

RESEARCH STATEMENT

Maria Gorlatova
mariaag@princeton.edu
Associate Research Scholar
Princeton University

My research focuses on the challenges and opportunities associated with adding connectivity and intelligence to every device big and small – the sub-area of communications and networking currently known as *the Internet of Things (IoT)*. My work, which involves the development of architectures, algorithms, and protocols for the IoT, crosses traditional discipline boundaries and requires thinking across multiple layers of system and protocol stacks.

At the core of my research approach are systematic multi-faceted cross-layer explorations of specific important emerging technologies within the broader space of the IoT, such as, for example, ultra-low-power energy harvesting IoT nodes and heterogeneous fog computing architectures for supporting the IoT. I conduct comprehensive theory-and-practice examinations of such technology areas to identify the most important issues that prevent the realization and scaling of the technologies, and address the identified “research bottlenecks” using a wide range of tools and techniques. The mathematical techniques I use include optimization, algorithm design, graph theory, and stochastic modeling. My work also incorporates a full range of applied components, such as instrumentation design, measurement studies, prototyping, testbed development, and performance evaluations. I develop strong collaborations across traditional discipline boundaries, and partner with industry to ensure broad impacts of my research.

As the rest of this statement demonstrates, my work has already touched upon many elements of the IoT, including energy sources, adaptive behavior, node and network architectures, and distributed heterogeneous computing. Over the next several years I will focus on developing adaptive intelligent distributed IoT systems supported by fog and edge computing, developing and deploying multi-modal IoT systems for understanding human interactions, and addressing the challenges posed by paradigm-changing futuristic IoT technologies such as augmented reality.

Contributions: Characterizations of Energy Sources for Energy Harvesting IoT Nodes

I examined the implications of the properties of two important Internet of Things energy sources, indoor light and object and human motion, on the design of IoT nodes and IoT algorithms [1]–[4].

Indoor light energy harvesting is a key enabling technology for many envisioned IoT devices and applications. To characterize the indoor light energy, I led a *first-of-its-kind long-term indoor irradiance (light energy) measurement study*, in which we designed and developed a system for long-term irradiance data collection, and deployed these systems in a set of locations for over 1.5 years. We used the indoor irradiance traces we collected to obtain insights into energy availability and characteristics. Our study demonstrates the feasibility of powering IoT nodes using indoor light energy, and provides practical guidance for designing indoor light energy harvesting nodes (e.g., determining appropriate solar cell and battery sizes). It also demonstrates and quantifies temporal and spatial variability in light energy availability, and provides guidance in choosing appropriate models for different indoor light energy conditions for energy harvesting adaptive algorithms [1], [2]. **Publication [2] received the 2016 IEEE Communications Society Young Author Best Paper Award.**

I also examined properties of energy associated with *object and human motion*, which are important for kinetic energy harvesting IoT nodes. In this study we conducted unique experiments with object motion in commonplace environments. We also studied human motion using a dataset with over 40 participants, and conducted an acceleration measurements campaign, collecting long-term traces of acceleration information corresponding to daily human routines. Our study demonstrates the range of harvested powers for different participants and activities, provides insights into variability and predictability of motion energy harvesting processes, and uniquely demonstrates the dependence of motion energy availability on human physiological parameters. **This study was featured in the MIT Technology Review in 2013 and 2014 and in the New Yorker Magazine in 2017.**

This work has been cited in publications of a wide range of technical communities, including IEEE Sensors, IEEE Transactions on Power Electronics, Chemical Society Reviews, and IEEE International Symposium on Circuits and Systems, in addition to the top venues of communication and networking communities. We made publicly available the unique measurement traces we collected in these studies [5], [6]. The traces, which can be used as inputs for energy harvesting system simulators and emulators, have been downloaded by researchers from more than 200 institutions world-wide.

Contributions: Energy Harvesting Adaptive Resource Allocation Algorithms

I designed adaptive resource allocation algorithms for energy harvesting IoT systems [1], [2], [7], which require nontraditional approaches. Energy harvesting shifts the nature of energy-aware protocols from prolonging the lifespan of a network to enabling perpetual operation. Moreover, different algorithmic approaches are required for supporting different kinds of energy storage, energy harvesting devices, and energy harvesting environments.

In this work, motivated by the needs of tracking and monitoring IoT applications to communicate consistently despite the dynamic energy conditions, we formulated resource allocation problems aimed at *allocating energy harvesting nodes' time-varying energy in a uniform way with respect to time*. To formulate these problems, we adapted frameworks typically used to achieve fairness among different network nodes to achieve fairness (uniformity) between the different time slots [1]. We developed algorithms for time-fair energy and data rate allocation in networks with different types of energy storage devices. We additionally designed and developed several simple algorithms, for some of which we provided performance guarantees, and evaluated the developed algorithms in simulations and in the energy harvesting network testbed that we designed and developed [7], [8]. This work is the first to model nonlinearity of capacitor-based energy storage systems, and the first to examine analytically the impact of different energy storage quantization levels on the performance of energy harvesting adaptive algorithms.

Contributions: Ultra-low-power Networked Nodes for the IoT

Jointly with collaborators from multiple disciplines, I designed, developed, and prototyped new ultra-low-power energy harvesting nodes for the Internet of Things [8]–[12]. This research was conducted within the framework of a large-scale Energy Harvesting Active Networked Tags (EnHANTs) project, which I contributed to from the very beginning, initiating and leading multiple research explorations.

We envision EnHANTs as small, flexible, self-powered wireless tags. The form factor of the EnHANTs will allow them to be attached to objects that are traditionally not networked (books, walls, doors, keys, clothing). The EnHANTs will be providing the infrastructure for novel tracking applications, for example, locating misplaced objects and continuous peer-based object monitoring. In order for the EnHANTs to communicate and network while harvesting low levels of environmental energy (for example, indoors, harvesting the energy of the indoor lights), they have to spend substantially less energy than the existing low-power communication technologies. Furthermore, the enabling technologies for the EnHANTs also have multiple paradigm-shifting implications on the design of higher-layer protocols [9], [10]. **Publication [10] received the 2011 IEEE Communications Society Award for Advances in Communications.**

I also led a multi-year cross-disciplinary multi-group effort in designing and developing EnHANTs prototypes and a prototype testbed. We presented different stages of the prototypes and the testbed at 6 different conference demonstration sessions and at over 30 other events. **Demonstration [12] received the 2011 ACM SenSys Best Student Demonstration Award.** The prototypes we developed are the first wireless devices that demonstrate multihop data forwarding over ultra-low-power ultra-wide-band physical layer. The testbed we developed is the first to support trace-based light energy control functionality, and the first to support repeatable experiments with energy harvesting adaptive communication protocols with real energy harvesting hardware [8], [11].

Contributions: Aiding Internet of Things with Fog and Edge Computing

My recent research at Princeton University has been focused on enhancing functionality of IoT systems with fog and edge computing, emerging computing paradigms in which computing and storage are placed at multiple locations between the endpoint IoT devices and the cloud. Towards the goal of improving the performance of the IoT with fog computing, we recently designed, developed, and evaluated *fog computing control-as-a-service architectures* for aiding responsive IoT applications [13]. We introduced an analytical framework to investigate optimal design choices for the placement of virtualized controllers along the “cloud-to-things continuum”, and used the framework to investigate latency-reliability tradeoffs in control-as-a-service architectures. We also developed a testbed for evaluating fog-specific decompositions of complex machine learning operations that are traditionally implemented via non-decomposable iterative algorithms. The testbed, which we presented at 3 conferences [14], includes heterogeneous local and cloud computing elements, and captures experimentally obtained latency, cost, and quality tradeoffs associated with different placements of computing functions in fog. We are currently conducting further experimental examinations of in-fog service architectures, elucidating multiple important properties of

fog computing, including, for example, appropriate probabilistic models for end-to-end response times of in-fog serverless computing. We are also preparing our code and data for future public release via GitHub and CRAWDDAD.

An important part of this work is *enabling technology transfer to industry* via our engagement with the OpenFog Consortium that accelerates the development of fog computing architectures. The OpenFog Consortium includes over 50 leading industrial members, such as Intel, CISCO, Microsoft, ARM, and Dell. As the 2016-2017 Co-chair of the OpenFog Consortium Communications Working Group, I contributed to its recently published reference architecture that is *expected to be adopted as an IEEE standard* by April 2018, and led the definitions of advanced fog communication and networking use cases. Via the OpenFog Consortium, we will translate the insights of our ongoing work to industry-standard architectures.

Research Plans

While individual connected objects are now commonplace (e.g., lightbulbs, toys, wearables), enabling the full potential of large-scale ubiquitous Internet of Things systems requires orders-of-magnitude improvements to many aspects of the IoT, including usability, security, node form factor, and system-wide autonomy, adaptiveness, and intelligence. Additionally, emerging IoT deployments and applications will continue to push the boundaries of modern network architectures, and will present multiple opportunities for important societal transformations.

My work involves systematic multi-faceted examinations of emerging technologies, and informs the design of both systems and algorithms. Multiple developments within the domain of the IoT require such approaches. Over the next 2–5 years I will be tackling the following specific challenges:

- *Restructuring IoT applications to make use of fog and edge computing*: Fog and edge computing represent an exciting frontier of convergence of large-scale cloud computing technologies and resource-constrained IoT systems, and as such offer important opportunities for orders-of-magnitude improvements in the performance of the IoT. My ongoing work [13], [14] has already hinted at several potentially transformative approaches to aiding IoT with fog computing. For instance, I already started examining *cloud-coordinated in-fog learning* for the IoT (which enables personalized intelligent behavior in the IoT) and *large-scale optimization problem decompositions in multi-level multi-timescale IoT-edge-cloud systems* (which dramatically improve the performance of the IoT). I have already established independent collaborations with researchers from Yale University and IBM Research to address some specific related questions (e.g., appropriate system design abstractions). Over the next 1–3 years I will develop these and related approaches, and will transfer the developed techniques to pioneering industry deployments via my ongoing engagement with the industry-wide OpenFog Consortium.
- *Making augmented reality deployments pervasive and scalable*: Developments in hardware and algorithms recently enabled the emergence of constrained proof-of-concept augmented reality (AR) devices and applications. However, making AR a practical commonplace multi-user technology requires solving a wide range of communications and networking challenges. Building on my expertise in creating integrated networked systems [8]–[12] and in characterizing non-traditional system properties [1]–[4], I have already started examining communications and computing components of AR at multiple levels (headsets, local computing nodes, cloud). Over the next 2–4 years I will characterize the key properties of AR and will design multi-level algorithms to support them. I will explore the use of energy harvesting to aid AR (for example, via using dispersed ultra-low-power energy harvesting nodes as semantic beacons for AR), and will develop new research tools required for working with large-scale distributed AR, such as end-to-end system-wide testbeds and emulators.
- *Using multi-modal IoT to study in-context human interactions*: My prior work on characterizing the behavior of IoT devices in environments the IoT shares with humans [1]–[4] has hinted at the richness of opportunities for non-invasive monitoring of humans in commonplace environments. At the same time, my organizational work made me realize the need for quantitative in-context understanding of *human interactions* (as opposed to tracking physiological states of individuals, as provided by existing IoT systems). Over the next 3–5 years I will design multi-modal integrated fog computing-aided IoT solutions for feature-rich understanding of in-context human interactions, and will use these solutions to inform studies that require quantitative fine-grained understanding of human contacts, such as studies of organizational dynamics.

Long-term I am excited about the convergence of the domains of communications and computing within the IoT, the promise of IoT-tailored advanced machine learning techniques, and the relentlessly continuing integration of IoT systems and day-to-day human lives. The fascinating research problems that arise due to these trends will continue requiring new multi-faceted approaches and field-transcending techniques.

REFERENCES

- [1] M. Gorlatova, A. Wallwater, and G. Zussman, "Networking low-power energy harvesting devices: Measurements and algorithms," in *Proc. IEEE Conference on Computer Communication (IEEE INFOCOM'11)*, Apr. 2011.
- [2] M. Gorlatova, A. Wallwater, and G. Zussman, "Networking low-power energy harvesting devices: Measurements and algorithms," *IEEE Transactions on Mobile Computing (IEEE TMC)*, vol. 12, Sept. 2013. **2016 IEEE Communications Society Young Author Best Paper Award.**
- [3] M. Gorlatova, J. Sarik, M. Cong, I. Kymissis, and G. Zussman, "Movers and shakers: Kinetic energy harvesting for the Internet of Things," in *Proc. ACM SIGMETRICS'14*, July 2014.
- [4] M. Gorlatova, J. Sarik, G. Grebla, M. Cong, I. Kymissis, and G. Zussman, "Movers and shakers: Kinetic energy harvesting for the Internet of Things," *IEEE Journal on Selected Areas in Communications (IEEE JSAC)*, vol. 33, pp. 1624–1639, Aug. 2015.
- [5] M. Gorlatova, M. Zapas, E. Xu, M. Bahlke, I. Kymissis, and G. Zussman, "Light energy measurements CRAWDAD dataset." crawdad.org/columbia/enhants/20110407, Apr. 2011.
- [6] M. Cong, K. Kim, M. Gorlatova, J. Sarik, I. Kymissis, and G. Zussman, "Human motion for the Internet of Things kinetic energy CRAWDAD dataset." crawdad.org/columbia/kinetic/20140513, May 2014.
- [7] M. Gorlatova, A. Bernstein, and G. Zussman, "Performance evaluation of resource allocation policies for energy harvesting devices," in *Proc. IEEE Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (IEEE WiOpt'11)*, May 2011.
- [8] R. Margolies, M. Gorlatova, J. Sarik, G. Stanje, J. Zhu, B. Vignraham, M. Szczodrak, L. Carloni, P. Kinget, I. Kymissis, and G. Zussman, "Energy harvesting active networked tags (EnHANTs): Prototyping and experimentation," *ACM Transactions on Sensor Networking (ACM TOSN)*, vol. 11, Nov. 2015.
- [9] M. Gorlatova, P. Kinget, I. Kymissis, D. Rubenstein, X. Wang, and G. Zussman, "Challenge: ultra-low-power energy-harvesting active networked tags (EnHANTs)," in *Proc. ACM Conference on Mobile Computing and Networking (ACM MobiCom'09)*, Sept. 2009.
- [10] M. Gorlatova, P. Kinget, I. Kymissis, D. Rubenstein, X. Wang, and G. Zussman, "Energy-harvesting active networked tags (EnHANTs) for ubiquitous object networking," *IEEE Wireless Communications*, vol. 17, Dec. 2010. **2011 IEEE Communications Society Award for Advances in Communications.**
- [11] M. Gorlatova, R. Margolies, J. Sarik, G. Stanje, J. Zhu, B. Vignraham, M. Szczodrak, L. Carloni, P. Kinget, I. Kymissis, and G. Zussman, "Prototyping energy harvesting active networked tags (EnHANTs)," in *Proc. IEEE Conference on Computer Communication (IEEE INFOCOM'13)*, Apr. 2013.
- [12] G. Stanje, P. Miller, J. Zhu, A. Smith, O. Winn, R. Margolies, M. Gorlatova, J. Sarik, M. Szczodrak, B. Vignraham, L. Carloni, P. Kinget, I. Kymissis, and G. Zussman, "Demo: Organic solar cell-equipped energy harvesting active networked tag (EnHANT) prototypes," in *Proc. ACM Conference on Embedded Networked Sensor Systems (ACM SenSys'11)*, Nov. 2011. **ACM SenSys'11 Best Student Demonstration Award.**
- [13] H. Inaltekin, M. Gorlatova, and M. Chiang, "Virtualized control over fog: Interplay between reliability and latency." *In preparation for submission to the IEEE Internet of Things Journal*. ArXiv:1712.00100, Nov. 2017.
- [14] T. Chang, L. Zheng, M. Gorlatova, C. Gitau, C.-Y. Huang, and M. Chiang, "Demo: Decomposing data analytics in fog networks," in *Proc. ACM Conference on Embedded Networked Sensor Systems (ACM SenSys'17)*, Nov. 2017.