*Proceedings of the 11th North American International Conference on Industrial Engineering and Operations Management, Milwaukee, Wisconsin, USA, June 9-11, 2026*

*Proceedings of the International Conference on Industrial Engineering and Operations Management*

Publisher: IEOM Society International, USA DOI: 10.46254/NA11.20260158
Published: June 9, 2026

# Agentic AI-Driven Microgrid Energy Management for Resilient Renewable Energy Communities

**Md Monirul Islam**
Assistant Professor
Industrial Management & Technology
Texas A&M University-Kingsville
Kingsville, TX, USA
md.monirul.islam@tamuk.edu

**Nuri Yilmazer**
Associate Professor
Electrical Engineering and Computer Science
Texas A&M University-Kingsville
Kingsville, TX, USA
nuri.yilmazer@tamuk.edu

**Abdallah Farraj**
Associate Professor
Electrical Engineering
Texas A&M University-Texarkana (RELLIS Campus)
Bryan, TX, USA
AFarraj@tamut.edu

**Baisravan HomChaudhuri**
Assistant Professor
Mechanical Engineering
Lamar University
Beaumont, TX, USA
bhomchaudhur@lamar.edu

## Abstract

A resilient renewable energy community relies on a community-scale microgrid that integrates PV power, wind generation, battery energy storage systems (BESS), and controllable loads. Achieving reliable and cost-effective operation under uncertainty (renewable intermittency, grid price volatility, forecasting error, and extreme events) requires both predictive intelligence and closed-loop operational decision-making. This paper proposes an agentic AI microgrid management framework which is built upon three major layers: Strategic Planner, Operational Supervisor, and Executor. A Large Language Model (LLM) acts as the central orchestrator that coordinates specialized agents for data retrieval and forecasting, dispatch optimization, execution, and monitoring. The framework includes key agents such as the Home Agent, Central Agent (Orchestrator), Data Search and Forecast Agent, Decision Agent, Execution Agent, and Monitor & Review Agent. The framework follows an agentic pattern where the LLM is not only a

parameter extractor but also a decision-maker across different stages of the workflow. In addition, the orchestrator can reason over system state, select tools based on the state, and coordinate multi-agent actions using hierarchical orchestration and tools calling, as well as multi-agent, tool-rich operational frameworks emphasizing standardized interfaces (e.g., MCP) and closed-loop execution. The framework will enable (i) adaptive, auditable dispatch decisions, (ii) human-in-the-loop policy review, and (iii) robust operation under uncertain renewable supply and grid conditions.

## Keywords
Microgrid operation, Agentic AI, LLM, Resilience.

## 1. Introduction
Building community-level microgrids can be a promising pathway to improve reliability and decarbonize local electricity supply through high penetration of distributed energy resources, including PV, wind generation, BESS, and controllable loads. However, microgrid operation is inherently complex due to the large system of systems architecture and operational interdependence. The controllers must continuously balance multiple objectives: minimizing operating cost and emissions, maintaining power quality, satisfying device constraints such as inverter limits and BESS state-of-charge, and ensuring resilience under disturbances. In practice, achieving such a balance requires repeated effort and strong domain expertise, especially because the management and engineers often need to navigate multiple tools and datasets for simulation, forecasting, optimization, and operational analysis.

Most of the existing microgrid energy management systems still follow a conventional static pipeline, where forecasting and optimization are executed sequentially in a single-level manner, and re-optimization is triggered manually or through heuristic rules (e.g., price shifts, increasing forecast errors, or measurement drift). This separation between stages becomes increasingly problematic in renewable communities where uncertainty is persistent and time-varying. These systems require adaptive decision updates, scenario evaluation, and the ability to learn from feedback over time. Moreover, although machine learning models are widely used for forecasting, classification, and operational intelligence, they often remain "model-in-the-loop" components and cannot coordinate end-to-end operational workflows to optimize overall system performance.

Recent advances in agentic AI offer an alternative design direction, where microgrid control can be implemented as an orchestrated set of specialized agents coordinated by an LLM. In this setting, the LLM functions as a workflow-level coordinator that supports goal-driven planning and tool execution, rather than serving only as a front-end interface. With strong reasoning and generalization capability, the LLM can coordinate across agents and adapt decisions through feedback-driven execution instead of relying on static, manually triggered re-optimization. However, deploying LLMs directly as standalone solutions in power and energy systems is challenging due to the sensitivity of grid operations, the risk of cyber-attacks, air-gapped equipment, strict constraint enforcement requirements, and the need for tool-grounded, verifiable workflows. As a result, modern agentic systems emphasize structured tool invocation and modular orchestration like the specialist agents handle domain-specific computations, while an orchestrator enforces global ordering and dependencies to improve scalability and robustness.

Recent studies support the need for this transition. As renewable penetration and DER-driven operations increase, automation and decision support in smart grids and community microgrids become essential because manual, tool-fragmented workflows cannot scale to dynamic, data-intensive operations (Omitaomu, 2021; Ahmadi et. al., 2026; Fathollahi, 2025). Surveys of AI methods in power systems show that data-driven models are increasingly applied for stability, reliability, and operational analytics, yet they are often deployed as standalone modules rather than fully integrated operational controllers (Ahmadi et. al., 2026; Fathollahi, 2025; Bose, 2017). Studies related to the AI based operational plan using RL and DRL demonstrate that learning-based controllers can handle nonlinear dynamics and high-dimensional decision spaces for voltage/frequency regulation and microgrid energy management, while safe RL highlights the importance of constraint satisfaction and risk-aware behavior before real-world adoption [Omitaomu, 2021; Glavic, 2019). Multi-agent systems enable distributed coordination among grid entities, but typical Multi-Agent System implementations still require explicit integration logic to connect forecasting, optimization, execution, and monitoring into a consistent closed-loop workflow (Izmirlioglu et. al., 2024). Digital twin research further shows that continuous fusion of historical and real-time data can support what-if simulation, predictive maintenance, and resilience, although interoperability, cybersecurity, and model validation remain key barriers for deployment at scale (Sifat,2023; Mchirgui et. al.,2024). Finally, explainable AI studies emphasize that interpretability and traceable decision rationales are necessary for operator trust, auditing, and governance—requirements that become even more important as autonomy increases (Machlev et. al., 2022).

Motivated by these gaps, this paper proposes an agentic AI-based microgrid control framework for resilient renewable energy communities. The framework adopts a three-layer architecture: Strategic Planner, Operational Supervisor, and Executor and integrates (i) a Data Search and Forecast Agent, (ii) a Decision Agent for dispatch optimization, (iii) an Execution Agent for safe actuation, (iv) a Monitor and Review Agent for structured logging and anomaly detection, and (v) a Central Agent (Orchestrator) that coordinates these modules using tool-grounded reasoning and human-in-the-loop policy review. The proposed design aligns with emerging agentic architectures that emphasize end-to-end autonomy at the workflow level by perceiving operational context, generating multi-step plans, executing actions through standardized tools, and revising decisions through feedback-driven adaptation, rather than relying on sequential, human-intensive procedures and manual re-optimization. The paper's goal can be summarized as follows:

I. Define a microgrid-oriented agentic AI architecture for a resilient energy community that plans and adapts the end-to-end operational workflow.
II. Ensure safe, risk-aware microgrid control management by preserving operational and device constraints and supporting dispatch decisions.
III. Deliver a transparent and deployable framework through a modular multi-agent design with human-in-the-loop oversight and an implementation plan.

## 2. Agentic AI-Based Microgrid Management Framework

The framework adopts a three-layer agentic architecture: Strategic Planner, Operational Supervisor, and Executor. It will transform community microgrid management from a static "first generate a single forecast, then perform a one-time optimization of the plan, and finally execute the dispatch" pipeline into a goal-driven, tool-grounded, closed-loop workflow (shown in Figure 1). Within this layered workflow, the system is implemented as a multi-agent architecture coordinated by a Central Agent, comprising six core agents: Home Agent, Data Search and Forecast Agent, Decision Agent, Execution Agent, and Monitor & Review Agent.

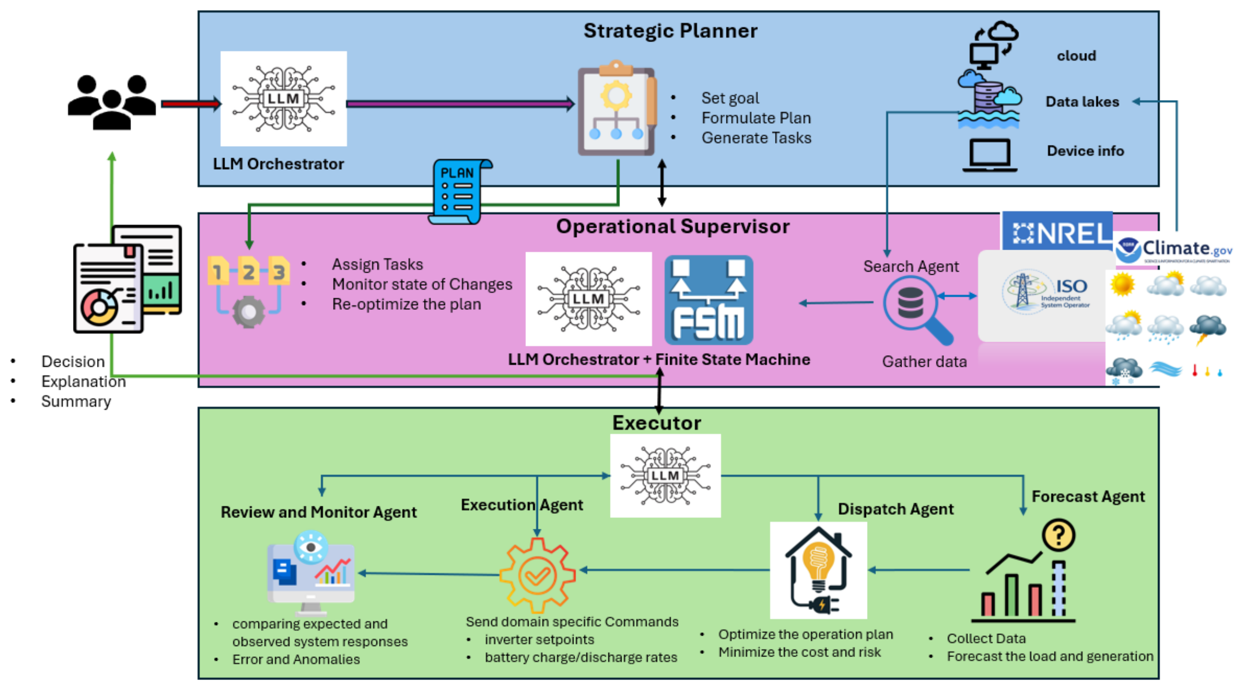


Figure 1. Agentic AI-Based Microgrid Management Framework

The framework treats the microgrid environment (such as telemetry, device controllers, and the community data lake, etc.) as external variables accessed through standardized tools, consistent with an MCP-style separation of reasoning from execution. This approach supports modularity, safety through validation, and transparency via structured logging, while enabling human-in-the-loop governance. The layers and corresponding agents are described below:

## 2.1 Strategic Planner

It is the planning layer where the goal will be determined, decompose the goal into smaller objectives, and develop the policy aware strategy. The Planner converts community objectives and operational context into a structured operational plan. It will begin with inputs from the Home Agent, including user preferences and constraints such as comfort bounds, battery reserve targets, islanding priorities, tradeoffs between cost and benefits, and resilience mode. The Central Orchestrator then constructs a plan using four functions:

- Intent formalization: This function translating objectives into operational targets (e.g., minimizing the total energy cost considering the constraint of SOC greater or equal than 10 percent reserve and zero load shedding for critical loads).
- Task decomposition: The decomposition function will identify what must be computed or retrieved (data updates, forecast refresh, scenario generation, dispatch solve, validation, execution).
- Model and parameter selection: This function will choose which forecasting model or version to run, the forecast or dispatch horizon, and whether robust or risk-aware settings are required (e.g., dispatch during storms, snow, etc.).
- Policy and constraint validation: It will check governance rules and safety boundaries before any action is executed.

## 2.2 Operational Supervisor

This is the coordination layer where the workflow is managed, system states are tracked, and re-optimization is initiated when the Strategic Planner requires the update plan. The Operational Supervisor operates the framework as a continuous operational loop rather than a one-shot pipeline. In this layer, the Central Orchestrator functions as the scheduler and supervisor, which ensures that the tasks are executed in the correct order, intermediate results are shared consistently, and the system adapts as conditions evolve. The core responsibilities include:

- Task Routing: Tasks are assigned to the appropriate agent based on the current need, such as forecasting, decision making, or monitoring.
- Operational Memory Management: The key intermediate outputs such as forecasts, selected scenarios, optimization results, feasibility flags, and execution acknowledgements are stored as short-term memory so that all agents rely on a consistent single source of truth.
- Quality-Gate Monitoring: Forecast drift, solver status, and potential constraint violations are checked before execution to prevent unsafe or unreliable actions.
- Trigger based Re-Planning: The system detects when the current plan must be revised (e.g., forecast error spikes, price shocks, SoC approaching bounds, islanding events, device unavailability, or telemetry anomalies) and then initiates the updated forecasting or prepares re-dispatch accordingly.

This layer addresses a key limitation of conventional energy management systems, manual re-optimization, by continuously evaluating the validity of the operational plan and automatically initiating re-optimization when defined triggers occur. The Operational Supervisor is implemented using a deterministic and auditable Finite-State Machine or event-driven workflow engine, where the operational cycle progresses through states such as ForecastReady, DispatchReady, Execute, Verify, and Replan.

## 2.3 Executor

The Executor performs domain-specific computations and executes physical or virtual actions through specialized agents and tools while enforcing safety and correctness checks before sending any command to devices. Execution is implemented as a closed-loop sequence aligned with the agents, where forecast outputs drive power dispatch decisions, dispatch decisions drive device actuation, and monitoring feedback triggers re-forecasting or re-optimization when demanded.

### 2.3.1. Forecast execution

In the forecasting stage, the data Search and forecast agent retrieves external signals (weather, irradiance, wind speed, prices, and frequency indicators) and generates short-horizon forecasts for renewable generation (PV and wind) and, power demand profile. Let denote $ex_t$ exogenous inputs (irradiance, wind speed, temperature, humidity, etc.) and

$pw_t$ denote renewable output (PV power or wind power). For a forecast horizon $H$ , the agent generates scenario-consistent point and probabilistic forecasts from exogenous inputs and historical measurements, enabling risk-aware dispatch and trigger-based re-optimization under uncertainty:

$$\hat{pw}_{t+1:t+H} = f_\theta(ex_{t-L+1:t}, \hat{pw}_{t-L+1:t})$$

where $L$ is the lookback window length and $f_\theta$ is the model parameterized by $\theta$. Model training minimizes an error objective across the horizon (e.g., MSE or MAE):

$$\min_\theta \sum_t \sum_{h=1}^{H} (pw_{t+h} - \hat{pw}_{t+h})^2$$

Beyond point forecasts, the agent returns forecast-quality indicators such as rolling error statistics, drift flags, and uncertainty proxies, which the Orchestrator and Operational Supervisor use to decide whether to accept the forecast, trigger a fallback baseline, or request a scenario envelope for robust dispatch. This stage is tool-grounded through external data retrieval, feature engineering, and forecast inference tools.

### 2.3.2. Dispatch execution

The Agent employs a hybrid dispatch strategy that combines deterministic mathematical optimization as the primary solver with risk-aware scenario-based extensions. Specifically, the dispatch problem is formulated as a Mixed-Integer Linear Program (MILP) and solved via optimization algorithm. The LLM orchestrator does not solve the optimization directly; instead, it acts as a workflow coordinator that constructs the problem parameters from forecast outputs and system state, invokes the solver tool, interprets the feasibility certificate, and routes results to the Execution Agent. If the solver returns infeasibility or the CVaR-penalized cost exceeds a policy threshold, the orchestrator triggers scenario re-evaluation or relaxes non-critical constraints under human-in-the-loop review. The decision agent maps forecasts and operational constraints into actionable control setpoints. The microgrid decision state is represented as

$$s_t = \{\hat{LF}_{t+1}, \hat{pw}_{t+1}, SoC_t, gp_{t+1}, \pi_t\}$$

where $\hat{LF}_{t+1}$ is the forecasted load profile, $\hat{pw}_{t+1}$ is the estimated renewable forecasts, $SoC_t$ is battery state-of-charge, $gp_{t+1}$ is grid price for time $t+1$ , and $\pi_t$ captures contextual variables (price of the electricity, time-of-day, weather features, grid event flags). The control action is defined as

$$a_t = \{P_t^{grid}, P_t^{BESS}, P_t^{PV-W}, C_t^{inv}\}$$

where $P_t^{grid}$ is grid import/export, $P_t^{BESS}$ is battery charge/discharge power, $P_t^{PV-W}$ is power drawn from renewable sources (PV and Wind), and $C_t^{inv}$ denotes inverter-level control commands (e.g., curtailment or reactive support). The decision is computed under physical and operational constraints, including battery dynamics:

$$SoC_{t+1} = SoC_t + \eta_c P_t^{BESS} \Delta t$$

$$SoC_{t+1} = SoC_t - \frac{1}{\eta_d} \left| P_t^{BESS} \right| \Delta t$$

State of charge for the BESS should follow the bound $SoC^{\min} \le SoC_t \le SoC^{\max}$ and BESS power flow should follow $P^{BESS,\min} \le P_t^{BESS} \le P^{BESS,\max}$ and inverter and feeder constraints.

A canonical deterministic objective for is

$$\min \sum_{t=1}^{T} \left( gp_t P_t^{grid} \Delta t + \lambda_{\deg} UC_{\deg} P_t^{BESS} + \lambda_{cur} UC_{cur} P_t^{cur} \Delta t \right)$$

subject to power balance,

$$P_t^{PV-W} + P_t^{grid} + P_t^{BESS} = LF_t + P_t^{cur}$$

For resilience, the framework supports risk-aware extensions such as penalizing tail-risk of violations (e.g., load shedding) using CVaR:

$$\min \mathbb{E}\left[\mathrm{Cost}\right] + \lambda_{risk} CVaR_{\alpha}$$

This stage uses optimization tools, risk-scenario evaluation tools for stress testing under uncertainty, and constraint-validation tools to produce setpoints, schedules, feasibility certificates, and brief explanations (binding constraints and trade-offs).

### 2.3.3. Actuation execution

In the actuation stage, the execution agent translates the dispatch solution into device-level commands, such as inverter setpoints and battery charge/discharge rates, and then issues commands through the underlying controller interface. As microgrid actuation is safety-critical, this stage enforces validation gates before transmission, ensuring the planned actions remain feasible under real device limits, ramp-rate constraints, and current operating mode (grid-connected versus islanded). The executed command $C_t^{inv}$ is an implementation of the dispatch action $a_t$ that must satisfy the same hard constraints used in planning (e.g., $SoC$ bounds and inverter limits), and it must also pass runtime checks for controller availability, communication integrity, and equipment readiness. Tool support includes command translation (dispatch to device mapping), API tools (send commands and receive acknowledgements), and safety interlock tools that enforce hard constraints and activate safe fallbacks when validation fails, including compatibility with segmented or air-gapped control networks.

### 2.3.4. Monitoring and reflection

After actuation, the Monitor and Review agent closes the loop by comparing expected and observed system responses using telemetry and records outcomes for auditing and governance. Let $om_t$ denote observed measurements (power flows, voltage/frequency indicators, SoC, device status), and $\hat{om}_t$ denote the expected response implied by the dispatch plan; the monitoring step evaluates deviations through residuals such as $e_t = om_t - \hat{om}_t$ and detects anomalies when deviations exceed operational thresholds or persist over time. In addition to anomaly detection, the agent maintains a structured audit trail that stores inputs (forecasts and constraints), outputs (dispatch decisions and commands), timestamps, acknowledgements, and explanation summaries. If deviations occur, the Monitor & Review Agent flags the Operational Supervisor to trigger re-forecasting, re-optimization, or policy review.

## 3. Conclusion

This paper presented an agentic AI framework for microgrid control management to build resilient renewable Energy communities based on a hierarchical multi-agent architecture where an LLM-based orchestrator coordinates data search and forecasting, dispatch decision-making, safe execution, and monitoring as a workflow-level decision-maker rather than only a natural-language interface. Standardized tool interfaces (MCP-style abstractions) are used to separate reasoning from environment execution, improving interoperability and system robustness. The framework can be implemented using cloud services, knowledge bases, modular tool servers for forecasting and decision modules, data pipelines, orchestrator logic, and human-in-the-loop features for scenario evaluation and policy review. This direction reflects modern agentic system practice, where structured tool registries and staged planning are essential to maintain orchestration reliability as workflows scale. The framework can deliver practical advantages for renewable communities by enabling closed-loop resilience through automatic re-forecasting and re-optimization under disturbances, and by supporting policy-aware, human-in-the-loop governance for high-stakes decisions. In addition, the framework is modular and extensible; new agents can be added without redesigning the pipeline. Future work will focus on stronger uncertainty modeling (probabilistic forecasts and risk-aware dispatch), broader validation in representative microgrid scenarios, and systematic evaluation of orchestration reliability and safety under real operational disturbances.

## References

Ahmadi, M., Aly, H. and Gu, J., A comprehensive review of AI-driven approaches for smart grid stability and reliability, *Renewable and Sustainable Energy Reviews*, vol. 226, no. xx, p. 116424, 2026.

Bose, B.K., Artificial intelligence techniques in smart grid and renewable energy systems—Some example applications, *Proceedings of the IEEE*, vol. 105, no. 11, pp. 2262–2273, 2017.

Fathollahi, A., Machine learning and artificial intelligence techniques in smart grids stability analysis: A review, *Energies*, vol. 18, no. 13, p. 3431, 2025.

Glavic, M., (Deep) reinforcement learning for electric power system control and related problems: A short review and perspectives, *Annual Reviews in Control*, vol. 48, no. xx, pp. 22–35, 2019.
Izmirlioglu, Y., Pham, L., Son, T.C. and Pontelli, E., A survey of multi-agent systems for smartgrids, *Energies*, vol. 17, no. 15, p. 3620, 2024.
Mchirgui, N., Quadar, N., Kraiem, H. and Lakhssassi, A., The applications and challenges of digital twin technology in smart grids: A comprehensive review, *Applied Sciences*, vol. 14, no. 23, p. 10933, 2024.
Machlev, R., Heistrene, L., Perl, M., Levy, K.Y., Belikov, J., Mannor, S. and Levron, Y., Explainable Artificial Intelligence (XAI) techniques for energy and power systems: Review, challenges and opportunities, *Energy and AI*, vol. 9, no. xx, p. 100169, 2022.
Omitaomu, O.A. and Niu, H., Artificial intelligence techniques in smart grid: A survey, *Smart Cities*, vol. 4, no. 2, pp. 548–568, 2021.
Sifat, M.M.H., Choudhury, S.M., Das, S.K., Ahamed, M.H., Muyeen, S.M., Hasan, M.M., … and Das, P., Towards electric digital twin grid: Technology and framework review, *Energy and AI*, vol. 11, no. xx, p. 100213, 2023.

## Biographies

**Dr. Md Monirul Islam** is an Assistant Professor and Graduate Program Coordinator in the Department of Industrial Management and Technology at Texas A&M University-Kingsville. He earned his Ph.D. in Systems Engineering from Missouri University of Science and Technology (2020) and his B.S. in Industrial and Production Engineering from the Bangladesh University of Engineering and Technology (2014). His research focuses on mathematical optimization, complex systems modeling, systems simulation, and integrating AI/ML methods into decision-making frameworks.

**Dr. Nuri Yilmazer** is an Associate Professor and Graduate Program Coordinator in the Department of Electrical Engineering and Computer Science at Texas A&M University-Kingsville. He received the B.S. degree in Electrical and Electronics Engineering from Cukurova University, Adana, Turkey in 1996, and the M.S. and PhD degrees in Electrical and Computer Engineering from University of Florida and Syracuse University in 2000 and 2006 respectively. He worked as a Post Doctoral Research Associate in Computational Electromagnetics Laboratory at Syracuse University from 2006 to 2007. His current research interests include adaptive array processing, signal processing, and smart antennas.

**Dr. Abdallah Farraj** is an Associate Professor and Program Coordinator of Electrical Engineering at Texas A&M University–Texarkana (RELLIS Campus). He earned his B.Sc. and M.Sc. degrees in Electrical Engineering from the University of Jordan and his Ph.D. degree in Electrical Engineering from Texas A&M University. His research focuses on the security and resilience of critical infrastructure, IT/OT, IoT, and industrial control systems; modeling and analysis of cyber-physical systems; and the stability, integration, and optimization of smart power grids.

**Dr. Baisravan HomChaudhuri** is an Assistant Professor in Mechanical Engineering. He earned his Ph.D. in Mechanical Engineering and his M.S. in Mechanical Engineering from the University of Cincinnati, and his B.E. in Electrical Engineering from Jadavpur University (India). Prior to this position, he served as an Assistant Professor of Mechanical and Aerospace Engineering at the Illinois Institute of Technology and held postdoctoral appointments in the Department of Electrical and Computer Engineering at the University of New Mexico and at the Clemson University International Center for Automotive Research. His research focuses on distributed optimization and control, model predictive control, stochastic and robust MPC, estimation methods, motion planning, reachability analysis, robotics, and connected and automated vehicle systems.