Application of AceCAST for the NYSM Profiler Data for Hyperlocal Weather Forecast Evaluation

Outline

- AceCAST
  - Background
  - Model performance
  - Model advancements
  - Future outlooks
- Data Assimilation (NYSM Profile network)
  - Radiometer
  - Project evaluation
  - Societal Impact

AceCAST

What is it?

- NWP model ported from CPU to GPU
- Uses the Weather Research and Forecasting Model (WRF)
- GPU accelerated weather forecasting
- Supports major WRF dynamics, physics schemes, and namelist options

GPU enables:

- Much faster time-to-solution
  - Compared to standard computing methods (CPU)
- Higher frequency capability (timesteps)
- Greater awareness of localized weather phenomena
- Reduced cost
  - Computational
  - Financial

CPU to GPU Porting for NWP Models Background

- Most work has been with porting parts of an NWP model such as microphysics schemes, advection, climate radiation schemes, and chemical kinetics/subroutines [1-6]
- Others have ported entire models for specific applications such as ocean modeling, climate modeling, etc. [7-10]
- Research gap with porting entire general purpose NWP models

Motivation

- The ever increasing need for computational speed and power in a world overflowing with data has given rise to a lag in capability

Performance

Test Case

- 500 × 500 grid at 2km resolution run on 3 different processors (CPU and GPU)

Speed (simulation time divided by real time):

- 1 GPU up to 4.3 times faster than 1 CPU (1a-1b)
- 2 NVIDIA A100 GPU's nearly 2 times faster than 4 Intel Xeon CPU's (2a-2b)

Cost per run*:

- GPU about half as expensive as CPU counterpart, up to 3.4 times cheaper
  - Through about 8 processors, beyond that workload is not enough for maximizing resource efficiency

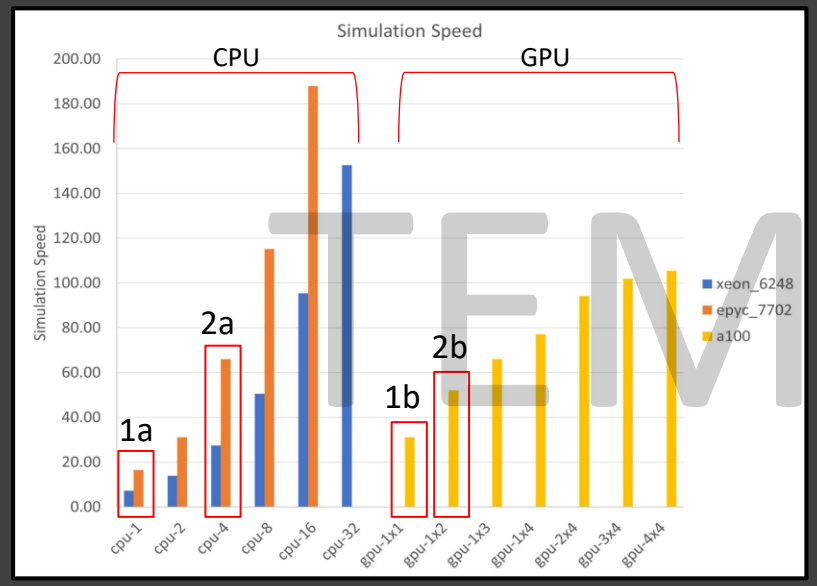

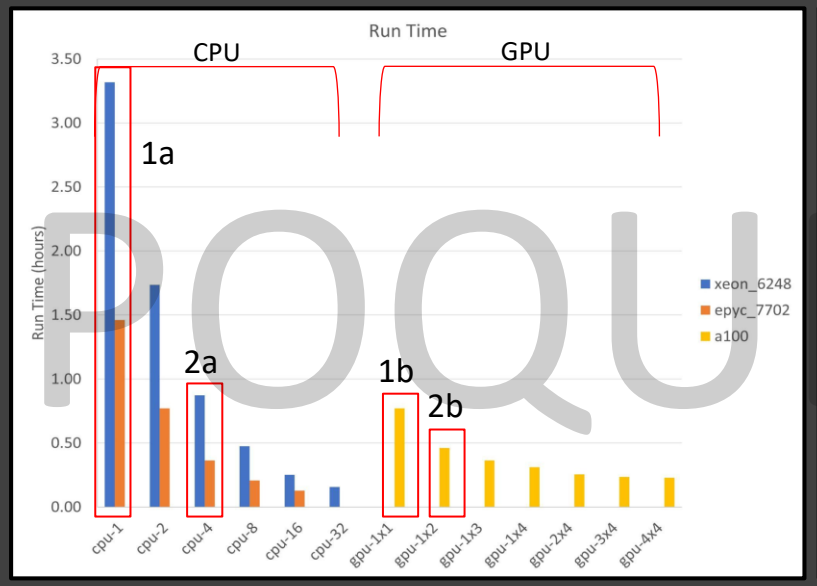

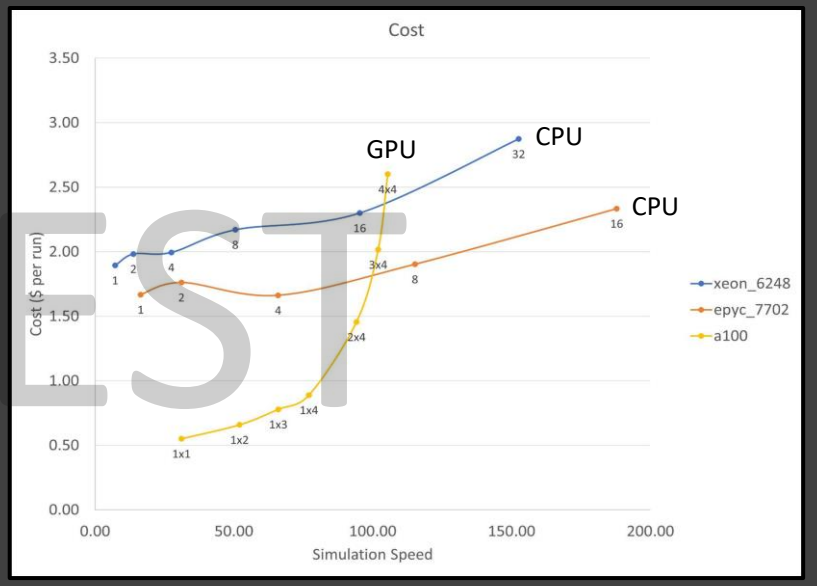

*Based on estimated node costs of $10,000 for Intel Xeon, $20,000 for Intel Epyc and $50,000 for NVIDIA A100 GPU

Verification and Validation

Extensive review process to ensure AceCAST meets high quality standards*

- For every release of AceCAST TempoQuest runs AceCAST through 1,000s of small simulations to test for different dynamics options and physics options
- Through a series of regression tests, TempoQuest also checks AceCAST:
  - Model performance across several practical temporal and spatial forecast ranges
  - Through 3 (main) different compiler builds
  - Different domain sizes and regions
  - Runtimes
  - Different GPU's (NVIDIA V100 and NVIDIA A100)
  - Support tests
  - Host-compile/exec. preservation tests
  - Bit-for-bit tests
  - False-positive support tests
  - etc…
- Ensure GPU simulations are consistent compared to CPU and that differences are within an acceptable tolerance

*TempoQuest does not directly develop the WRF model, thus any issues in the WRF model are automatically transferred to AceCAST. WRF issues are reported to the development team.

AceCAST Advancements In The Last Year

- Upgraded to base WRF v4.4.2
  - Latest version – December 2022
- Upgraded to AceCAST v3
  - More supported features
  - Expected end of January
- Support for AMD Instinct GPU's
  - Anticipated early June
- AceCAST OnDemand release
  - Cloud-Based weather forecasting
- Major updates to AceCAST visualization software (WSV3)
- *Continuous* addition of new physics

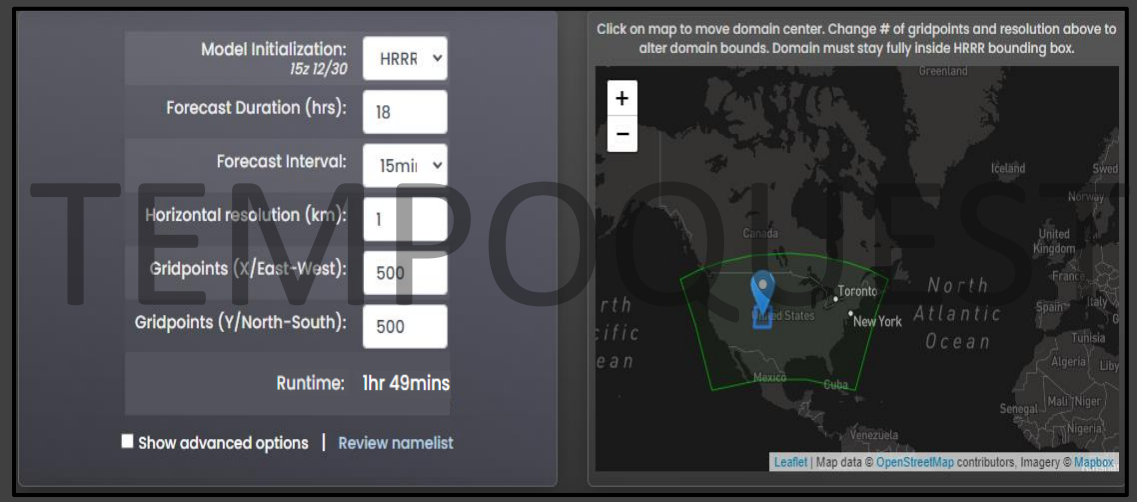

AceCAST (WRF) OnDemand Start Page

WSV3 Visualization

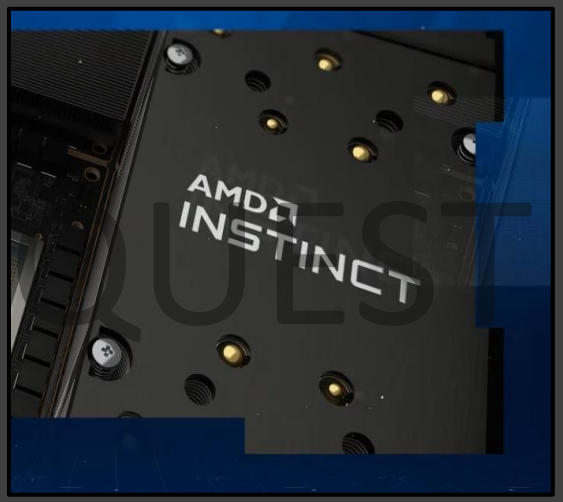

AMD Instinct GPU

Future Outlooks

- Fully porting the Data Assimilation (WRF-DA) and Chemical (WRF-CHEM) Models into AceCAST
  - TQI has already successfully implemented computationally inexpensive FDDA such as:
    - Surface analysis
    - Analysis nudging
  - And plans to soon include:
    - Spectral nudging
    - Observational nudging
  - Radiometer observations case study
- Creating hybrid WRF-Artificial intelligence weather forecasts

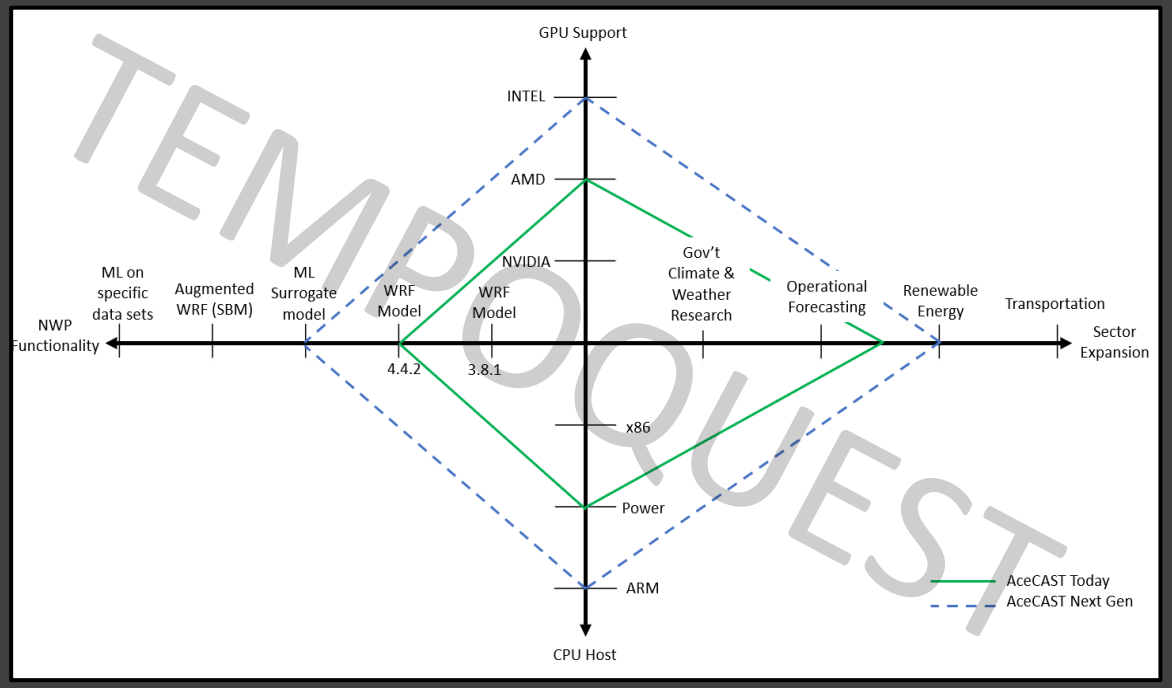

TempoQuest current roadmap

Radiometer*

- The New York State Mesonet (NYSM) profiler network observes boundary layer thermodynamics and winds
  - Network of Radiometers
  - Remote sensing observational platform
- Accuracy similar to radiosondes
- A 200 (EW) x 100 (NS) km area including six NYSM profiler sites, provides an ideal test bed for evaluation of high spatiotemporal resolution data for hyperlocal (1 km) weather forecasting
- Assimilation of NYSM profiler data at 5-min intervals
- Evaluate parallel process hyperlocal forecast latency and accuracy

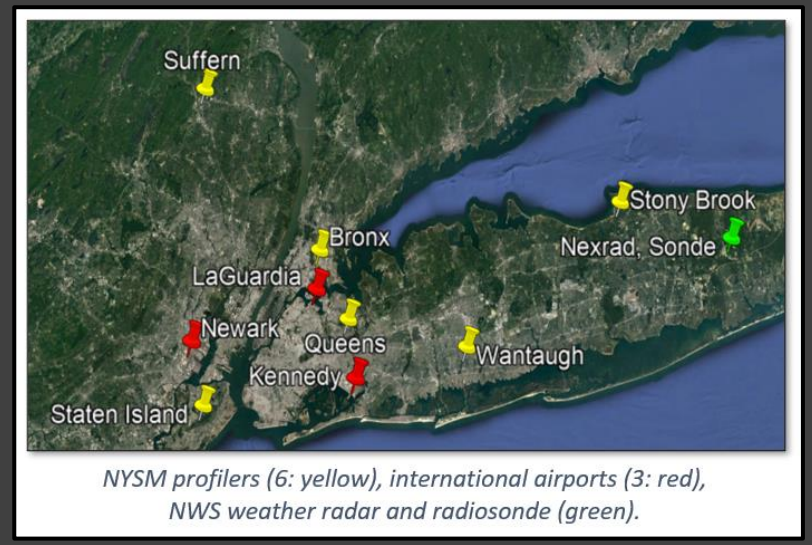

NYSM profilers (6: yellow), international airports (3: red), NWS weather radar and radiosonde (green).

Radiometer Locations near New York, New York

Radiometer Device

*Courtesy of Radiometrics (Boulder, CO)

WRF Data Assimilation

Three different methods:

- Variational – Most common
  - Includes Objective Analysis AND Initialization step
  - 3D/4D-VAR – Time independent/dependent
- Objective Analysis – Simplified
  - Additional Steps for processing needed
  - Older method and noisier method
- Newtonian Relaxation – Nudging
  - Nudging the simulation as it runs
  - Analysis
    - Towards difference of simulated and spatial/temporal interpolations of reference state at each grid point
  - Spectral
    - Fourier decomposition
    - Difference of modeled and reference states
    - Longest waves are relaxed toward the equivalent wavelengths in reference state

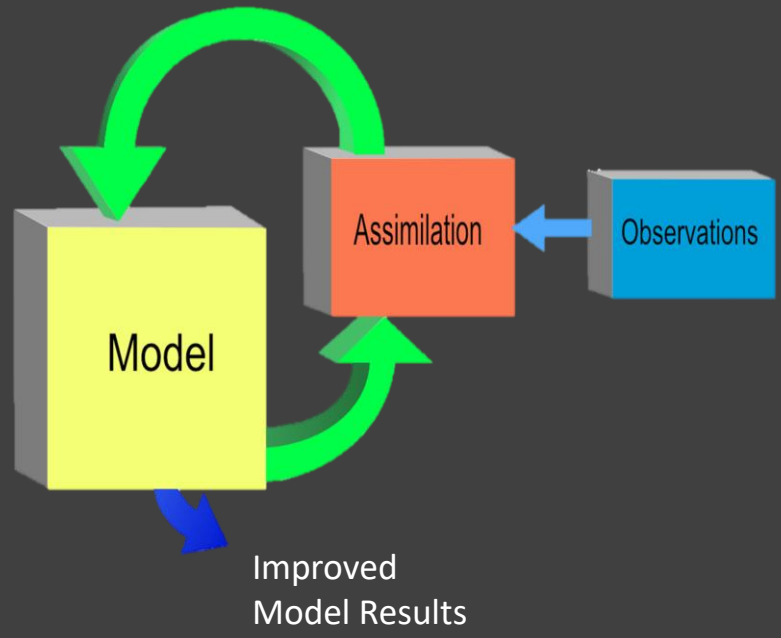

Methodology

- Assimilation input data
  - NYSM Radiometer data
    - Temperature
    - Humidity
  - Background weather data
    - Nested North American Model (3km)
- Process
  - Develop WRF simulations/forecasts using hybrid accelerated weather forecasting techniques
  - Initially CPU variational Data Assimilation
  - Use of GPU's enables increased assimilation speed and higher resolution

Expected Project Outcome Evaluation

Forecast evaluation*:

1. Convective initiation, development, timing and location
2. Low-level winds
3. Freezing isotherm altitudes during cold seasons
4. Additional weather forecast parameters

Applications:

1. Space launches
2. Crewed and Uncrewed Aerial Vehicle operations
3. Solar and wind energy management
4. Electric load forecasting
5. Numerous other weather-sensitive operations

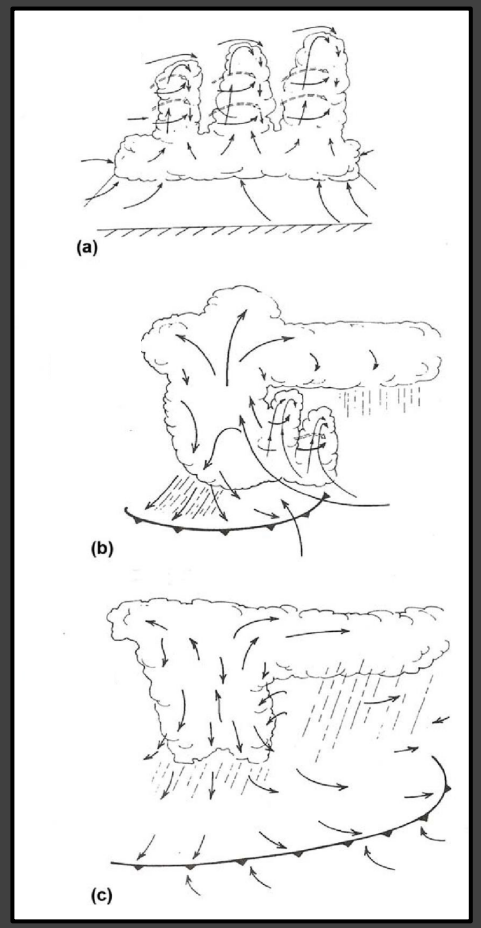

*When comparing simulated results to all available observations

Project Impacts

- High resolution results have shown improved forecast accuracy
  - In reference to spatial and temporal modeling and profiling (see red boxes)
- Significant potential to improve forecast

High-Resolution (1km) Accelerated Weather Forecasting

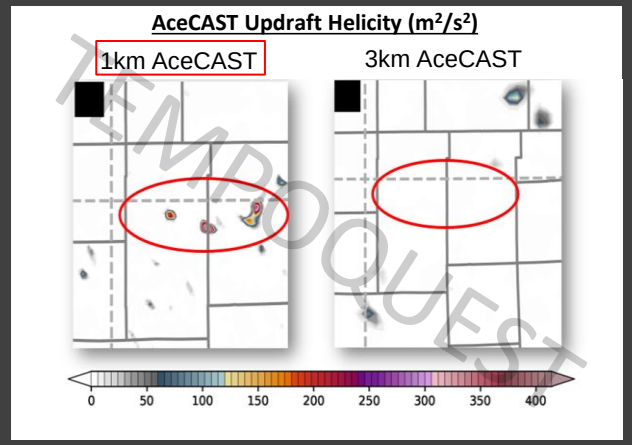

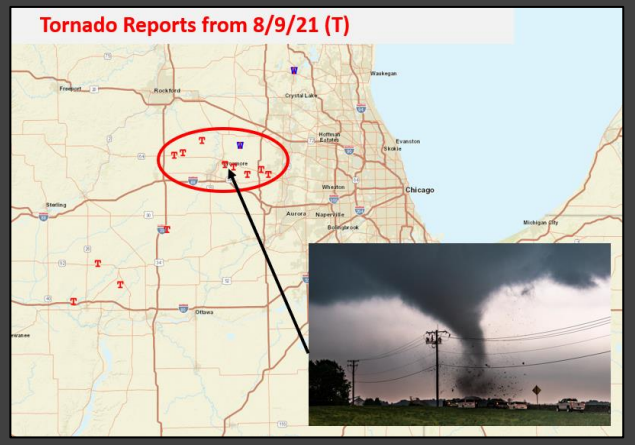

3km simulation underperforms when it comes to significant severe weather potential

High Frequency Radiometer Data

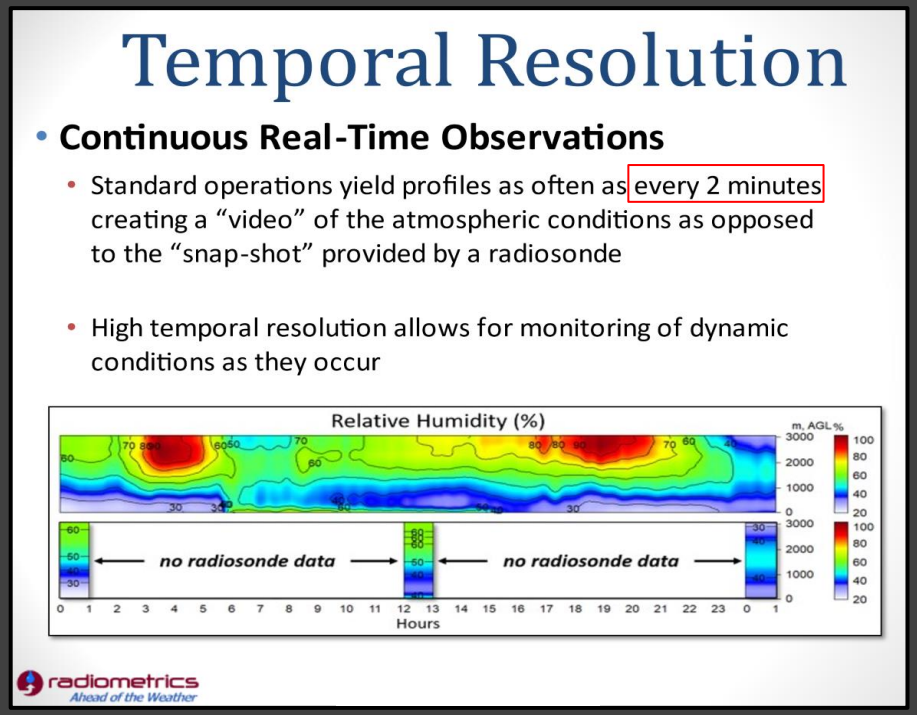

Preliminary Results –Analysis Increment for Temperature

- Corrections to the first guess (background)
- Observations adding value to the model analysis by providing corrections and improvements in the analysis
- Most significant over the western Long Island region

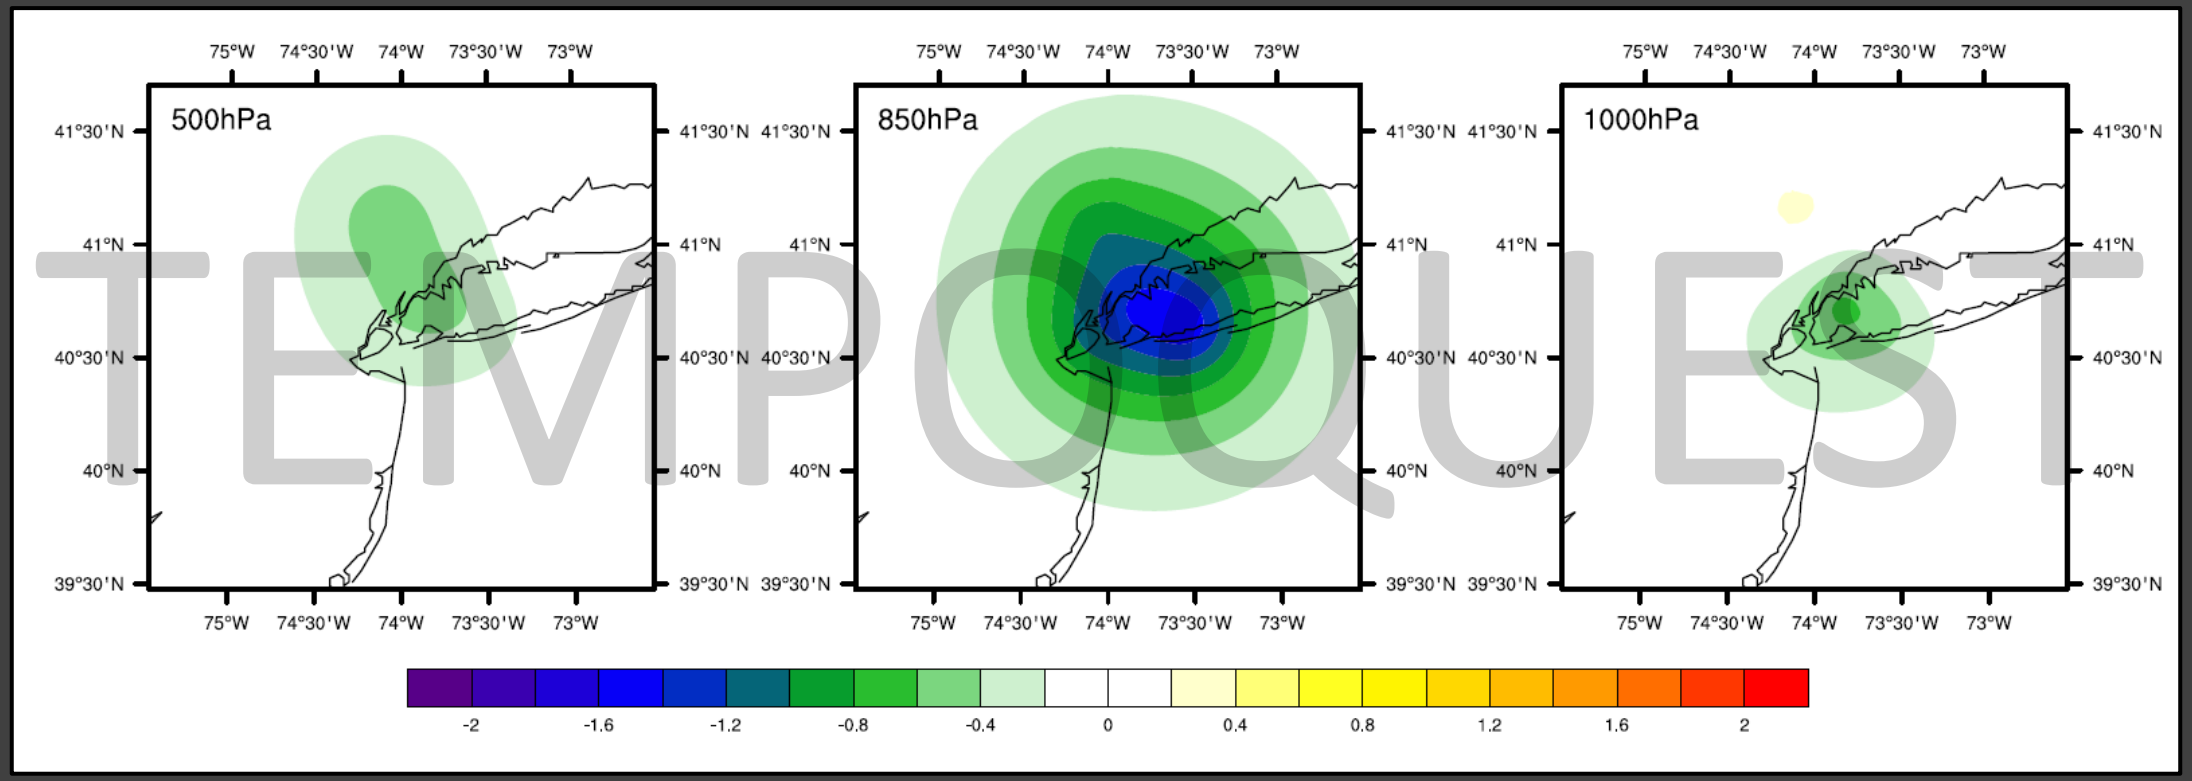

Preliminary Results –Analysis Increment for Relative Humidity

- Corrections to the first guess (background)
- Observations adding value to the model analysis by providing correction and improvement in the analysis
- Most significant over the western Long Island region

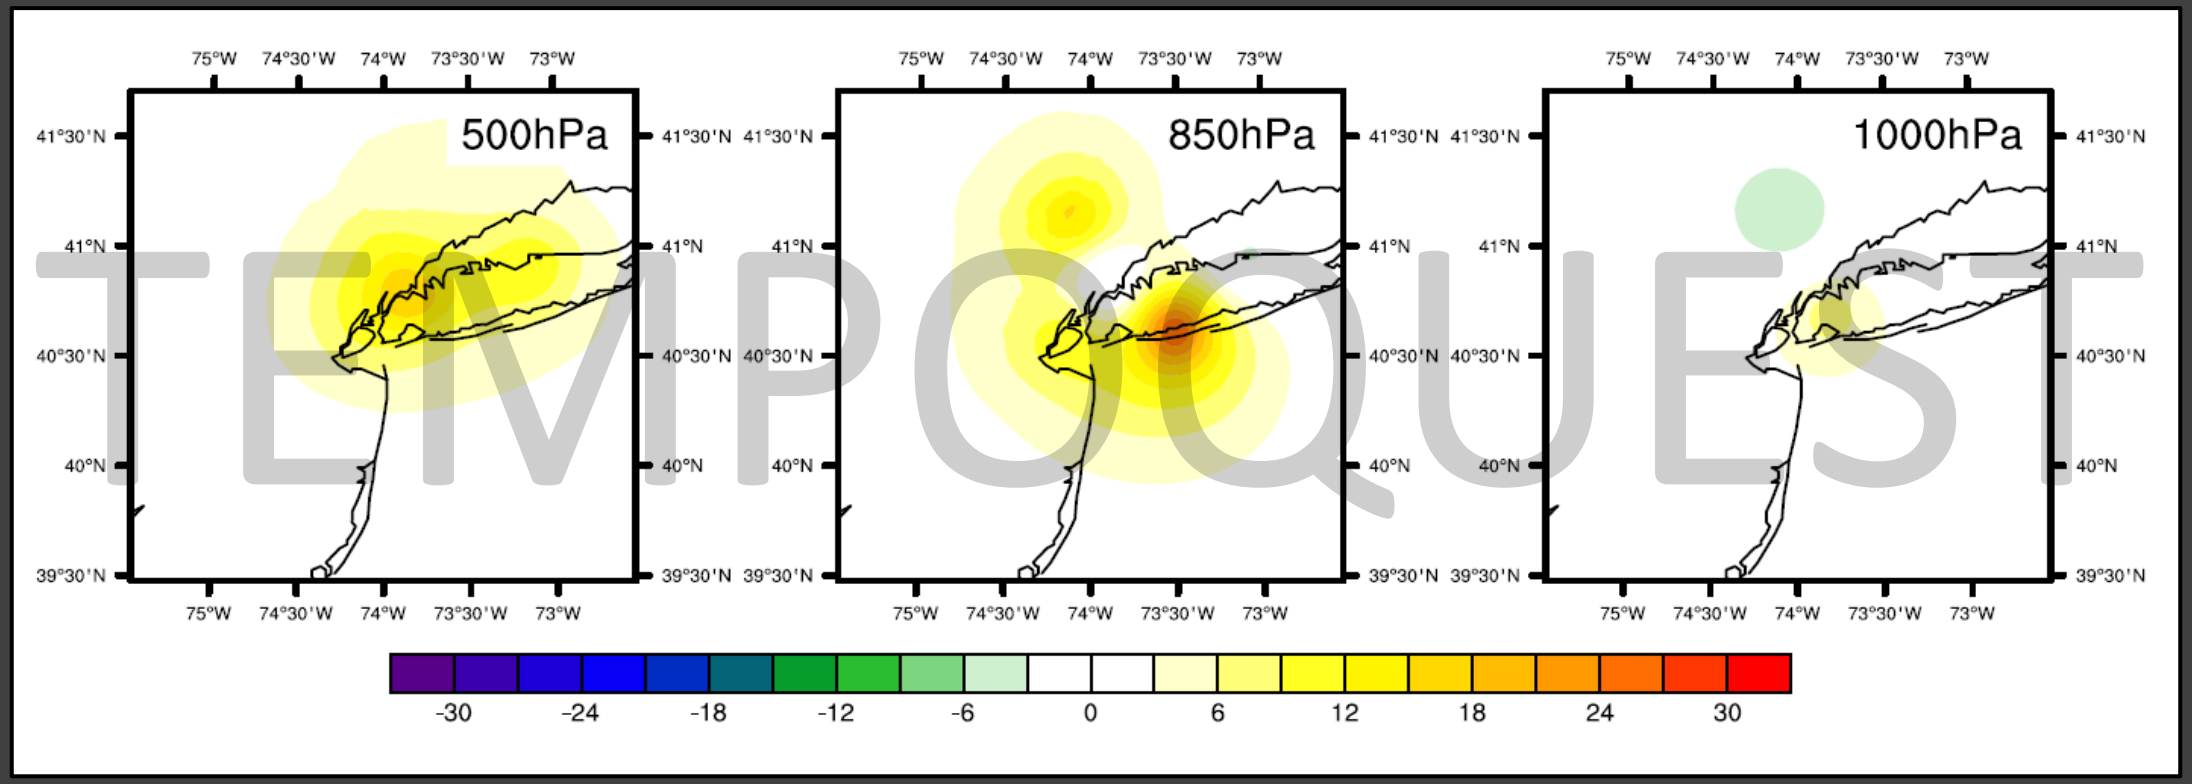

In a Nutshell

- TempoQuest has accelerated WRF using GPU's
  - Regional forecast modeling
  - Computational benefits compared to CPU counterpart
- AceCAST serves as a drop-in replacement for existing WRF configurations
- GPU-acceleration enables:
  - Performance benefits
    - faster time-to-solution
    - higher resolution
    - deeper insights
- Using Radiometer data AceCAST has the potential to provide higher accuracy hyperlocal weather forecasts much faster compared to conventional methods

Thank You For Your Attention

Longer Question?

Let’s start a conversation!

Email us at: support@tempoquest.com

References

1. Sharma and Kulkarni (2015)
2. Wang et. al. (2011)
3. Michalakes and Vachharajani (2008)
4. Alvanos and Christoudia (2017)
5. Linford et al. (2009)
6. Korwar et. al. (2014)
7. Xu et. al. (2015)
8. Shimokawabe et. al. (2010)
9. Leutwyler et. al. (2016)
10. Giorgetta et. al. (2022)