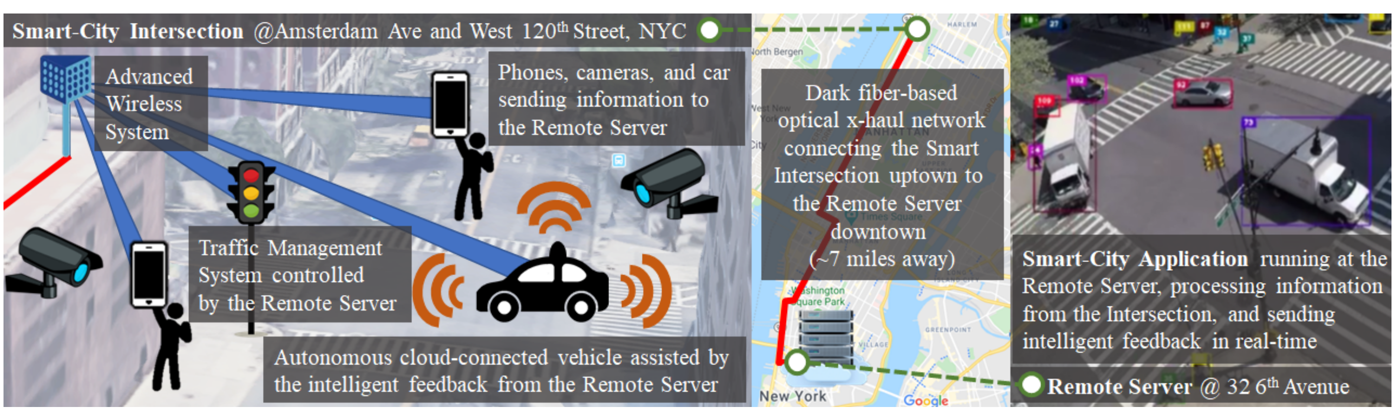

Figure 1: Illustration of a Smart-City Intersection leveraging the (actual) deployment of cameras and Software-Defined Radios (SDRs) in West Harlem and the data center in Lower Manhattan.

full-duplex chips [2], mmWave systems [1,15] (e.g., 77 GHz FMCW radar and 28 GHz channel sounder), FPGA-enabled SDRs [3,16,17], and small embedded computing platforms [18,19] (e.g., Zybo Z7 FPGA and Raspberry Pi). Another important aspect of my research are collaborations. The ongoing projects on full-duplex [2, 20], mmWave [1,15], DSA [17,21], and opportunistic weather sensing are active **collaborations with researchers from various universities and companies on domains spanning drone swarms, weather monitoring, Smart-Cities, and microchips.** Next, I describe selected projects and future research plans.

Wireless Networks for Time-Sensitive Applications: Theory and Systems

This project addresses communication networks that carry time-sensitive information. These networks are relevant for applications such as: IoT with several sensors transmitting data to a remote monitor, autonomous vehicles exchanging position information, and cooperative robots in a factory sharing status information. In such application domains, it is essential to *keep the information fresh, as outdated information loses its value and can lead to system failures and safety risks.* The **Age-of-Information (AoI)** metric, proposed by Kaul et al. in [22], captures the freshness of the information (generated at the source) from the perspective of the destination. To keep the information fresh, i.e. minimize AoI, it is necessary to simultaneously provide: (i) low latency; (ii) high data rate; and (iii) service regularity. Notice that optimizing one of these three performance metrics can be a challenge in itself. My research was the **first to develop network control algorithms with provable performance guarantees in terms of information freshness** for wireless networks with shared and unreliable channels [8, 10] and it was also the **first to implement and evaluate a freshness-based scheduling algorithm in a real-world wireless network** [16, 18].

Rigorous Theory. In prior work, I have developed different types of algorithms (e.g., distributed/centralized) with different performance requirements (e.g., optimal/near-optimal with performance guarantees) and different levels of computational complexity. Selected contributions include:

- obtaining a high/low-complexity AoI-optimal transmission scheduling algorithm for general/symmetric single-hop wireless networks using Dynamic Programming/Stochastic Coupling techniques in [11]/ [10], respectively. These were the **first AoI-optimality results in the literature** for wireless networks with unreliable links. Obtaining a low-complexity algorithm that is AoI-optimal for general networks is still an open problem;
- deriving lower bounds on the achievable AoI performance for different network settings and deriving upper bounds on the AoI performance of various scheduling algorithms [3–8]. By comparing lower bounds with the associated upper bounds, I established performance guarantees in terms of AoI. For these contributions, [6] received the **Best Paper Award at IEEE INFOCOM 2018**;
- evaluating the combined impact of traffic load, queueing discipline, and scheduling algorithm on the AoI of the network [4, 5]. For these contributions, [4] was a **Best Paper Finalist at ACM MobiHoc 2019**.

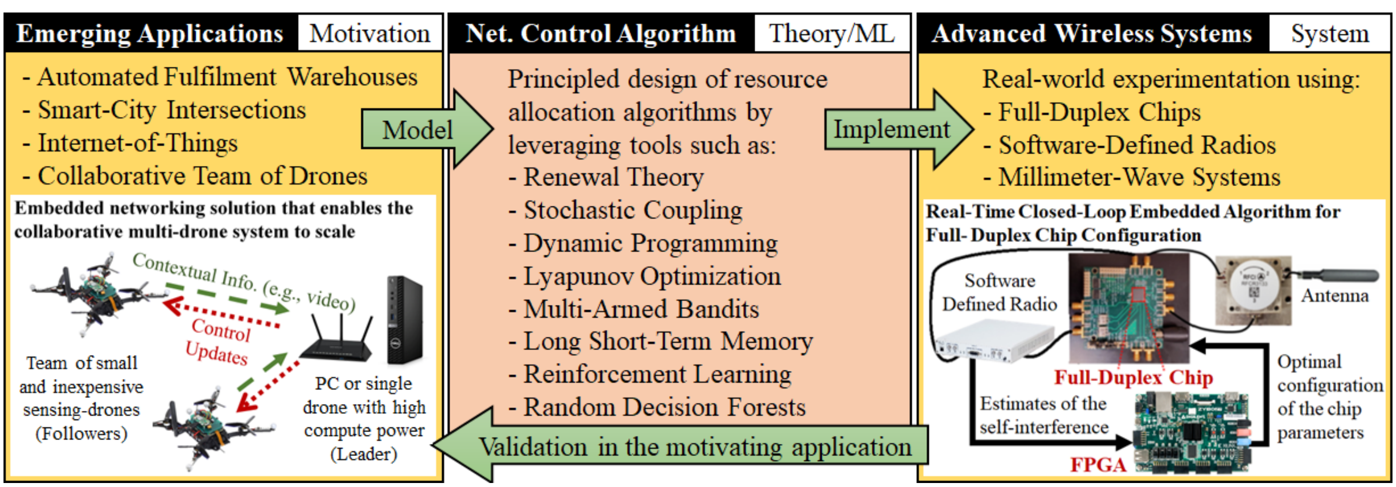

Figure 2: Research approach. To design networking solutions for a given emerging application, I start with theory. This allows me to remove technological constraints and work on the core problem, capturing fundamental limits and trade-offs. Then, I develop principled solutions with a mix of theory and ML. Finally, I use programmable wireless systems to prototype and validate the impact of my solutions in real operating scenarios. The pictures of the emerging application and advanced wireless system are from the (actual) hardware utilized in two different publications: Collaborative Team of Drones [19] and Full-Duplex Chip Configuration [2].

Bridging Theory and Systems. Leveraging the theoretical results in [5,7,9], I proposed WiFresh [16,18]: an unconventional network architecture that achieves near optimal information freshness in wireless networks *of any size, even when the network is highly congested.* To do so, the architecture combines three elements: Last-In First-Out queueing discipline, Polling Multiple Access mechanism, and freshness-based transmission scheduling. I implemented WiFresh on the CNRG testbed composed of eleven FPGA-enabled SDRs [18] and on a mobility tracking application [19] (illustrated in Fig. 2) in which a central controller is receiving video streams from multiple drones and then transmitting trajectory updates, allowing the drones to collaboratively track multiple moving objects on the ground (video of flight tests). **This was the first work to implement a freshness-based scheduler in a real network and also the first work to demonstrate the impact of an underlying AoI-optimal network on the performance of a real-world time-sensitive application.** The experimental results in [19] show that the embedded networking solution can improve information freshness by a factor of 109x/48x and tracking accuracy by a factor of 4x/6x, when compared to standard WiFi-UDP/WiFi-TCP networks, respectively. Most importantly, our results also showed that the performance gains of our embedded solution increase as the system scales and/or the traffic load increases.

Beyond-5G Wireless Systems: Millimeter-Wave and Full-Duplex Chips

Predictive Weather-Aware Reconfiguration of mmWave Networks. A major challenge associated with mmWave networks is their susceptibility to weather conditions. In particular, *rain may cause severe signal attenuation*, which can significantly degrade the network performance. In [1], I developed a network control algorithm that combines classical and ML approaches, using historical data to predict the future condition of the communication links in order to prepare the network (ahead of time) for imminent disturbances. The algorithm has two components: (i) *ML-based attenuation prediction* that employs an encoder-decoder LSTM model to estimate future attenuation levels of each link in the network. (ii) *Network reconfiguration* that leverages these predictions to dynamically optimize routing and admission control decisions aiming to maximize network utilization, while preserving max-min fairness. I trained, validated, and evaluated the algorithm using a dataset containing over 2 million measurements collected from a real-world city-scale backhaul network in Gothenburg, Sweden, by Ericsson AB. The results showed that the algorithm: (i) predicted attenuation with high accuracy, with an RMSE of less than 0.4 dB for a prediction horizon of 50 seconds; and (ii) improved the instantaneous network utilization by more than 200% when compared to *reactive* network reconfiguration algorithms that cannot leverage information about future disturbances. A patent including some of the results is pending [23]. In future work, I plan to consider alternative ML-approaches for predicting attenuation (e.g., Transformer models) and for dynamic network reconfiguration (e.g., Reinforcement Learning).

FlexICoN project: Real-Time Closed-Loop Configuration of Full-Duplex Chip. Existing wireless systems such as Wi-Fi and cellular networks operate in half-duplex mode, where radios transmit and receive either at different times or different frequencies. Full-duplex radios can transmit and receive at the same time and frequency, thus significantly improving throughput and latency. A major challenge in full-duplex is to cancel the *extremely strong* and *time-varying* self-interference (SI) signal. Recent advances in microchips enabled SI cancellation using compact programmable chips [20]. Consider the full-duplex system illustrated in Fig. 2. The SDR is continuously sending estimates of the SI signal to the FPGA which is running an embedded chip configuration algorithm. The goal of this algorithm is to find the combination of chip parameters that minimizes SI. Two challenges are: (i) the extremely large number of possible combinations of chip parameters ($> 10^{19}$), which makes brute force solutions too computationally costly and time-consuming; and (ii) the non-idealities of the hardware, which may lead to a significant mismatch between the expected and actual chip responses. In [2], I developed a low-complexity algorithm that captures chip non-idealities using a data-driven model and searches for a near-optimal configuration using a heuristic Greedy-like algorithm. The experimental results in [2] showed that the chip provided an SI cancellation of 31 dB over a bandwidth of 80 MHz at a center frequency of 800 MHz. In ongoing work, I have developed a new algorithm inspired by Orthogonal Matching Pursuit that incorporates expert knowledge of the chip (e.g., "nearly" equivalent subsets of configuration parameters) to improve performance. Experimental results (using the same chip from [2]) showed that the new embedded algorithm can achieve SI cancellations of 41 dB and 33 dB over bandwidths of 80 MHz and 160 MHz, respectively. **These results clearly show how algorithms can be critical for the operation of chips as they become more powerful and complex.** In the future, I plan to consider ML-based chip configuration algorithms (e.g., Deep Neural Networks) that can be trained offline thus, further reducing the online run time.

Future Research Directions

Emerging applications with ever more stringent performance requirements and advanced wireless systems with increasing complexity will continue driving the need for novel networking solutions. I plan to continue working at the intersection of theory and systems, using rigorous analysis to support the development of practical networked systems that are tailored to mission-critical applications. In particular, I am interested in pursuing the following research programs:

1) Networking for Collaborative Multi-Agent Systems. Consider a swarm of drones, a fleet of autonomous vehicles, or an automated warehouse. In these systems, agents with incomplete information collaborate towards a goal that can be monitoring an emergency scenario, safely driving passengers and/or manipulating objects. Communication is key for collaboration, and it can significantly impact the system's performance and functionality. My research so far has been focused on creating mechanisms for agents to share status information using the network infrastructure. I want to explore other application domains, with possibly different types of information (e.g., local point cloud maps), different requirements (e.g., hard-deadline constraints), and no infrastructure (i.e., ad-hoc networks). I intend to continue designing and implementing networking solutions tailored to the application, as opposed to developing general mechanisms for maximizing throughput and minimizing latency. To better understand the application requirements and validate my networking solutions, I plan to actively seek collaborations with domain experts in academia and industry.

2) Networking for Machine Learning Applications. Consider the Smart-City Intersection in Fig. 1 with different sensors transmitting time-sensitive information over the wireless network to multiple ML edge-cloud servers that detect and track objects and then send real-time feedback to automated vehicles. The main challenges imposed by the environment are the large number of vehicles moving at different speeds, the erratic behavior of pedestrians, the varying conditions of the weather, and the frequent obstruction of sensors and communication links. Reliable Smart-City applications require advances in a number of technologies. I intend to develop network control algorithms that can manage the available communication and computation resources by dynamically adapting the wireless network (possibly containing full-duplex and mmWave links) to the time-varying conditions of the environment and, at the same time, intelligently allocating the computational tasks to the appropriate edge-cloud server. I plan to implement and evaluate these algorithms using the COSMOS Smart-City Platform which contains sensors, edge-cloud servers, optical x-haul, and SDRs.

3) Machine Learning for Advanced Wireless Systems. Major challenges in developing solutions for full-duplex, mmWave, and DSA are computational complexity and non-idealities of the hardware implementation. In [1, 2], I used a combination of data-driven models and classical optimization methods to propose networking solutions that addressed those challenges. In ongoing work, I am building upon [17, 21] to propose Reinforcement Learning-based DSA techniques and implement them on the NSF PAWR COSMOS testbed. In another ongoing project, I am using classical Signal Processing techniques together with Random Decision Forests to perform opportunistic weather sensing using mmWave signals. I intend to continue leveraging and combining theory, ML, and advanced wireless systems to develop solutions to next-generation communication networks.

References

[1] **I. Kadota**, D. Jacoby, H. Messer, G. Zussman, and J. Ostrometzky, "Switching in the rain: Predictive wireless x-haul network reconfiguration," *ACM Measurement and Analysis of Computing Systems*, 2022. [ACM SIGMETRICS'23 issue][Acceptance rate 18.3%].

[2] A. Nagulu, S. Garikapati, M. Essawy, **I. Kadota**, T. Chen, A. Natarajan, G. Zussman, and H. Krishnaswamy, "Full-duplex receiver with wideband multi-domain FIR cancellation based on stacked-capacitor, N-path switched-capacitor delay lines achieving $> +54$db SIC across 80MHz BW and $> +15$dBm tx power handling," in *Proc. IEEE ISSCC*, 2021.

[3] **I. Kadota** and E. Modiano, "Age of information in random access networks with stochastic arrivals," in *Proc. IEEE INFOCOM*, 2021. [Acceptance rate 19.9%].

[4] **I. Kadota** and E. Modiano, "Minimizing the age of information in wireless networks with stochastic arrivals," in *Proc. ACM MobiHoc*, 2019. **[Best Paper Award Finalist]** [Acceptance rate 23.7%].

[5] **I. Kadota** and E. Modiano, "Minimizing the age of information in wireless networks with stochastic arrivals," *IEEE Transactions on Mobile Computing*, 2021.

[6] **I. Kadota**, A. Sinha, and E. Modiano, “Optimizing age of information in wireless networks with throughput constraints,” in *Proc. IEEE INFOCOM*, 2018. **[Best Paper Award]** [Acceptance rate 19.2%].

[7] **I. Kadota**, A. Sinha, and E. Modiano, “Scheduling algorithms for optimizing age of information in wireless networks with throughput constraints,” *IEEE/ACM Transactions on Networking*, 2019.

[8] **I. Kadota**, A. Sinha, E. Uysal-Biyikoglu, R. Singh, and E. Modiano, “Scheduling policies for minimizing age of information in broadcast wireless networks,” *IEEE/ACM Transactions on Networking*, 2018.

[9] Y. Sun, **I. Kadota**, R. Talak, and E. Modiano, *Age of Information: A New Metric for Information Freshness*. Morgan & Claypool, 2019. **[Book]**.

[10] **I. Kadota**, E. Uysal-Biyikoglu, R. Singh, and E. Modiano, “Minimizing the age of information in broadcast wireless networks,” in *Proc. IEEE Allerton*, 2016.

[11] **I. Kadota**, “Transmission scheduling of periodic real-time traffic in wireless networks,” Master’s thesis, MIT, 2016.

[12] K. S. Kim, C.-P. Li, **I. Kadota**, and E. Modiano, “Optimal scheduling of real-time traffic in wireless networks with delayed feedback,” in *Proc. IEEE Allerton*, 2015.

[13] R. Talak, **I. Kadota**, S. Karaman, and E. Modiano, “Scheduling policies for age minimization in wireless networks with unknown channel state,” in *Proc. IEEE ISIT*, 2018.

[14] E. Atay, **I. Kadota**, and E. Modiano, “Aging wireless bandits: Regret analysis and order-optimal learning algorithm,” in *Proc. WiOpt*, 2021.

[15] M. Kohli, A. Adhikari, G. Avci, S. Brent, J. Moser, S. Hossain, A. Dash, **I. Kadota**, R. Feick, D. Chizhik, J. Du, R. Valenzuela, and G. Zussman, “Outdoor-to-indoor 28 GHz wireless measurements in manhattan: Path loss, location impacts, and 90% coverage,” in *Proc. ACM MobiHoc*, 2022. [Acceptance rate 19.8%].

[16] **I. Kadota**, M. S. Rahman, and E. Modiano, “Poster: Age of information in wireless networks: from theory to implementation,” in *Proc. ACM MobiCom*, 2020.

[17] D. Stojadinovic, P. Netalkar, C. E. C. Bastidas, **I. Kadota**, G. Zussman, I. Seskar, and D. Raychaudhuri, “A spectrum consumption model-based framework for DSA experimentation on the COSMOS testbed,” in *Proc. ACM WiNTECH*, 2021.

[18] **I. Kadota**, M. S. Rahman, and E. Modiano, “WiFresh: Age-of-information from theory to implementation,” in *Proc. IEEE ICCCN*, 2021.

[19] V. Tripathi, **I. Kadota**, E. Tal, M. Rahman, A. Warren, S. Karaman, , and E. Modiano, “WiSwarm: Time-sensitive wireless networking for a collaborative team of UAVs,” in *Proc. IEEE INFOCOM*, 2023.

[20] T. Chen, S. Garikapati, A. Nagulu, A. Gaonkar, M. Kohli, **I. Kadota**, H. Krishnaswamy, and G. Zussman, “A survey and quantitative evaluation of integrated circuit-based antenna interfaces and self-interference cancellers for full-duplex,” *IEEE Open Journal of the Communications Society*, 2021.

[21] P. Netalkar, A. Zahabee, C. Bastidas, D. Stojadinovic, **I. Kadota**, G. Zussman, I. Seskar, and D. Raychaudhuri, “Large-scale dynamic spectrum access with spectrum consumption models,” 2022. [Submitted].

[22] S. Kaul, R. Yates, and M. Gruteser, “Real-time status: How often should one update?,” in *Proc. IEEE INFOCOM*, 2012.

[23] J. Ostrometzky, G. Zussman, H. Messer, D. Jacoby, and **I. Kadota**, “Predictive weather-aware communication network management,” US Patent Application No. 17/551,643. December, 2021.