FORCES
FOUNDATIONS OF RESILIENT
CYBER-PHYSICAL SYSTEMS

Welcome to the FORCES newsletter

Summer 2017

Summer at the FORCES project provides opportunities for research and collaboration. This issue of the newsletter explores how FORCES research has evolved, including new projects that have arisen from the initial work of FORCES partner universities as well as current applications in industry environments. Lillian Ratliff, who was previously included in the FORCES project at Berkeley and is now at the University of Washington, describes how her work is continuing and expanding at her home institution--while collaborating with current FORCES members. Another article from a group at Vanderbilt outlines their development of a hardware-in-the-loop (HIL) testbed and its application to communication networks. We also learn about a study at Berkeley that examines how routing applications used on a smartphone impact urban traffic, while researchers at MIT offer details about aircraft fuel burn and its effect on aircraft performance. All these projects provide insight into how research being conducted through the FORCES project has impact beyond just the partner universities.

Thanks for taking time to read the FORCES summer newsletter, and many thanks to the National Science Foundation for its generous support of this work.

As always, if you have feedback, comments, and suggestions I'd appreciate hearing from you.

Sincerely,

S. Shankar Sastry
Professor and Dean of Engineering
University of California, Berkeley

RESEARCH SPOTLIGHT

Urban Traffic and Routing Applications

by Jérôme Thai (University of California, Berkeley)

A significant number of drivers now use routing apps such as Google Maps, Waze, INRIX, or Apple Maps. This widespread adoption is fueled by the increasing penetration of smartphones, the rapid expansion of Mobility-as-a-Service systems such as Uber and Lyft, and in the future, will be used by unmanned vehicles following shortest path algorithms. Even though these tools optimize route choices and decrease travel time, their impact on road traffic and urban congestion is not well-studied and understood. Cities bordering major highways in the U.S. have noticed an increase of traffic demand on their networks, presumably due to application users leaving highways to avoid congestion. This alleged flow transfer is a challenge for public policy. When an affected city's infrastructure, mostly financed by and for local taxpayers, experiences increased traffic demand, new policy must be developed citywide to address the problem. In collaboration with scientists at Lawrence Berkeley Lab, we are developing a game-theoretical framework to describe heterogeneous traffic in which a varying percentage of drivers use the navigation apps. Using the Los Angeles road network from OpenStreetMap composed of 14,617 nodes, 28,376 links (*Figure 1*) and 99,097 Origin-Destination demand pairs from Southern California Association of Governments, we applied a Frank Wolfe algorithm parallelized on 5 NERSC computing nodes.

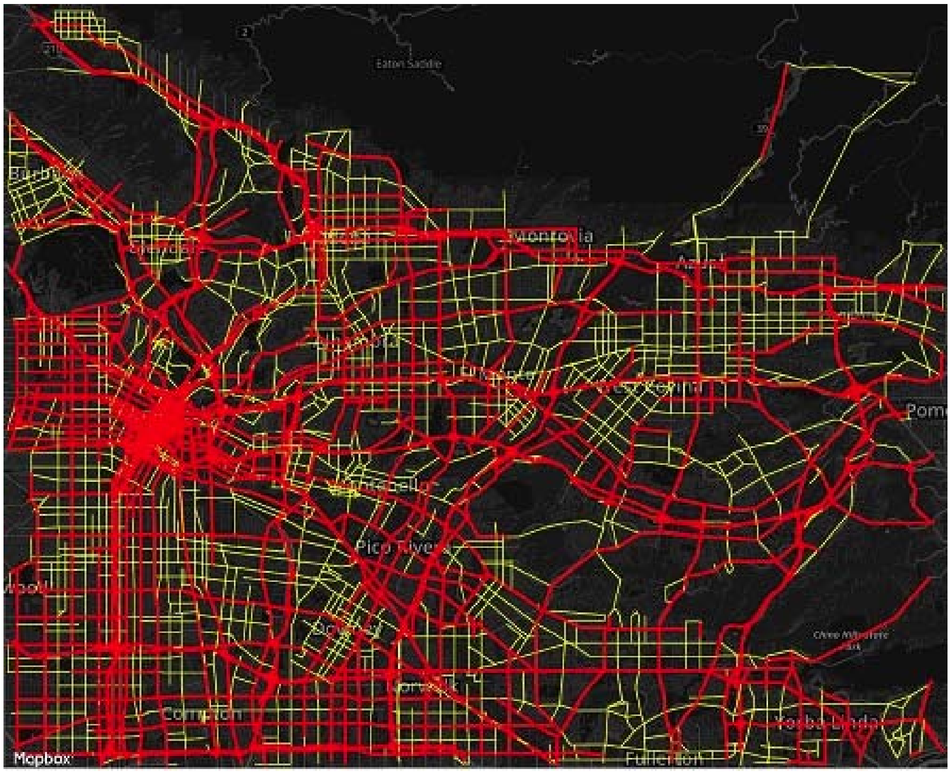

Figure 1 - Los Angeles road network

Preliminary results show app-based routing can potentially increase the Vehicle-miles Traveled (VMT) on local roads threefold, while there is only a 10% decrease in VMT on high-capacity roads (*Figure 2*).

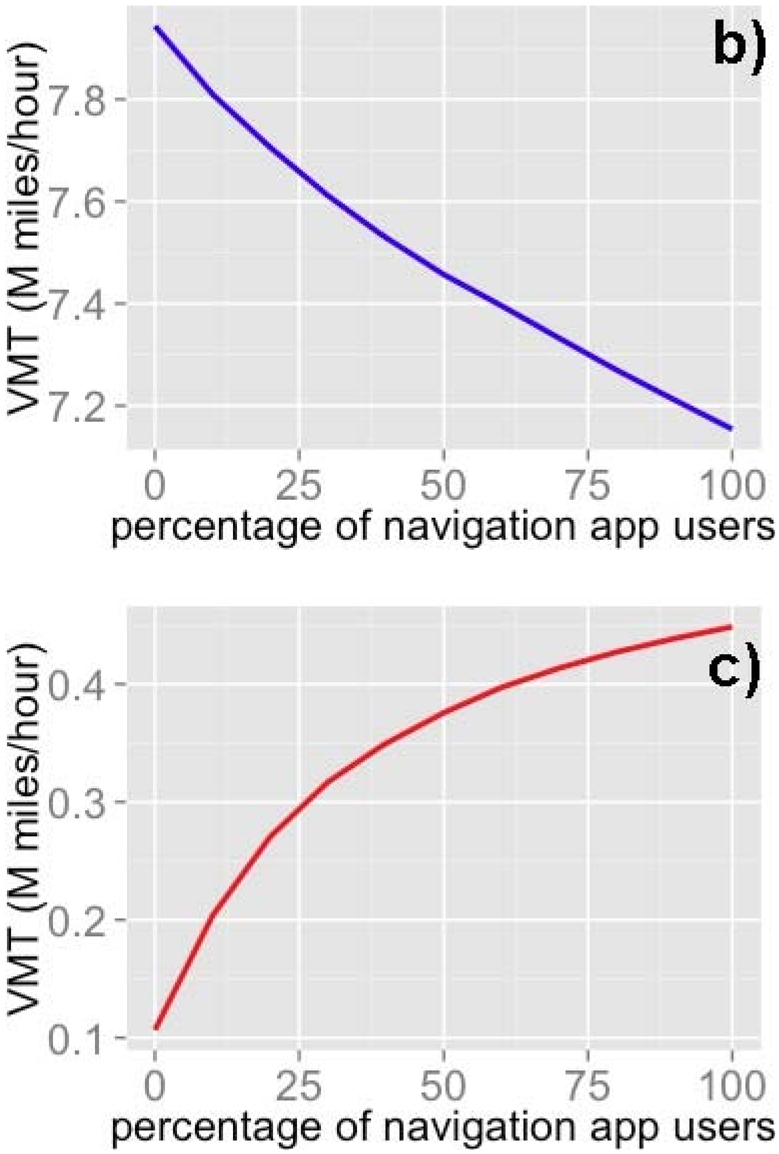

Figure 2 - Vehicle miles traveled comparison

Design and Analysis of Secure CPS Utilizing a Hardware-in-the-Loop Testbed

by Bradley Potteiger (Vanderbilt University), William Emfinger (Vanderbilt University), Gabor Karsai (Vanderbilt University), and Xenofon Koutsoukos (Vanderbilt University)

Cyber-physical systems (CPS) exist in a wide variety of critical infrastructure. While applications have traditionally only been susceptible to physical attacks, the introduction of communication networks between sensors, actuators, and computational platforms, as well as the synchronization of data with the cloud, introduces a new layer of cyber domain threats. Cyber-attacks have been demonstrated in applications including medical devices, vehicles, and national defense infrastructure. With the future shift to smart cities, autonomous vehicles, and remote patient monitoring, the inclusion of cybersecurity principles in the design and development of CPS is of critical importance.

Our work focuses on two areas: developing a hardware-in-the-loop (HIL) testbed for implementing cybersecurity experiments on distributed CPS [1], and looking at the effects of moving target defenses on the protection and operation of CPS. When conducting CPS cybersecurity analysis through simulations, it is hard to realize the effects of system dependent attacks such as distributed denial of service (DDOS) and code injection attacks. To solve this problem, we developed an HIL testbed from a cluster of embedded computing nodes to emulate CPS controllers and communications consistent to deployed devices in the field. Additionally, the testbed is capable of emulating various communication protocols in the network. By developing experiments on the HIL testbed, more reliable predictions can be made about expected deployment reliability.

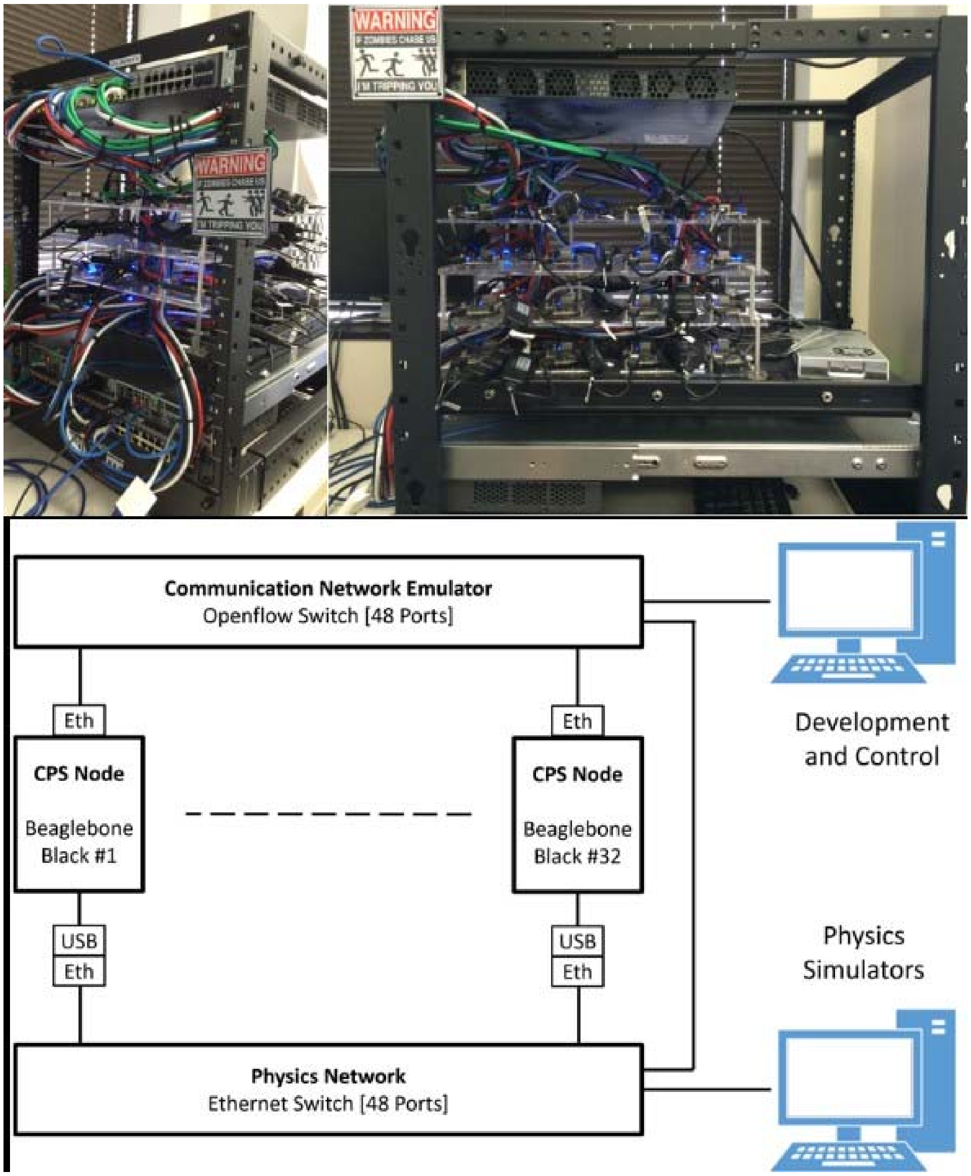

Figure 1 - CPS Security HIL Testbed

Our group has also developed a software development framework called ROSMOD in support of the HIL testbed for rapidly deploying CPS security experiments. ROSMOD is a graphical model integrated computing (MIC) tool that utilizes component-based design principles with the Robot Operating System (ROS) middleware for representing the distributed nature of a CPS. By using this tool, component-specific code and communication networks can be implemented to build up a complex model of a CPS. ROSMOD is also able to integrate component and network cyber-attacks to aid in the comprehensive cybersecurity analysis of a system. We have used ROSMOD to develop cybersecurity experiments in traffic, railway, autonomous vehicle, and satellite applications.

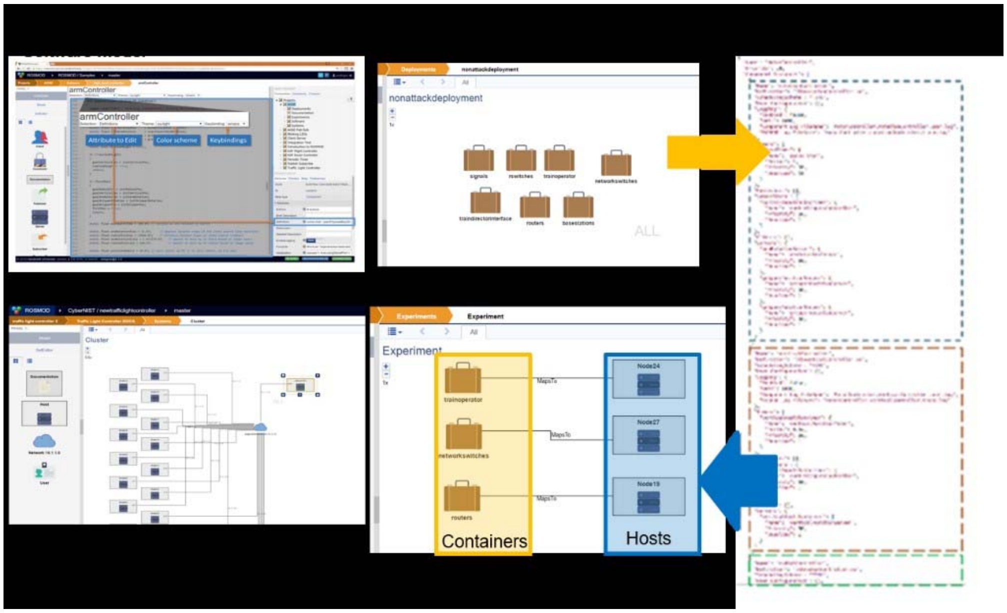

Figure 2 - ROSMOD Tool Graphics

The last part of our work focuses on analyzing the effects of moving target defenses on preventing cyber-attacks in CPS. In the autonomous vehicle and connected car domain, one large threat vector includes code injection attacks through buffer overflow vulnerabilities. We proposed a framework based on instruction set randomization (ISR) techniques to prevent the execution of injected payloads on a system. Additionally, a recovery mechanism was integrated into the framework that leverages diversity principles by switching between controller versions in the event of attacks. An autonomous vehicle case study was implemented utilizing this framework leveraging the operation of a neural network and a GPS waypoint controller.

[1] B. Potteiger, W. Emfinger, H. Neema, X. Kotsoukos, C. Tang, and K. Stouffer. Evaluating the Effects of Cyber-Attacks on Cyber Physical Systems using a Hardware-in-the-Loop Simulation Testbed. In Resilience Week (RWS), 2017, IEEE, Wilmington, DE, September 18, 2017

Bi-Directional Engagement: Beyond FORCES

by Lillian Ratliff (University of Washington)

Infrastructure systems are continuously integrating new technologies that enable people, companies, and service providers to be at once consumers and producers of resources which may be physical- or information-based. This bi-directional engagement of participants has unveiled interesting challenges and new research directions, many of which have been identified or have evolved out of FORCES fundamental research. Currently, FORCES alumni Lillian Ratliff (University of Washington) and Roy Dong (University of California, Berkeley) are working on several research projects that include graduate students Tanner Fiez (University of Washington), Eric Mazumdar (University of California, Berkeley), and Tyler Westenbroek (University of California, Berkeley), as well as faculty member Shankar Sastry (University of California, Berkeley). This article offers an overview of some problems arising from these bi-directional exchanges.

At the center of many of the research projects explored by the group is the fact that now prolific smart sensors, ranging from personal devices to more traditional purpose-built sensors, may be owned by a multitude of sources, and can produce qualitatively different data streams which can be combined to make inferences. This is partly driven by the observance that many companies and service providers now solicit data from users in exchange for quality of service or monetary incentives. We are working on developing learning algorithms to model and understand dynamic decision-making under uncertainty when agents (e.g., human participants) are risk-sensitive and learning about exogenous and endogenous uncertainties. In addition, we are developing new data market (i.e., markets in which data is a commodity) models where competition arises between market participants at both the service provider or firm level as well as the user or data source level. Our recent work investigates scenarios in which strategic firms solicit data from a common pool of data sources in order to perform an estimation task. We show that inference quality depends on the level of competition and that free-riding can be a significant issue which results in a loss of social welfare. This motivates regulation and policy to improve market efficiency. As part of our ongoing work, we are developing a dynamic equilibrium model where firms adaptively learn, e.g., demand and design, incentives to encourage market participation by data sources. We hope to understand how competition and data quality impacts decisions such as market research versus technology/production investments while also developing new metrics for socio-economic impact that are relevant to the application domain.

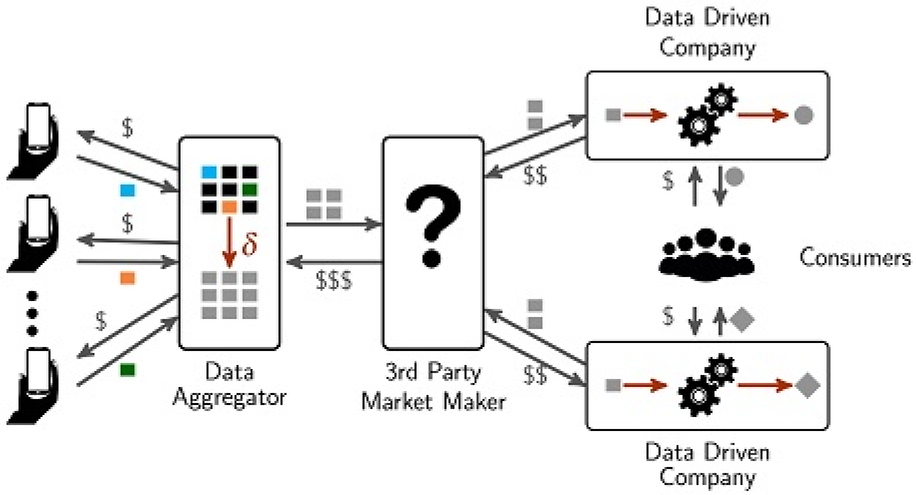

Data Market Model

These efforts leverage real data streams from the City of Seattle as well as existing open data sets such as the NYC Taxi data set. In our work with the Seattle Department of Transportation, Lillian Ratliff and Tanner Fiez are conducting in-situ experiments to adjust transportation policies in the U-District near the UW campus and understand the impact of new ride-sharing services and infrastructure investments (e.g., light-rail expansion to the area).

References

E. Mazumdar, L. Ratliff, T. Fiez, S. Sastry. Gradient-Based Risk-Sensitive Inverse Reinforcement Learning with Applications. Submitted to IEEE CDC, 2017.

T. Westenbroek, R. Dong, L. Ratliff, S. Sastry. Statistical Estimation with Strategic Data Sources in Competitive Settings. Sumbmitted to IEEE CDC, 2017.

Statistical Modeling of Aircraft Fuel Burn

by Yashovardhan Chati (Massachusetts Institute of Technology); Hamsa Balakrishnan (Massachusetts Institute of Technology)

Fuel burn is a key driver of aircraft performance, and contributes to airline costs and emissions. Over the next two decades, passenger enplanements are forecast to grow by 2.3% per year. In order to achieve the FAA's NextGen program objective of sustained aviation growth, we need fuel burn models that can assess the environmental impact of current and future traffic. In recent work under FORCES, we have leveraged techniques from identifying models for cyber-physical systems to develop statistical models of a flight's fuel burn, given observations of its physical position (i.e., its trajectory). Such models provide both mean estimates of the fuel consumed and also quantify the underlying uncertainty, a marked improvement over current aircraft performance models that provide point estimates of an inherently stochastic quantity. We have also leveraged a physical understanding of aircraft and engine dynamics to conduct more effective feature selection.

We have applied our modeling approach, based on machine learning techniques such as Gaussian Process Regression (GPR), to several aspects of aircraft performance. Under an Airport Cooperative Research Program (ACRP) Graduate Research Award supported by the Transportation Research Board (TRB) of the National Academies, we have developed data-driven models for the low-altitude (climb out and approach) phases of flight. Our statistical models of fuel burn are found to be over 74% more accurate in climb out, and over 61% more accurate in approach, when compared to current state-of-the-practice aircraft performance models. We are also applying these methodologies to develop models of airport surface fuel burn, in an FAA-funded collaboration with the MIT Lincoln Laboratory. Finally, our FORCES-supported research on estimating the takeoff mass of an aircraft from surveillance data [1}, received the Best Paper Award in “Trajectory Prediction” at the Air Traffic Management R&D Seminar (ATM-2017), co-organized by the FAA and Eurocontrol, which was held in Seattle, WA in June 2017.

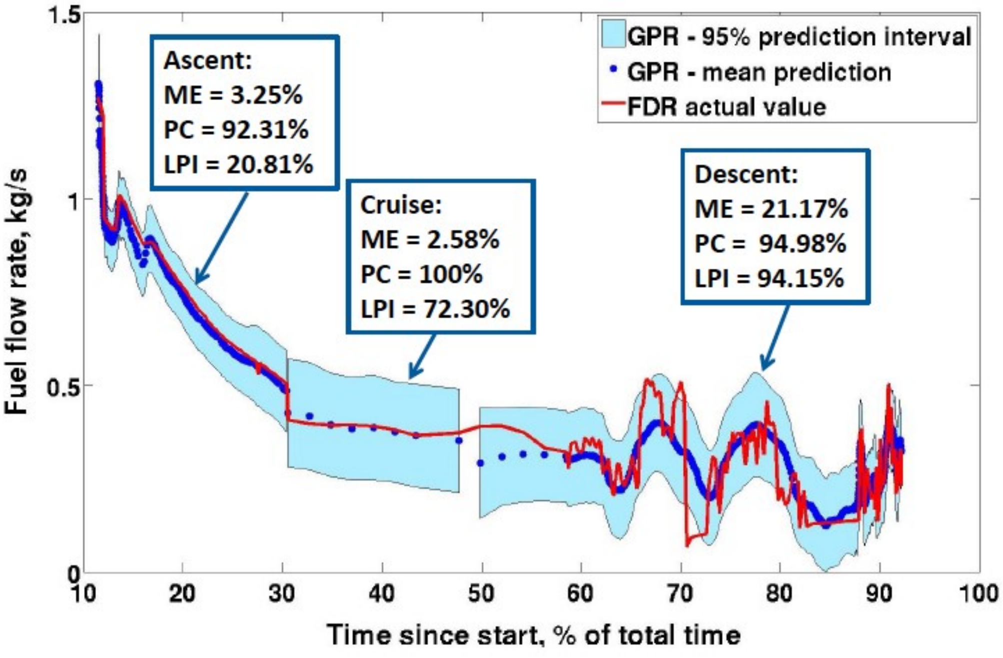

Figure 1. Example of Gaussian Process Regression (GPR)-based model predictions of the fuel flow rate in

different phases of flight, for a trajectory flown by an A321-111 aircraft. ME denotes the mean absolute error of the model, PC denotes the percentage of observations that fall within the 95% prediction interval, and LPI denotes the normalized length of the prediction interval (i.e., the length of the PC divided by the mean fuel flow rate). The red line shows the ground truth, namely the actual fuel flow rate reported by the Flight Data Recorder (FDR).

[1] Y.S. Chati and H. Balakrishnan. Statistical Modeling of Aircraft Takeoff Weight, USA/Europe Air Traffic Management Research and Development Seminar, June 2017.

PROJECT NEWS

Upcoming Events

FORCES All Hands Meeting
August 23-24, 2017
University of California, Berkeley