

Research Statement

Mohan Krishnamoorthy

☎ (703) 989 6434

✉ mkrishn4@gmu.edu

🏠 cs.gmu.edu/~mkrishn4

Introduction

The big data era has dramatically influenced almost every aspect of our modern lives. It has allowed engineers to build compact and accurate models of real-world processes thereby enabling an engineer to visualize the current state of the process and interact with these models at unprecedented levels. A number of different approaches have also made use of the data-intensive computing and storage systems available today to perform analysis and optimization over these models to get actionable recommendations. This is a top priority for industries as they look towards ways to reduce cost and increase efficiency of operations to maintain competitiveness in the global economy. However, today's solutions lack the generality to make the process models reusable. The computing world has made a lot of advances with the representative and query models and languages. But these advances have not been extrapolated to solutions that model real-world processes. This lack of reusability among models, even within the same organization, has led to non-holistic piece-by-piece analysis and optimization practice that has resulted in sub-optimal solutions and recommendations. Additionally, these solutions : (1) tend not to be extensible to additional aspects of machines, processes and non-functional metrics such as qualitative and quantitative parameters covering the QoS of the system, and (2) perform "silo" optimization, which would not achieve a system-wide optimum if an extended underlying system needs to be optimized. This becomes even more challenging as the scale of the system increases to hundreds/thousands of processes deployed at different geographic locations.

My research is focused on the design and development of reusable and flexible (ease of deployment, and ease-of-use) models of real-world processes that would lend a hand to encoding standardization, and developing analysis and optimization solutions for these models across different domains, extensible to emerging computing platforms such as cyber-physical systems and the Internet of Things (IoT). Specifically, my Ph.D. research is driven by complexities inherent in modeling real-world processes and inefficiencies of current optimization algorithms over stochastic models. My research addresses the need for efficient and flexible approaches to overcome these complexities and inefficiencies. By performing extensive analysis of a wide range of real-world manufacturing processes to examine performance issues, designing rigorous models, and managing tradeoffs, my research aims at improving the efficiency and usability of real-world processes with a broad focus on practical and user-centric metrics.

A key goal of my research is to have practical impact and innovative solutions for real-world problems. During my Ph.D., I have successfully applied my research methodology to model and optimize several real-world stochastic systems and enable decision guidance support for scheduling and resource allocation problems over these systems. My investigations of these problems have led to novel and effective solutions. For example, the stochastic optimization algorithm based on deterministic approximations [2,3,11] that I developed is, according to the experimental study performed, an efficient and scalable approach to get actionable recommendations using white box code analysis of the process model. My research has appeared in a number of premier conferences and journals in computer science, decision support, and bioinformatics, specifically, in HICSS IEEE-BigData, JDS, ICS, IJAMT, and BMC-bioinformatics. Overall, I have comprehensive experience in the research process and have worked in close collaboration with multi-disciplinary research scientists at the National Institute of Standards and Technology (NIST) and the Los Alamos National Laboratory (LANL). My vision for the future of my research is to expand the horizon of analyzing complex processes using smart, robust algorithms and software, as well as to balance efficiency of analysis against flexibility of models.

Highly reusable performance models of real-world processes

Advances in computing have permitted flexible and robust data models, language APIs, automated learning schedules, and query languages to parse large amounts of data, visualize, perform analysis and to provide actionable recommendations. Extrapolating these advances to solutions that model real-world processes in complex domains such as manufacturing and supply-chain management however, remains elusive. More specifically, the analysis and optimization solutions in these domains are typically implemented *de novo*, following a linear methodology. Not only does this lead to cost and time intensive development cycles for models and algorithms, but also make them difficult to modify, extend, and reuse.

My research tackles the above issues using advanced data modeling techniques including a noSQL database and query language. The goal of this work is to design and develop complete, encapsulated and reusable models that describe complex real-world hierarchical processes which can, in turn, be used to compose more complex models and perform analysis on them. To support this work, I develop a reusable, modular, and extensible knowledge base (KB) of process performance models for manufacturing [7,9]. These models are supported by a standard interface of (a) input parameters and controls, (b) output metrics and constraints, and (c) arithmetic closed form equations that transform the inputs to outputs i.e., the white box of the model. I exemplify the robustness and reusability of the KB by systematically designing the

repositories for (a) highly descriptive, varying, and non-linear unit manufacturing processes (UMPs) [6], (b) supply chain components such as contract manufacturers and suppliers [5], (c) dense composite service networks consisting of two or more UMPs and supply chain components [5], and (d) interval and batch schedulers over user-defined/time horizons for service networks consisting of inventories and randomness [1,8]. Further, these models are used to perform a number of different analysis and optimization tasks such as prediction, scheduling, what-if analysis, resource allocation, and stochastic optimization.

Prototype system for composition, analysis and optimization of reusable models

A key contributor to the reason why models describing real-world processes are generally implemented anew is the diversity of computational tools, each designed for a different task such as data manipulation, statistical learning, data mining, optimization, and simulation. Because of this diversity, modeling using computational tools often requires the use of specialized low-level mathematical abstractions. As a result, the same knowledge is often modeled multiple times using different specialized abstractions, instead of being modeled once using a uniform abstraction.

To overcome these deficiencies the prototype system provides the end-user with a portal to perform composition, analysis, and optimization of process performance models from the KB. This provides an innovative way to ensure ease of use, simplicity, and cost-effectiveness by leveraging process performance models from the KB to compose more complex hierarchical models and performing optimization operations on them [10]. To achieve this, we build a software framework that allows hierarchical composition, analysis and optimization of UMPs, supply-chain components, service networks, and schedulers consisting of real-world processes from the KB of process performance models [4].

Stochastic optimization algorithms based on deterministic approximations

Stochastic optimization has typically been performed using simulation-based optimization techniques. However, the general limitation of simulation-based approaches is that simulation is used as a black box, leaving the underlying mathematical structure under utilized. From the literature relating to work on Mathematical Programming (MP) we know that, for deterministic problems, utilizing the mathematical structure can lead to significantly better results compared to simulation-based approaches in terms of optimality of results and computational complexity. Stochastic optimization approaches that make use of the mathematical structure of the original problem extract this structure using samples from a black-box simulation, which is computationally expensive, especially for real-world problems composed of complex process networks.

Instead of extracting the mathematical structure using black-box simulation, my work uses the extraction of mathematical structure from a white-box simulation code analysis as part of a heuristic algorithm to solve a stochastic optimization problem. This algorithm is based on deterministic approximations where the mathematical structure is used to approximate a candidate set of solutions by solving a series of deterministic MP problems that approximate the stochastic simulation. Then, stochastic simulations are run on the candidate set using optimal simulation budget allocation methods to refine the set of collected candidates and find the optimal one among them [2,3,11].

The efficiency of this approach is demonstrated by performing an experimental study of running this algorithm over a 21-process real-world composite use-case with 88 decision variables and 12 stochastic constraints. The experimental study demonstrates that this approach to stochastic optimization not only scales well but also significantly outperforms other popular simulation-based stochastic optimization algorithms in terms of optimality of results and computation time [1].

Conclusion and Future Directions

My research is concentrated on practical problems of decision guidance systems with focus on (a) designing general and sound descriptions of real-world processes, and (b) development of efficient and scalable analysis and optimization algorithms. This has allowed me to gain a deep understanding of this research avenue and has equipped me with the skills required to solve these and other problems that are of high impact. Looking forward, I would like to expand the horizon of modeling and analyzing complex real-world processes by designing systems with high efficiency and flexibility. Additionally, I aim to extend my understanding and gain expertise in emerging computing areas/platforms such as serverless computing, cyber-physical systems, IoT, ubiquitous computing, and artificial intelligence so that I can leverage my experience in decision guidance towards challenges involving scalability, efficiency, and flexibility. In the following, I discuss several future directions that I am particularly interested in.

1. My ongoing and previous work has shown that there exists great potential to extend current analytical techniques to representations of large scale processes in both engineering as well as the computing domains. I am interested in exploring these domains to enrich my understanding of the objectives, constraints and challenges involved in modelling and analyzing

- real-world artifacts. Additionally, I am interested in using my experience to enhance the scalability and efficiency of artifact representations and analysis algorithms thus providing actionable recommendations contributing to increased profits.
2. I am interested in exploring algorithmic techniques such as machine learning to automate modeling and analysis efforts for fellow researchers and engineers. Some areas that would be very interesting to explore are:
 - (a) Building a data store of metadata descriptions of real-world processes that are flexible and reusable.
 - (b) Construction of a software framework with a capability to search models and algorithms that solve a variety of analysis and optimization tasks.
 - (c) Building a software architecture that would allow for easy composition, visualization and interaction with the process models and analysis routines.
 - (d) Building a crowd-sourced engineering information store of well described, sound, and composable artifacts, that is modular, extensible, and reusable.
 3. Computing has expanded beyond the Internet and become ubiquitous everywhere in the physical world. It would be worthwhile to rethink software models and analysis in the context of cyber-physical systems, IoT, mobile and wearable devices. Future paradigms of software for these devices would involve complex interactions with users and require the right balance of resources such as CPU, memory, storage, energy, and non-functional metrics. To provide effective and efficient infrastructure support, I am interested in exploring the tradeoffs among all possible objectives including flexibility of models, efficiency of analysis, real-time autonomic sensing, planning and executions, reducing deployment and operation costs, and improving reliability and energy efficiency.
 4. Serverless computing as a new lighter-weighted computing paradigm has transformed how engineers and developers interact with software systems and applications. Compared to traditional VM-deployment based cloud computing, and elastic scaling services, serverless computing simply allows tenants to run fine-grained functions in a reliable, elastic, and *serverless* manner, finally delivering the *pay-as-you-go* promises of cloud computing. This piques my interest in studying and understanding the challenges, objectives, and constraints in cloud service architectures from the perspective of serverless computing. Additionally, I would like to research next generation infrastructure support that aims to improve both real-time and long term decisions dependent on resource/energy efficiency, usability/flexibility, and performance in serverless platforms.

Selected Publications

- [1] **M. Krishnamoorthy**, A. Brodsky, and D. Menascé. Scheduling Stochastic Processes with Inventories over Multiple Intervals. Under preparation. May 2018.
- [2] **M. Krishnamoorthy**, A. Brodsky, and D. Menascé. Stochastic Decision Optimization based on Deterministic Approximations of Processes described as Closed-form Arithmetic Simulation. In *the Journal of Decision Systems*, May 2018.
- [3] **M. Krishnamoorthy**, A. Brodsky, and D. Menascé. Stochastic Optimization for Steady State Production Processes based on Deterministic Approximations. Under review. Dec 2017.
- [4] A. Brodsky, M. O. Nachawati, **M. Krishnamoorthy**, W. Z. Bernstein, D. A. Menascé. Factory Optima: A Web-based System for Composition and Analysis of Manufacturing Service Networks based on a Reusable Model Repository. Under review. Dec 2017.
- [5] A. Brodsky, **M. Krishnamoorthy**, M. O. Nachawati, and W.Z. Bernstein, D Menascé. Manufacturing and Contract Service Networks: Composition, Optimization and Tradeoff Analysis based on a Reusable Repository of Performance Models. In *Proceedings of the 2017 IEEE International Conference on Big Data*, Boston, MA. Dec 2017.
- [6] A. Brodsky, **M. Krishnamoorthy**, W. Z. Bernstein, M. O. Nachawati. A System and Architecture for Reusable Abstractions of Manufacturing Processes. In *Proceedings of the 2016 IEEE International Conference on Big Data*, Washington DC. Dec 2016.
- [7] A. Brodsky, G. Shao, **M. Krishnamoorthy**, A. Narayanan, D Menascé, and R. Ak. Analysis and Optimization in Smart Manufacturing based on a Reusable Knowledge Base for Process Performance Models. In *the International Journal of Advanced Manufacturing Technology*, April 2016.
- [8] **M. Krishnamoorthy**, A. Brodsky, and D. Menascé. Modular Modeling & Optimization of Temporal Manufacturing Processes with Inventories. In *Proceedings of the 2016 Hawaii International Conf. on System Sciences*, Kauai, HI. Jan 2016.
- [9] A. Brodsky, G. Shao, **M. Krishnamoorthy**, A. Narayanan, D Menascé, and R. Ak. Analysis and Optimization in Smart Manufacturing based on a Reusable Knowledge Base for Process Performance Models. In *Proceedings of the 2015 IEEE International Conference on Big Data*, Santa Clara, CA. Oct 2015.
- [10] D. Menascé, **M. Krishnamoorthy**, A. Brodsky. Autonomic Smart Manufacturing. In *the Journal of Decision Systems, Special Issue on Integrated Decision Support Systems*, June 2015.
- [11] **M. Krishnamoorthy**, A. Brodsky, D. Menascé. Optimizing Stochastic Temporal Manufacturing Processes with Inventories: An Efficient Heuristic Algorithm Based on Deterministic Approximations. In *Proceedings of the 2015 INFORMS Computing Society Conf.*, Richmond, VA. Jan 2015.