RRoCCET21

Research Running on Cloud Compute & Emerging Technologies

Held virtually on August 10th - 12th, 2021
from 9 AM - 1:30 PM PST

Recordings:
https://cloudbank-project.github.io/cb-resources/rroccet/videos-2021/

This conference is brought to you by CloudBank, an organization whose mission is to facilitate and accelerate research computing on the cloud. It is a collaboration between UC San Diego, the University of Washington and UC Berkeley together with Strategic Blue. CloudBank is sponsored by and originated from the National Science Foundation Computer and Information Science and Engineering (CISE) Directorate.

The purpose of this conference is to inspire you to consider utilizing the cloud in your research. Presenters will share how cloud adoption in their domain enabled them to move their research computing ahead. We hope the proceedings which follow give you an idea of what is possible, and can act as a "recipe book" for mapping powerful computational resources onto your own field of inquiry. We also hope you will take away an appreciation of CloudBank as a cloud access facilitation program. For more information on CloudBank please check out our website and reach out to us at help@cloudbank.org.

The CloudBank Philosophy

CloudBank is an academic program supporting academic research. We are "**guides** not **ticket-takers**", viewing cloud adoption as a journey and a process. We're here to help you with that when the cloud makes sense for your research program. In practice, NSF has begun to offer cloud support to awardees and CloudBank manages this process. This includes:

- **Project Arc**
  We will work with you and with the NSF to optimize your spend, deal with post-project

requirements (Data Management Plans and more), and close gaps between planned spend and actual spend. Should you over- or under-estimate your cloud spend in the NSF grant process: That's fine; CloudBank will work with you and NSF to accommodate your project needs.

- **Training / Building**
  Our cloud training approach is based on a research-oriented perspective. We emphasize the core skills you need to work effectively on the cloud. (We know you likely do not have time to develop a vast cloud engineering skill set.) To this end, we provide:
  - Live support including direct consulting, webinars and events like RRoCCET21
  - Curated learning material from the best available resources
  - CloudBank Solutions and Protocols: Reference materials for cost optimization, design/build patterns for common data science tasks, plans for anticipating the future direction and needs of your research program, case studies, example projects, community forums and more
  - Facilitating connection with and support from cloud vendors, specifically AWS, Azure, Google Cloud and IBM Cloud.

- **Financial Management**
  The CloudBank partnership includes Strategic Blue, a cloud brokerage company that helps us stretch cloud dollars and provide participating researchers with tools to access and manage cloud accounts. CloudBank works with NSF, academic administrations and cloud vendors to minimize cost in other ways as well. For example we do not levy overhead or indirect fees on cloud spend.

Concerning other cloud adoption paths
CloudBank's learning resources are openly available. We encourage using them to understand the utility of cloud computing.

Concerning Students
CloudBank recognizes the student community as an essential part of the academic research ecosystem. In view of the rapid expansion of data science, CloudBank supports collegial instances of Berkeley's ***Data 8*** Foundations of Data Science course. Placing course infrastructure on the cloud means the standardized course environment can be made available to hundreds of students. We also note that students often adapt readily to cloud technology; and can actively contribute to research program migration to the cloud.

The ***RRoCCET21 conference broadsheet*** will be sent to attendees via email. It contains some useful links because (a) we hope the conference inspires you to look into the cloud and (b) we intend the broadsheet to be your springboard to the logical continuation via hands-on activities.

Wildfire modeling with TPUs

Dr. Qing Wang and Dr. Yi-fan Chen

Author Bio:

Dr. Qing Wang joined Google as a software engineer after he received his PhD degree from Stanford University in 2018. His current research focuses on developing computational fluid dynamics libraries based on the tensor processing unit (TPU) infrastructure, which is used to simulate geophysical flow related applications such as wildfires and clouds. Before joining Google, he worked on model development for large-eddy simulations with focuses on multiphase and combustion related problems. Dr. Yi-fan Chen works as a software engineer at Google, leading a team on developing high performance scientific computing libraries and applications. His previous work at Google includes development of recommendation systems and other Google infrastructures. Before joining Google, he worked at Intel and ASML. He got his PhD degree in 2005 in applied physics at Cornell University.

Abstract:

Increasing rates of catastrophic wildfires has become a major concern for the economy and environment globally. Nearly 8 billion tons of CO2 is released by wildfires every year, which accounts for ~10% of the global CO2 emissions. Additionally, wildfires cause an annual loss of 220 billion dollars and 775 thousand residences in the US. It has been shown that high fidelity numerical simulations are useful for investigating and predicting wildfires. These simulations provide insights for fire risk assessment and mitigation, fire management and firefighting strategy optimization, and land-use and wildland-urban interface planning. These simulations are also potentially useful for generating synthetic data, which can be used in developing improved fire prediction models (including ML based ones). We present a computational fluid dynamics (CFD) simulation framework based on clusters with Tensor Processing Units (TPU). In these clusters, up to 1024 chips are connected through a dedicated high-speed,low-latency, two-dimensional toroidal inter core interconnect (ICI) network, which provides higher computational efficiency than a CPU based high performance computing framework.The simulation framework solves the 3D Navier-Stokes equation along with constitutive models for combustion, pyrolysis, heat-transfer, and other thermochemical processes. With this framework, we have been able to produce a high fidelity wildfire dataset based on the FireFlux II experimental setup. This dataset considers a broad range of physical factors including wind speed, terrains, fuel density, and moisture. Additionally, we have simulated the full event of the

Tubbs fire that happened historically in California in 2017. These high fidelity simulations are generated by the TPU simulation framework at a significantly lower cost than by CPU based CFD softwares, which provides us a foundation for studying the physical insights of wildfire propagations scientifically.

Case Study Summary:

The scientific problem:

We want to investigate wildfire propagation at the unprecedented resolution, potentially coupling the local cloud etc weather phonenoma

We implemented physical models that are based on first-principles by solving governing conservation equations for mass, momentum, energy, species transport in conjunction with constitutive models for combustion, pyrolysis, heat-transfer, and other thermochemical processes. These models solve for the 3D wildfire behavior that is two-way coupled to the atmospheric flow field dynamics.

Based on this simulation framework, we performed a large scale simulation of the Tubbs fire, which happened historically in California, 2017. This simulation represents a domain of 200 thousand acres mountainous terrain with a mesh of more than 800 million degrees of freedom. The full duration of the fire is simulated with 2.5 days of runtime.

The computational methods:

These are finite difference large scale distributed simulation of CFD + thermodynamics + combustion modeling. We are also exploring using ML + non-ML approaches together in understanding the problem.

The set of governing equations is solved on a Cartesian coordinate system employing an equidistant mesh in each direction. All dependent variables are stored on a collocated mesh, and spatial derivatives are discretized using a finite-difference scheme. We use Tensorflow to program the described computational methods on TPU.

The cloud resources:

TPU clusters which provide great scalability through the dedicated high bandwidth low-latency inter-chip connections up to 1024 chips.

Our work is conducted on clusters with Tensor Processing Units (TPU). We are using Google internal clusters at present for the convenience of development. The framework we built is transferable to Google cloud.

TPU are application-specific integrated circuits that have been specifically designed for machine-learning applications. A TPU v3, considered in this work, consists of a TPU board with four independent chips. Each chip has two tensor compute cores that are optimized for vectorized operations. Up to 1024 chips are connected through a dedicated high-speed, low-latency, two-dimensional toroidal inter core interconnect (ICI) network, forming a TPU pod.

Each TPU core has 16 GiB on-chip high-bandwidth memory (HBM), and consists of a scalar, vector and two matrix units (MXU), each MXU supporting 128×128 multiply-accumulate (MAC) operations per cycle. Each MXU has a raw peak throughput of 22.5 teraflops.

The TPU board is connected through the PCIe-bus to the host server, which sends the instruction to the TPU for execution. The ICI network directly connects the cores within a TPU pod, so that the communication between TPU cores go directly through this dedicated network without involvement of PCIe-bus and host.

For further information:
https://research.google/research-areas/general-science/

Towards Facilitating Empathic Conversations in Online Mental Health Support

Ashish Sharma

Author Bio:
Ashish Sharma is a PhD student at the Paul G. Allen School of Computer Science & Engineering, University of Washington. His research focuses on building computational techniques for encouraging empathic conversations in text-based mental health support. Ashish holds a dual degree (B.Tech. + M.Tech.) in computer science from the Indian Institute of Technology, Kharagpur and has worked as a Research Fellow at Microsoft Research, India.

Abstract:
Online peer-to-peer support platforms such as TalkLife enable conversations between millions of people who seek and provide mental health support. If successful, web-based mental health conversations could improve access to treatment and reduce the global disease burden. Psychologists have repeatedly demonstrated that empathy, the ability to understand and feel the emotions and experiences of others, is a key component leading to positive outcomes in supportive conversations. Computational methods that can empower non-expert peer-supporters with empathy-based feedback and training has the potential to help them express higher levels of empathy and in turn, improve the effectiveness of online support platforms.

In this talk, I will present our work on understanding and improving empathy in online mental health support conversations. First, we design a novel framework and dataset of empathic conversations and develop a RoBERTa-based bi-encoder model for identifying empathy in conversations. Using this model, we demonstrate that highly empathic conversations are rare on online mental health platforms. Therefore, to improve empathy in conversations, we introduce Empathic Rewriting, a new task that aims to transform low-empathy conversational posts to higher empathy. We propose PARTNER, a deep reinforcement learning agent that learns to make sentence-level edits to conversations in order to increase the expressed level of empathy while maintaining conversation quality through fluency, specificity, and diversity.

Case Study Summary:

The scientific problem we tackled:
Access to mental health care is a global challenge that online peer support platforms like TalkLife can help mitigate. Millions of people seek and provide support online but they struggle to have effective conversations. We developed computational methods to help

facilitate empathic interactions in peer support platforms through intelligent and actionable feedback. This research helped us understand the role of empathy in peer support and how to empower peer supporters in writing more empathic responses.

The computational methods we used:
We trained Natural Language Processing and Reinforcement Learning models for identifying and improving empathy in peer support conversations. Using these models, we have designed interactive tools for providing intelligent and actionable real-time feedback to peer supporters and for suggesting concrete ways in which the conversations can be made more empathic.

The cloud resources we used:
We used Azure GPUs for training our models and deploying them in the real world.

The differences we've observed between locally-provided and cloud-provided resources:
Cloud-provided resources provide significantly more flexibility and scale than locally-provided resources. Training large AI models requires GPUs with varying memory needs. Azure facilitates instantiating different types of GPUs. Also, for the models to optimally learn these complex tasks, we needed to perform a hyperparameter tuning using grid search. This requires training the same model several times. While this is infeasible on a local machine, we can scale this hyperparameter search by creating multiple Azure GPUs.

For further information:
http://bdata.uw.edu/empathy/

COVID-19 Genomic Analysis

Dr. Niema Moshiri

Author Bio:
Niema Moshiri is an Assistant Teaching Professor in the Computer Science & Engineering Department at the University of California, San Diego (UCSD). His research interests lie in computational biology, with a research focus on viral phylogenetics and epidemiology. He also places a heavy emphasis on teaching, namely on the development of online educational content, primarily Massive Adaptive Interactive Texts (MAITs).

Abstract:
Throughout the COVID-19 pandemic, UC San Diego has implemented a "Return to Learn" Initiative (https://returntolearn.ucsd.edu/), which consists of mass testing and sequencing of SARS-CoV-2 samples across the county. Importantly, this large-scale effort is yielding massive amounts of raw viral amplicon sequencing data (~10,000 samples per day) that need to be analyzed in real-time, posing a significant challenge with regard to computational infrastructure. In response to this challenge of scalability, I have developed a novel tool called ViReflow (https://github.com/niemasd/ViReflow) that leverages the Reflow system of incremental data processing developed by GRAIL (https://github.com/grailbio/reflow) in order to implement an elastic, massively-parallelizable, and massively-scalable (e.g. country-wide scale) pipeline for executing standard amplicon sequence data → variant call + consensus sequence workflows.

Case Study Summary:

The scientific problem we tackled:
Throughout the COVID-19 pandemic, UC San Diego has implemented a "Return to Learn" Initiative (https://returntolearn.ucsd.edu/), which consists of mass testing and sequencing of SARS-CoV-2 samples across the county.

The computational methods we used:
This large-scale effort is yielding massive amounts of raw viral amplicon sequencing data (~10,000 samples per day) that need to be analyzed in real-time, posing a significant challenge with regard to computational infrastructure. In response to this challenge of scalability, I have developed a novel tool called ViReflow (https://github.com/niemasd/ViReflow) that leverages the Reflow system of incremental data processing developed by GRAIL (https://github.com/grailbio/reflow) in order to implement an elastic, massively-parallelizable, and massively-scalable (e.g. country-wide scale) pipeline for executing standard amplicon sequence data → variant call + consensus sequence workflows.

The cloud resources we used:
We are using Amazon Web Services (AWS) Elastic Compute Cloud (EC2) for conducting the analyses, and we are using AWS Simple Storage Service (S3) for data storage.

The differences we've observed between locally-provided and cloud-provided resources:
The key benefit between locally-provided and cloud-provided resources for our purposes is the ability to scale the computational resources we utilize to precisely match what is needed for a specific sequencing run. It doesn't matter if we have 10 samples vs. 1000 samples: the execution remains the same, and the allocated AWS resources dynamically scale, meaning we experience reduced walltime due to massively-parallelized execution while maintaining low compute costs.

For further information:
https://niema.net/
https://github.com/niemasd/ViReflow
https://github.com/niemasd/ViReflow-Paper/blob/main/scripts/Figures.ipynb

Data storage for neuroscience at a massive scale

Dr. Satrajit Ghosh

Author Bio:
Satrajit Ghosh is a Principal Research Scientist at the McGovern Institute for Brain Research at MIT and an Assistant Professor of Otolaryngology at Harvard Medical School. He directs the Senseable Intelligence Group whose research portfolio comprises projects on spoken communication, brain imaging, and informatics to address gaps in scientific knowledge in three areas: the neural basis and translational applications of human spoken communication, machine learning approaches to precision psychiatry and medicine, and preserving information for reproducible research.

Abstract:
With the advent of the brain Initiatives in the US and other countries, the scale and diversity of neuroscience data continues to grow. DANDI (https://dandiarchive.org/) is a US BRAIN Initiative supported data archive and collaboration platform that is housed using several cloud services. The intent of the archive is to make data accessible to others, but also to transform the way people think about computing and data in labs. The platform provides access to cellular timeseries and imaging data that ranges from MBs to TBs per file, and growing. This diversity challenges all of us to rethink where and how we should store and compute data. In this talk, I will briefly cover the infrastructural components, the use of different cloud resources, and the technical developments needed to standardize data for search, and computation, and to provide easy access to pieces of data. I will also present a set of social and technical challenges that will need to be tackled by the neuroscience community and other disciplines as the field evolves to co-design and generate scientific output around the planet.

For further information:
https://dandiarchive.org/

SMRF: a Cloud-Based Social Media Research Framework

Dr. D. Hudson Smith

Author Bio:
Dr. Hudson Smith has a Ph.D. in physics from the Ohio State University. His current research interests include a focus on Natural Language Processing techniques for the analysis of large language corpora in the social sciences.

Abstract:
Clemson University's Watt Artificial Intelligence program has supported a wide range of research projects involving the collection and analysis of social media data. Despite the diverse application areas including political science, brand management, and cybersecurity, these projects have each involved a standard set of steps on the road from concept formation to the generation of research outputs. Though there are many tools available for specific components of the social media analytics research process, integration of these tools remains an extreme challenge for many researchers in the social sciences. To date, the Watt AI program has developed bespoke solutions to meet the needs of specific researchers. This approach does not scale well beyond a small number of engagements, limits the potential for independent activity on the part of domain experts, and increases the burden of long-term project maintenance. We designed the Social Media Research Framework (SMRF) to be a reusable, low-maintenance toolset for automating many of the rote activities encountered in social media analytics research such as data collection, data labeling, model deployment, and model monitoring. In this presentation, we discuss the challenges of software-intensive research in the social sciences, describe the thinking behind the design of SMRF, and share our researcher-friendly cloud-deployment model using IBM Cloud CodeEngine. Finally, we present a case study applying SMRF to the analysis of all Tweets from members of the US House of Representatives.

Case Study Summary:

The scientific problem we tackled:
We would like to better understand the behavior patterns of US political elites on social media and the impact that those behaviors have on others using social media platforms. For example, we would like to know the extent to which the use of polarizing language by politicians teaches other social media users to resort to polarizing language.

The computational methods we used:
Large corpora of social media messages are collected and analyzed on an ongoing basis. Types of analysis include topic and sentiment analysis of individual messages

using natural language processing (NLP) methods and graphical analysis of the relationships between social media users and messages.

The cloud resources we used:
We make use of generic cloud-based services like managed SQL databases, cloud object storage, and virtual servers for running applications that continuously collect and analyze data. We use IBM Cloud's serverless utility (IBM Cloud Functions) to run periodic data management jobs. We have also made use of IBM Watson's NLP capabilities as part of our content analysis of social media messages. We currently process around 2,000 messages per day with plans to scale this 10-fold in the near future.

For further information:
http://www.clemson.edu/cecas/departments/ece/faculty_staff/faculty/hsmith.html

Analysis Ready Data in the Cloud

Charles Stern

Author Bio:
Charles is a Data Infrastructure Engineer in the Ocean Transport Group at Lamont-Doherty Earth Observatory. His work focuses on Pangeo Forge, an open source tool for data Extraction, Transformation, and Loading (ETL). The goal of Pangeo Forge is to make it easy to extract data from traditional data repositories and deposit it in cloud object storage in analysis-ready, cloud-optimized (ARCO) format. He is endlessly curious about elegant, open-source tools that help us understand our changing planet.

Abstract:
Analysis-ready, cloud-optimized (ARCO) data formats allow us to leverage elastic scaling and distributed parallel processing to dramatically accelerate scientific discovery. In this live demo, we will implement these principles to analyze and visualize a massive sea surface height dataset from the Copernicus Marine Environment Monitoring Service (CMEMS). The presentation will feature Pangeo's open-source software stack including Jupyter, Xarray, Dask, and Zarr and will conclude with a look at Pangeo Forge, a new open-source toolkit for transforming your own data from archival formats into ARCO data stores.

For further information:
https://pangeo-forge.readthedocs.io/en/latest/

Privacy-preserving machine learning

Dr. Martine De Cock

Author Bio:
Martine De Cock is a Professor at the School of Engineering and Technology, University of Washington–Tacoma, USA and a guest professor at Ghent University, Belgium. She has over 190 peer-reviewed publications in international journals and conferences on artificial intelligence, machine learning, information retrieval, web intelligence, and logic programming. Her current research interests are privacy-preserving machine learning (PPML) and machine learning for cybersecurity. She holds a patent on cryptographically secure machine learning, and her team won Track IV of iDASH2019, the most significant competition in privacy-preserving analysis of genomic data in the world.

Abstract:
Machine learning-based applications provide a lot of convenience but rely on our personal data to work well. Can we get the benefits of using them without disclosing sensitive information about ourselves? We present cryptographic protocols based on Secure Multiparty Computation (SMC) that let machine learning-backed applications provide their services without demanding our privacy in return. The protocols run on Microsoft Azure virtual machines in a machine-learning-as-a-service setup while keeping the data as well as the model parameters encrypted. We present use cases of privacy-preserving classification of personal text messages and privacy-preserving recognition of people's emotions in videos, both with state-of-the-art deep learning models.

Case Study Summary:

The scientific problem we tackled:
Machine learning-based applications provide a lot of convenience but rely on our personal data to work well. Can we get the benefits of using them without disclosing sensitive information about ourselves?

The computational methods we used:
We use cryptographic protocols based on Secure Multiparty Computation (SMC) that let machine learning-backed applications provide their services without demanding our privacy in return. Use cases include privacy-preserving classification of personal text messages and privacy-preserving recognition of people's emotions in videos, both with state-of-the-art deep learning models.

The cloud resources we used:
The protocols run on Microsoft Azure virtual machines in a machine-learning-as-a-service setup while keeping the data as well as the model parameters encrypted.

The differences we've observed between locally-provided and cloud-provided resources:
Benchmarking our protocols in the cloud is important in the research group and in the broader research community. It enables reproducibility of research results, and consistent apples-to-apples comparisons of the efficiency and scalability of different cryptographic protocols on the same kind of hardware.

For further information:
http://faculty.washington.edu/mdecock/

Generative Models for High Energy Physics simulation

Dr. Sofia Vallecorsa

Author Bio:
PhD in Physics, Sofia Vallecorsa is a CERN physicist and the AI and Quantum research lead at CERN openlab. Dr. Vallecorsa has large experience as software developer in the field of High Energy Physics with significant accumulated expertise across the full research chain from real data analysis to simulation workloads. Extensive background on Quantum Machine Learning and classical Deep Learning architectures, frameworks, and methods for distributed training and hyper-parameters optimization, on different environments ranging from commercial clouds to HPC. Dr Vallecorsa is today the coordinator for Quantum Computing within the CERN Quantum Technology Initiative.

Abstract:
With the increasing number of Machine and Deep Learning applications in High Energy Physics, easy access to dedicated IaaS represents a requirement for fast and efficient R&D. This work explores different types of cloud services to train a Generative Adversarial Network (GAN) in a parallel environment, using a Tensorflow data parallel strategy. More specifically, we parallelize the training process on multiple GPUs and Google Tensor Processing Units (TPU) and we compare two algorithms: the TensorFlow built-in logic and a custom loop, optimised to have higher control of the elements assigned to each GPU worker or TPU core. The quality of the generated data is compared to Monte Carlo simulation. Linear speed-up of the training process is obtained, while retaining most of the performance in terms of physics results. Additionally, we benchmark the aforementioned approaches, at scale, over multiple GPU nodes, deploying the training process on different cloud vendor service offerings, seeking for overall efficiency and cost-effectiveness. The combination of data science, cloud deployment options and associated economics allows to burst out heterogeneously, exploring the full potential of cloud-based services.

Case Study Summary:

The scientific problem we tackled:
In particle physics the simulation of particle transport through detectors requires an enormous amount of computational resources, utilizing more than 50% of the resources of the current CERN Worldwide Large Hadron Collider Grid infrastructure. This challenge has motivated the investigation of different, faster approaches for replacing the standard Monte Carlo simulations. Deep Learning Generative Adversarial Networks are among

the most promising alternatives and, indeed, multiple prototypes are being explored for the simulation of different particle detectors. In this context, access to heterogenous infrastructure resources, such as cloud-based hardware accelerators, is essential to enable the development and then, deployment of complex models.

The computational methods we used:
We demonstrated the use of different approaches to distributed DNN training on diversified hardware (TPUs and GPUs). We also compare different approaches to cloud provisioning and orchestration methods: Kubeflow-based on GCP vs MLaaS offering on Azure cloud.

The cloud resources we used:
Multiple GPUs nodes on GCP and Azure cloud and TPUs on GCP.

The differences we've observed between locally-provided and cloud-provided resources:
There are cloud-based services (e.g TPUs) that we can't access locally on-premise not at the required scale (hundred of GPUs) what makes it very interesting to explore a hybrid cloud model.

For further information:
https://openlab.cern/ml-da

Studying COVID-19 Pandemic Misinformation and Protest Using Social Media Data

Dr. Michael Alvarez

Abstract:
Social media data is now widely used by many academic researchers. However, long-term social media data collection projects, which most typically involve collecting data from public-use APIs, often encounter issues when relying on local-area network servers (LANs) to collect high-volume streaming social media data over long periods of time. We present a cloud-based data collection, pre-processing, and archiving infrastructure, and argue that this system mitigates or resolves the problems most typically encountered when running social media data collection projects on LANs at minimal cloud-computing costs. We show how this approach works in different cloud computing architectures, and how to adapt the method to collect streaming data from other social media platforms.

For further information:
https://www.rmichaelalvarez.com/

Public transit monitoring in Baltimore

Dr. Vanessa Frias-Martinez

Author Bio:
Vanessa Frias-Martinez is an associate professor in the iSchool and UMIACS, and an affiliate associate professor in the Department of Computer Science at the University of Maryland (UMD) where she also leads the Urban Computing Lab. Frias-Martinez's research areas are data-driven behavioral modeling and spatio-temporal data mining. Her research focuses on the use of large-scale ubiquitous data to model the interplay between human mobility patterns, social networks and the built environment. Specifically, Frias-Martinez develops methodologies to model and predict human behaviors in different contexts as well as tools to aid decision makers in areas such as poverty, natural disasters or urban planning. Before coming to UMD, she spent five years at Telefonica Research developing algorithms to analyze mobile digital traces. Frias-Martinez is the recipient of a National Science Foundation (NSF) CAREER Award and a La Caixa Fellowship. She received her PhD in Computer Science from Columbia University.

Abstract:
In this presentation I will give a high-level overview of BALTO, a recently funded NSF SCC project that focuses on the design, development, deployment and evaluation of a privacy-respectful toolkit to identify and characterize the multi-factorial challenges typical of complex trips often times endured by low-income residents in Baltimore City; and to drive bottom-up, crowdsourced-informed actionable solutions via community conversations and a decision support system.

Case Study Summary:

The scientific problem we tackled:

Lack of understanding of the barriers that low-income residents in Baltimore City face when using public transit. Collecting accurate data about their trips and their quality of service perceptions is challenging. We propose a toolkit to collect door-to-door mobility experiences and quality-of-service perceptions.

The computational methods we used:
The project combines quantitative machine learning approaches to model mobility experiences from spatio-temporal data so as to identify public transit barriers; together

with focus groups to explore the design of data collection tools that are privacy respectful so as to enhance adoption.

The cloud resources we used:
Currently, we are developing the mobile app to carry out the door-to-door mobility experience collection. The app requires login and GPS/survey data storage. We are currently using AWS cognito, DynamoDB and S3 through AWS Amplify.

The differences we've observed between locally-provided and cloud-provided resources:
Many functionalities are already implemented in AWS and would require implementation from scratch in a locally-provided cloud e.g., app login.

For more information:
https://balto.umd.edu/

High Fidelity Simulation of Supercritical CO2 Power Cycles

Dr. Raghu Kancherla and Dr. Fahad Ahmad Khan

Author Bios:
Dr. Raghu Kancherla obtained his Ph.D. in Mechanical Engineering from University of Central Florida and current working as Postdoctoral Scholar. He has his masters from Indian Institute of Technology, Madras, India. After his Master, he served Nissan Auto India as a CFD engineer for four years. His research mainly focuses on applied and fundamental combustion modeling and analysis.

Dr. Fahad Ahmad Khan is a Research Cyberinfrastructure Facilitator with the Graduate Research IT team at UCF. He facilitates researchers by advising them on their cyberinfrastructure needs which include advanced computing resources, data storage, networking, and security & privacy. He is an alumnus of UCF (PhD CE 2019) and UET Lahore (MSc EE 2014, BSc EE 2007). His current research interests include sensor networks, approximate computing (BDD minimization), and applied machine learning (biomedical imaging & natural language processing).

Abstract:
Our group (Raghu Kancherla, Fahad Khan, Javier Gioia, and Subith Vasu) has been significantly contributing to make direct-fired supercritical CO2 power cycles a practical reality. Supercritical CO2 power cycles attracted global energy sector's attention due to its remarkable promise of efficiency, economics, and environmental friendliness. However, the operation of cycle in supercritical state is unconventional to modern power sector. In supercritical state the working fluid is neither a liquid nor a gas, but it will have densities like liquid and properties like a gas. As per the current state-of-the-art, the pressure of the supercritical CO2 cycle combustor is 300 bar. It is around fifteen times higher pressure than a typical combustor in existing power plants. Any experimentation at these pressures is expensive, time consuming, and even dangerous. Thus, design and development of these combustors heavily rely on high-fidelity computational tools such a CFD.

Combustion CFD solvers are computationally intensive as the nature of the combustion ordinary differential equations are more stiff, coupled, and non-linear. On top of that, supercritical state adds more complex thermal, and transport property models. Therefore, supercritical CFD simulation can't typically fit into the memory of a single machine and requires computation to be shared on multiple communicating nodes. The effective communication between the computing

nodes and the availability of quick computing resources are two primary challenges for us to compete with other researchers in this area. Currently, we have utilized Azure HPC for these intense simulations. Our presentation discusses our experience with Azure HPC, our preliminary combustor development based on our high-fidelity CFD, and various challenges and lessons learnt in provisioning & deploying resources in UCF's Azure subscription while working with various technical teams.

For more information:
https://cater.cecs.ucf.edu/