

**Project Report
TIP-146**

**Mobile Capabilities for Micro-
Meteorological Predictions: FY20
Homeland Protection and Air Traffic
Control Technical Investment Program**

**M.D. McPartland
T.A. Bonin**

26 March 2021

Lincoln Laboratory
MASSACHUSETTS INSTITUTE OF TECHNOLOGY
LEXINGTON, MASSACHUSETTS



DISTRIBUTION STATEMENT A. Approved for public release. Distribution is unlimited.

This report is the result of studies performed at Lincoln Laboratory, a federally funded research and development center operated by Massachusetts Institute of Technology. This material is based upon work supported by the United States Air Force under Air Force Contract No. FA8702-15-D-0001. Any opinions, findings, conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the United States Air Force.

© 2021 Massachusetts Institute of Technology

Delivered to the U.S. Government with Unlimited Rights, as defined in DFARS Part 252.227-7013 or 7014 (Feb 2014). Notwithstanding any copyright notice, U.S. Government rights in this work are defined by DFARS 252.227-7013 or DFARS 252.227-7014 as detailed above. Use of this work other than as specifically authorized by the U.S. Government may violate any copyrights that exist in this work.

**Massachusetts Institute of Technology
Lincoln Laboratory**

**Mobile Capabilities for Micro-Meteorological Predictions:
FY20 Homeland Protection and Air Traffic Control
Technical Investment Program**

*M.D. McPartland
T.A. Bonin
Group 43*

Project Report TIP-146

26 March 2021

**DISTRIBUTION STATEMENT A. Approved for public
release. Distribution is unlimited.**

Lexington

Massachusetts

This page intentionally left blank.

EXECUTIVE SUMMARY

Existing operational numerical weather forecast systems are geographically too coarse and not sufficiently accurate to adequately support future needs in applications such as Advanced Air Mobility, Unmanned Aerial Systems, and wildfire forecasting. This is especially true with respect to wind forecasts. Principal factors contributing to this are the lack of observation data within the atmospheric boundary layer and numerical forecast models that operate on low-resolution grids. This project endeavored to address both of these issues. Firstly, by development and demonstration of specially equipped fixed-wing drones to collect atmospheric data within the boundary layer, and secondly by creating a high-resolution weather research forecast model executing on the Lincoln Laboratory Supercomputing Center.

Some success was achieved in the development and flight testing of the specialized drones. Significant success was achieved in the development of the high-resolution forecasting system and demonstrating the feasibility of ingesting atmospheric observations from small airborne platforms.

This page intentionally left blank.

TABLE OF CONTENTS

	Page
Executive Summary	iii
List of Illustrations	vii
1. METEOROLOGICAL DRONE DEVELOPMENT	1
1.1 Airframe Development	1
1.2 Wind Estimate Preparation	2
1.3 Water Vapor and Temperature Measurement	6
2. HIGH-RESOLUTION NUMERICAL WEATHER PREDICTION MODELING DEVELOPMENT	9
2.1 Terrestrial Data Improvements	9
2.2 Setup and Configuration of WRF	11
3. METEOROLOGICAL FLIGHT TESTING	21
4. ASSIMILATION OF DRONE MEASUREMENTS INTO WRF	27
5. CONCLUSION	33
Glossary	35
References	37

This page intentionally left blank.

LIST OF ILLUSTRATIONS

Figure No.		Page
1	Magnetometer calibration testing on custom turntable.	3
2	Calibration verification of compass #3. Includes short-term (T0=zero seconds) and long-term (T30=30 seconds) post rotation measurement errors.	4
3	Calibration verification of compass #1. Includes short-term (T0=zero seconds) and long-term (T30=30 seconds) post rotation measurement errors.	5
4	Progressive (wind-up) heading errors evident in compass #1. Each measurement is taken after a 360 degree heading rotation in either direction.	6
5	An HTY271 embedded water vapor and temperature sensor after harness attachment.	7
6	Water vapor and temperature sensor installation for flight. Placement under the root of port wing protects from solar heating, propeller wash, and is also isolated from internal heating via the foam fuselage.	8
7	Example test platform in flight with mounted sensors.	8
8	High-resolution, low-resolution, and the difference of terrain data (upper panels, left to right respectively). Bottom panels are zoomed-in on the region indicated by the boxes in upper panels.	10
9	Simplified diagram showing flow of WRF system.	11
10	Example nested WRF domains (D01, D02, D03) centered around Boston. The geographical dimensions of a domain are three times larger than its next inner domain.	13
11	Comparison of HRRR and 90-m WRF wind forecast at 10-m around the Hudson River Valley region.	16
12	Winds at 400 ft with vertical motion shaded for a sea breeze over Boston on July 29, 2020.	17
13	Modeled cloud cover from the HRRR, high-resolution WRF, and observed visible satellite over San Francisco on the morning of July 31, 2020. The WRF and HRRR forecasts are 4-hr forecasts initialized at 12:00 UTC.	18

LIST OF ILLUSTRATIONS (Continued)

Figure No.		Page
14	Wind forecast centered on Paradise, CA on November 8, 2018. The forecast is valid at 13:15 UTC, which is 75 minutes after it was initialized at 12:00 UTC.	19
15	Flight testing site at Fort Devens, Devens, MA.	22
16	Negative effect of loss of control due to communications failure. Complete separation of motors from wing mounts and cracked fuselage. Creative overnight reconstruction produced a flyable platform.	23
17	Field replacement of platform motors and power control units. Customized airframe modification was necessary for installation of parts not originally intended for unit.	24
18	Successful meteorological data collection flight. platform proceeded to 3995 ft, which was the maximum allowable flight altitude for the restricted airspace R-4102.	25
19	Meteorological data, water vapor and static temperature, as recorded at varying pressures (altitudes) along the platform's ascent and descent.	26
20	Comparison of the temperature and humidity profiles over Fort Devens, MA after assimilating meteorological drone observations using WRFDA.	28
21	Temperature difference of the model initial conditions of before and after assimilating the meteorological drone observations using WRFDA.	29
22	ASOS temperature and dew point one minute observations (OMO) taken at KFIT compared to forecasts from HRRR, WRF, and WRF run with data assimilated from drone observations (DO) at the same location.	31

1. METEOROLOGICAL DRONE DEVELOPMENT

1.1 AIRFRAME DEVELOPMENT

In the original vision of the program, up to 10 fixed-wing drones were to be constructed and equipped as meteorological sensing platforms. These systems were to be deployed for data collection at altitudes and distances that would have required a Certificate of Authorization (COA) or other authorization from the Federal Aviation Administration (FAA) to conduct Beyond Visual Line Of Sight (BVLOS) operations.

The systems were to be constructed using the same basic platform that has been used for many other flight tests and experiment out of our Small Unmanned Aerial System (sUAS) Test Bed. Experience with this platform, called the MTD-2, proved that it was capable of the required speeds, capacity and endurance, up to 2 hours, needed for flight tests. This fixed-wing, twin engine platform, which has a wing span of 2 meters, was not available at the beginning of this program and the model was discontinued by the manufacturer. This became a major impediment because obtaining a suitable replacement turned-out to originally be almost impossible. A Department of Defense (DoD) memorandum in 2018 title “Unmanned Aerial Vehicle Systems Cybersecurity Vulnerabilities” restricted the use of any UAS systems to those from domestic suppliers unless a special waiver was granted. Each year the Laboratory requests, and needs to justify to the DoD why, a waiver should be granted for a specific list of drones the Laboratory wants to operate. The MTD-2 was one of those on the very short list. The process to obtain the waiver is non-trivial, is performed only once a year, and has a long process time. Finding a suitable airframe to replace the MTD-2 was not difficult, but obtaining approval to use the new airframe would be.

The Laboratory did not possess more than 4 MTD-2 airframes and each in various states of “readiness.” The program had planned to grow the number of MTD-2 platforms in support of the small UAS testbed as well as satisfy the need of this program. After making a concerted attempt to acquire any available MTD-2 still available for sale, an alternate solution was needed. Multiple airframes that have similar performance characteristics to the MTD-2 were available. Unfortunately all of these airframes, like the MTD-2, were manufactured in China. Based on the Laboratory’s understanding of the DoD memorandum, and the fact that the new platforms were not on the waiver list, the program was in a quandary. How could further suitable platforms be obtained?

Like the MTD-2 platform, the newly identified airframes were a pair of carbon tubes used for the wing spars and a series of Expanded Polyolefin foam components that made up the fuselage and wing surfaces. When purchased outright, it was our professional opinion that these platforms themselves did not constitute a drone, but were just airframes. However, the legal counsel of the Laboratory’s interpretation of the DoD memorandum was that all of these airframes were drones. After studying the memorandum, it was clear to us that the airframes of interest should not be restricted under the memorandum regulation, which was explicitly concerned with cybersecurity risks of UAS systems that were not domestically manufactured. Since the airframes that were desired to be acquired had no electronics of any type, they could not pose a

cybersecurity risk. All the electronics that are used to operate and communicate with the MTD-2 and similar platforms use domestically sourced components or foreign components on an existing waiver.

After many discussions and lengthy deliberations with Laboratory Legal counsel, it was agreed that the DoD memorandum should not restrict the Laboratory from acquiring UAS fuselages and workings from non-domestic suppliers as long as there are not electronic components. Counsel agreed to write new guidance to distribute to the Laboratory to this point. We were thus given approval in written form that we could acquire the airframes of interest.

By the time we gained permission to acquire new airframes and then build them up, we were well into the project timeline. There was simply not enough time and money to build-up the multitude of platforms nor obtain a COA to operate across various sites of interest. Instead, we would take advantage of the four existing MTD-2 platforms we had at the Laboratory. These aircraft were to be used in flight tests at Fort Devens, where a COA was already in place, for a separate program to evaluate communication via Bluetooth in support of the Alert and Collision Avoidance System X unmanned (ACAS Xu) program.

The existing flight test plan for this program was modified to help reach our sensing goals by conducting flights well above the originally planned altitudes, in this case up to 4000 ft. This is the COA limit specified for the restricted airspace around Fort Devens. At altitudes of 4000 feet, we performed BVLOS operations as the airframe will be hardly visible by the unaided human eye. The vertical limit provided access to only the lower portion of the atmospheric boundary layer and the shape of the restricted airspace enabled access to a limited lateral region. The four platforms would fly simultaneously staggered across the permissible altitudes. Even with these limitations, practical knowledge and experience would be obtained. The meteorological data collected would be later assimilated into the high-resolution forecast model with promise to provide a positive impact on a regional forecast.

In addition to building and outfitting the four platforms with specialized humidity and temperature sensors and the software components to control and downlink their data, a critical aspect of each platform was to collect airspeed, ground speed, and attitude data such that local horizontal wind estimates could be obtained. A critical factor of this is proper measurement of instantaneous magnetic heading.

1.2 WIND ESTIMATE PREPARATION

Previous experience with these specific platforms on a separate program had shown that the magnetometer data, or the interpretation of the data, led to inaccurate and varying magnetic heading estimates that are central to a wind estimation capability for the UAS platform. The lack of accuracy was not a factor for the particular program at that time and it was then considered that the poor performance was due to poor calibration. Because the same platforms were to be used in these flight testing, significant attention was paid to the calibration and verification of the platforms' magnetometers.

The magnetometer calibration of the platforms was conducted in the parking lot adjacent to the parking ramp. The multi-axis calibration procedure was conducted multiple times on each system paying

careful attention to obtain complete and accurate raw data. A purpose-built non-ferromagnetic turntable was constructed to validate the calibration, shown in Figure 1. A precision orientation compass was affixed approximately 5 feet from the platform center to validate the reported magnetic heading post calibration.



Figure 1. Magnetometer calibration testing on custom turntable.

The initial validation tests showed that the reported magnetic heading post calibration did not reflect the actual magnetic heading at all. Repeating the calibration process multiple times and assuring the correct process was followed did not proffer better results. It was suspected that the placement of the magnetometers were affected by active electronics in the platform, so the precision compass was used to identify magnetic fields in the region of the magnetometers (there are three independent magnetometers of three axes each).

With the compass, we did in fact find an unexpected magnetic field near the primary magnetometer. At some earlier date, several rare earth magnets were embedded in the foam of the fuselage to help secure the platform's foam canopy, which was also embedded with pairing rare earth magnetics. One magnet was located less than two inches from the primary magnetometer.

With the magnets removed from each platform and no other detectable magnetic fields identified, the systems were recalibrated. Unfortunately, the resulting performance was still disappointing. For each of the systems, the reported magnetic heading had the following characteristics:

- Large pointing errors
- Varies with heading
- Non-repeatable
- Time varying
- Varies with progression

In Figure 2 and Figure 3, we see the post calibration errors for two of the onboard compasses. The indications of T0 or T30 refer to the indicated compass values taken either zero seconds after arriving at the new test heading or 30 seconds after arriving at the heading. As we can clearly see in Figure 2 regarding compass #3, there is a time varying component in the compass error as can be seen by the separation of the blue and orange lines. This characteristic does not appear in compass #1 (see Figure 3) but the post calibration spot checks do demonstrate some of this behavior.




Figure 2. Calibration verification of compass #3. Includes short-term (T0=zero seconds) and long-term (T30=30 seconds) post rotation measurement errors.




Figure 3. Calibration verification of compass #1. Includes short-term ($T0=$ zero seconds) and long-term ($T30=30$ seconds) post rotation measurement errors.

Additional errors called wind-up errors were also observed. These error are measured after rotating the vehicle 360 degrees in either direction returning to the original heading after each rotation and measuring the difference between the reported and the actual. As can be seen in Figure 4, the error varied with the number of rotations. The error was decoupled with time as the $T0$ and $T30$ measurements show little difference.




Figure 4. Progressive (wind-up) heading errors evident in compass #1. Each measurement is taken after a 360 degree heading rotation in either direction.

The end finding of this experiment was that none of the platforms were equipped with suitable components that could provide accurate estimates of the platforms' magnetic heading. True heading, which is a function of magnetic heading, is a required input to perform wind estimates. Thus, the existing platforms' system could not be used for generating local wind estimates during flight without the application of new sensory and software components, which the remaining schedule and program funds did not permit.

1.3 WATER VAPOR AND TEMPERATURE MEASUREMENT

The platforms also needed to be outfitted with temperature and water vapor sensors. A myriad of such devices were evaluated. Considerations of the performance characteristics included, range, accuracy, frequency response, hysteresis, data transfer modality, size, cost, weight, and form factor. After consideration of multiple candidates, and an evaluation that included performance testing, a small form factor low-profile candidate was selected. The device, the HTY271 and manufactured by Innovative Sensor Technology, embeds both sensors, water vapor, and temperature on a single substrate.

Like all the candidates evaluated, communication with the embedded subsystems was through a synchronized Serial Peripheral Interface (SPI). SPI ports are available as a programmable communication port on RaspberryPi processors, which was the compute and communication unit used for both collecting sensor data and communicating to the ground over a radio link. A set of specialized fine-wire harnesses had

to be crafted for the sensors. The harnesses had to be sturdy enough for flight and attached to the substrate without damaging the embedded components. As the components are very prone to suffer heat damage, significant attention was given to the soldering process of components. An example of an assembled unit is presented in Figure 5.



Figure 5. An HTY271 embedded water vapor and temperature sensor after harness attachment.

One of these units were fastened to the port side of the platform fuselage underneath the wing root as shown in Figure 6. This placement kept it protected from solar heating and in the free air slip stream away from the propeller wash, which can increase the local air temperature while under thrust. Figure 7 shows one of the platforms in flight with the meteorological sensor visible on the port side, as well as a vertical dorsal-side antenna used for ground-link communications.




Figure 6. Water vapor and temperature sensor installation for flight. Placement under the root of port wing protects from solar heating, propeller wash, and is also isolated from internal heating via the foam fuselage.



Figure 7. Example test platform in flight with mounted sensors.

2. HIGH-RESOLUTION NUMERICAL WEATHER PREDICTION MODELING DEVELOPMENT

For the modeling component of this program, a high-resolution numerical weather prediction (NWP) model capability has been developed to run on the Lincoln Laboratory Supercomputing Center (LLSC). For this, the open-source Weather Research and Forecasting (WRF) model was leveraged. The WRF model is the state-of-the-art mesoscale NWP system designed for both atmospheric research and operational forecasting applications. It was initially developed in the late 1990s as a collaborative effort between the National Center for Atmospheric Research (NCAR), the National Oceanic and Atmospheric Administration (NOAA, represented by the National Centers for Environmental Prediction [NCEP] and the Earth System Research Laboratory), the U. S. Air Force, and Naval Research Laboratory, the University of Oklahoma, and the FAA. The WRF is currently maintained by NCAR, and updates are released periodically incorporating updates that have been developed by its users and are fully documented and vetted before wider release. The version that has been used in this program is WRF V4.1.2, which was released in July 2019 and is available on GitHub. The codebase is primarily written in FORTRAN and C.

2.1 TERRESTRIAL DATA IMPROVEMENTS

The WRF system is comprised of the WRF Preprocessing System (WPS), WRF Data Assimilation (WRFDA), the WRF model, and a suite of post-processing programs and utilities. From a modeling standpoint, the objective of this program was to stand-up a high-resolution (90-m) weather modeling capability that incorporates recent observations collected from the meteorological drone. For successful implementation of this capability, substantial changes to the static terrestrial data files used by WPS were necessary to properly resolve and model atmospheric flows, as many of the default geography files that came with the WRF package were coarser than the objective 90-m resolution. The static terrestrial files contain information on the terrain elevation, sub-grid orography (mountain topography), land use characteristics, soil type, albedo, and greenness fraction. Some of the files (i.e., greenness fraction and albedo) are based on climatological statistics that differ by month, reflecting seasonal variation in vegetation. For the short-term high-resolution forecasting capability, the coarseness of the sub-grid orography and soil type are acceptable, as the sub-grid orography is typically small since most terrain features are adequately explicitly resolved and the soil type is more impactful for long-range simulations. Thus, the focus was to obtain higher-resolution terrestrial files for elevation, land-use, and albedo, which have significant impact on the accuracy of the targeted simulations.

The default terrain elevation in the WRF package is 30 arc seconds (approximately 1 km), which is inadequate for high-resolution modeling, particularly of the low-altitude wind in areas with complex terrain where local channeling effects can occur. To improve on this, high-resolution terrain elevation (down to ≈ 30 -m) from the Digital Terrain Elevation Data (DTED) dataset were leveraged. The very high-resolution (≈ 30 -m) data are only available over small select regions. However, the entire United States and significant other portions of the world are completely covered by the DTED Level 1 data, which has a resolution of

≈90-m and was used to improve the terrain data used in WRF. The DTED Level 1 data were obtained from the Lincoln Library's GeoServer through a query to the database, in which a GeoTIFF image was returned. Because WPS ingests static terrestrial data files in a Geogrid binary format consisting of many files, each containing a number of tiles, the GeoTIFF images were converted into the binary format by using an open source utility (https://github.com/openwfm/convert_geotiff). After the data were converted to the correct format, the DTED Level 1 data were then integrated into WRF through the WPS. Sample terrain elevation data over the New York and New Jersey border demonstrating the improvement with the DTED1 data are shown in Figure 8.




Figure 8. High-resolution, low-resolution, and the difference of terrain data (upper panels, left to right respectively). Bottom panels are zoomed-in on the region indicated by the boxes in upper panels.

The default land use categorization used with the WRF package is at 30" resolution (≈1 km) and contains 20 different land use types as derived from Moderate Resolution Imaging Spectroradiometer (MODIS) data. The land use types consist of different types of forests, shrubland, grassland, wetlands, croplands, urban areas, tundra, snow and ice, and water. Proper characterization of the land use is necessary to model the evaporation and transpiration rates, affecting the humidity and temperature fields. The local land use also affects the wind profile, as the wind increases more rapidly with height over smooth surfaces (e.g., water, grassland) and more slowly with height over rough surfaces (e.g., urbanized, forests). Higher-resolution data (9", ≈250-m) using the National Land Cover Database (NLCD) was obtained from the WRF users page containing supplemental data files. This dataset contains 40 land use classes, providing more

detailed information on the exact characteristics particularly in urban areas where developed areas are further classified based on their intensity/density. The NLCD has a resolution of 30-m. Due to a maximum size limit of the geogrid format, the native resolution was not used here since covering the United States at such a high resolution exceeds the maximum number of tiles allowed. For future high-resolution runs, the 30-m NLCD data can be downloaded online (<https://www.mrlc.gov/viewer/>) for a selected area, and the GeoTIFF can be converted to the geogrid binary format to cover a small domain.

The albedo determines the amount of solar radiation that is absorbed by a surface and the portion that is reflected. Different options for providing high-resolution albedo data were explored, as the default WRF albedo dataset is given at 3' (≈ 5 km) resolution with a different value for each month. This was a lower priority compared to the land use and elevation data, as the albedo is generally more homogeneous. Satellite-derived observations can be used to estimate albedo, but no high-resolution climatological databases were found that could replace the default albedo dataset. Future development efforts could take advantage of these satellite observations to provide high-resolution recent albedo maps used within WPS and WRF.

2.2 SETUP AND CONFIGURATION OF WRF

The WRF codebase was obtained through its official GitHub repository (<https://github.com/wrf-model>), including WPS, the WRF model itself, and WRFDA and placed on LLSC. Before each component of WRF could be compiled, additional supporting libraries of MPICH2, NetCDF, JasPer, libpng, and zlib were all obtained. After these dependencies were installed, each component of WRF were individually built following the instructions described in the WRF User's Guide (the latest version can be found at https://www2.mmm.ucar.edu/wrf/users/docs/user_guide_v4/contents.html) using the Distributed-Memory Parallelism (DMPAr) option, so that it could run efficiently using multiple cores on LLSC. A simplified diagram of the WRF system is provided in Figure 9, showing each component and all inputs and outputs. Note that many of the individual processes and intermediate steps are not shown for brevity.




Figure 9. Simplified diagram showing flow of WRF system.

Once built and all the high-resolution terrestrial data were obtained, gridded meteorological data are needed to provide initial and boundary conditions. The gridded meteorological data commonly stem from other numerical weather model output, such as the Global Forecasting System (GFS) or the High-Resolution Rapid Refresh (HRRR) model, which provide information on the general weather conditions in the forecast domain as well as any changing weather (e.g., fronts, precipitation, etc.) that enter the boundaries of the domain during the run. If the high-resolution model were cycling constantly, such as once per hour, then forecasts from previous WRF model runs can serve to provide the initial conditions for the next model run. For example, a 1-hour forecast from a WRF forecast initialized at 12:00 UTC can be used to initialize conditions at 13:00 UTC. Within this program, the capability to use either GFS or HRRR model output was established to initialize the WRF model. The GFS was chosen since it provides forecasts across the globe so that a high-resolution domain could be set up anywhere, and the HRRR is primarily used to initialize the WRF model in the United States where it has coverage and since it is the highest-resolution operational forecast. Archived GFS data are archived on LLSC back to 2016 and is continually updated, while HRRR archives are sourced from the HRRR archive maintained by the University of Utah. For real-time runs, the HRRR data are downloaded from the NCEP website where files from the most recent two days are stored.

After the geophysical data has been sourced, the WPS is run to set up the domains and resample the terrestrial and meteorological data to the desired grids. The desired grids are set up by specifying a center reference latitude, longitude, and type of map projection used along with other properties of the map projection. Since many of the simulations will occur within the contiguous United States at mid-latitudes, a Lambert projection was chosen with a standard longitude equal to the center reference longitude. While it is possible to set up only one grid for WRF, it is best to set up multiple nested domains, of equal sized grid dimensions, for high-resolution simulations, each at increasing resolution so that atmospheric processes are better resolved and fed into the smaller inner domains. Since the inner most resolution of our simulation is 90-m, larger outer domains are also set up at 450-m and 2250-m resolution, all centered on the same location. An example of the relationship in size is shown in Figure 10. Example nested WRF domains (D01, D02, D03) centered around Boston. The geographical dimensions of a domain are three times larger than its next inner domain.

Thus in our implementation, each outer domain's geographical dimensions are three times larger than its nested domain. This ensures an outer domain fully encompasses its inner domain and can adequately spin up atmospheric processes advecting between nested grid boundaries. With the WPS parameter file set up, it can be run to produce the initial and boundary conditions for WRF.




Figure 10. Example nested WRF domains (D01, D02, D03) centered around Boston. The geographical dimensions of a domain are three times larger than its next inner domain.

Optionally, the initial conditions for WRF can be updated by assimilating observations using WRFDA, yielding a more accurate estimate of the initial conditions as long as the observations are of high quality. Many types of observations can be assimilated, including surface, aircraft, weather radar, satellite, radiometer, and profiler observations, among others. While meteorological drone observations are assimilated using WRFDA during this program for a select forecast to show the capability and impact as described later, these observations were not available for many of the high-resolution modeling forecasts that were made. For all other high-resolution modeling results shown, no data were assimilated into WRF. Rather, the forecasts were run at a higher resolution (90-m) than the operational HRRR (3 km) to demonstrate the capability and resolve small-scale atmospheric processes and turbulence that cannot be resolved by a coarser NWP model.

Once the initial conditions are finalized either after the completion of the WPS or WRFDA, the WRF model itself can be run once its parameter file has been correctly configured. The parameter file for WRF contains many options, several of which must match the forecast time and domain parameters set for WPS.

Additionally, there are many more options such as: forecast interval, number of forecast levels, time step, several atmospheric radiation options, and numerous physics and dynamics options including parametrizations of processes that are not completely resolved depending on the resolution (e.g., cumulus cloud processes, diffusion, boundary-layer processes, gravity wave drag, etc.). All of these options are described in the WRF Users Guide, and the parameterizations need to be carefully chosen as there are known incompatibilities. Since WRF is being run over three different resolutions, different parameterization options were chosen for each nested domain considering the processes that are implicitly resolved. Additionally, it took considerable time and many trials to identify a suitable suite of parameterizations to use, as numerical instabilities of flow over steep terrain with certain sets of options led to WRF crashing. Eventually, a final set of parameters were determined that are appropriate for the resolution, correctly simulated intended features such as topographic shading, and mitigated numerical instabilities at such high resolutions.

Due to the significant amount of human effort to set up the parameter files and run the WRF system (of which many of the individual steps were not described for brevity), wrapper code was written in Python to automatically generate the appropriate WPS and WRF parameter files and sequentially run all of the processes to make a forecast. The Python wrapper also downloads any HRRR or GFS data that is necessary for the selected forecast, and moves files and creates symbolic links as necessary. The only inputs into the Python wrapper are: domain and test names, center latitude and longitude, resolution of the inner domain, east-west and north-south extent of the inner high-resolution domain in meters, model initialization time, forecast interval, and forecast horizon. With only a small number of variables that a user needs to change, a high-resolution version of the WRF can be easily set up and run over any domain and time period of interest for which there are archived gridded meteorological data. In future work, this framework could be utilized to continually cycle over an area of interest to provide real-time forecasts.

2.2.1 WRF Execution and Postprocessing

After the parameter files are set up, WRF is run by executing a series of programs to process the output from WPS and generate the weather forecast. Running WRF is computationally expensive, particularly when operating at a high-resolution. For one of the tests run, to generate a 24-hr forecast at 100-m resolution covering an 80 by 80 km area (801x801x40 grid points) with model forecasts saved at 5-min forecast intervals, it took approximately 10,000 core hours, using 512 xeon-e5 slots on LLSC. This took 19.5 hours of wall clock time. Simply saving the data at each forecast interval took approximately 18 seconds to complete. While it is currently not practical to operationally run a rapidly updating model at such a high-resolution, this model could be run at a different resolution or cover a smaller domain and make execution practical depending on the end users' needs.

To cover a given domain, the computational resources necessary typically scales by three orders of magnitude for each order of magnitude improvement in spatial resolution. This is due to the fact that the number of grid points increases respectively in the meridional and zonal direction and the time step is directly proportional to the resolution to maintain computational stability. Thus, to cover the same 80 by 80 km area at 1 km resolution, instead of 100-m, a forecast would only take about 10 core hours. Similarly,

an acceptable tradeoff between resolution and time-to-completion might be running the model at 250-m resolution out to an 8-hour forecast. Using the same computing resources mentioned above, it is estimated that this could be performed in about one hour of wall clock time. This fine-scale resolution would still capture many flow-features for UAS and Advanced Air Mobility (AAM) operations but could be run in an operational manner. Going forward, the trade space of aerial coverage, spatial resolution, forecast output interval, and forecast horizon will need to be evaluated to determine the best compromise to provide operationally-relevant data to end users depending on the application and mission needs.

As the WRF model runs and generates forecasts at a specified interval (nominally every 5 min for this program) in either NetCDF or GRIB2 format, the output can be visualized, postprocessed, and used for other downstream purposes leveraging open-source software. During this effort, the model output was read into Python, visualized, and resampled to specified altitude layers using the open-source wrf-python package. Many other utilities are also freely available to display and read into the output in a variety of languages for use in other downstream systems.

2.2.2 Sample High-Resolution Model Run Output

To demonstrate the high-resolution WRF capabilities, the model was run over different domains of interest that are applicable to relevant mission areas. One of the primary application fields is aviation, particularly the emerging UAS and AAM fields where low-altitude wind and turbulence information are critically important given the slow vehicle airspeeds and battery constraints. In one of the initial tests of the high-resolution WRF model, a domain was set up centered just north of New York City on the Hudson River Valley area. This location was chosen due to the large number of airports in the region for traditional aviation applications, as well as anticipated demand for AAM activities. The weather in this area is also challenging, particularly with the steep terrain surrounding the Hudson River. A snapshot comparing output of the HRRR and the high-resolution WRF models, run at 90-m, is shown in Figure 11. For this model run, the HRRR output was used as the initial and boundary conditions, and both forecasts were initialized as 00:00 UTC on January 30, 2020. The relative coarse resolution of the HRRR is apparent when focusing on such a localized region. This example demonstrates that the high-resolution WRF is able to resolve many flow features that are not fully, or even partially, captured by the HRRR. This is particularly evident along the New York and New Jersey border, where the Hudson River Valley is located. The high-resolution WRF shows much stronger winds in this region, due to channeling of the flow through the valley and the smoothness of the surface over the water. While there are unfortunately no observations in this region for validation, the results are physically realistic and justified. Additionally, other flow features such as acceleration over mountains in the northwest region of the domain, a deceleration over the Tappan Zee Bridge (officially named the Governor Mario M. Cuomo Bridge) from flow obstruction, and turbulent streaks aligned with the flow (streamwise).




Figure 11. Comparison of HRRR and 90-m WRF wind forecast at 10-m around the Hudson River Valley region.

The high-resolution WRF was also run over the Boston area for a summer case, to demonstrate its capability to resolve local sea breezes. The sea breezes are generated on days when there is a strong temperature differential between the warm land and cool ocean, resulting in the formation of a direct circulation where the denser cooler air over the ocean pushes inland displacing the warmer air. When this happens, a region of low-level convergence on the sea breeze leads to strong coherent rising motions. This can be seen in the high-resolution model output in Figure 12. In this figure, wind barbs are plotted at 400 ft, which is a potentially important flight level for UAS operations. The vertical velocities are also useful for UAS operations for energy management, as UASs may be able to conserve energy along coherent updrafts. Additionally, the vertical motions indicate the strength of turbulence, important for vehicle control. Within Figure 12, the sea breeze is evident by the rising motions indicated in red where the south-easterly winds predominantly over the ocean interface with the south-westerly winds over the land. The turbulence, indicated by the amount the vertical velocity varies in a local area, is also much stronger to the west of the sea breeze and weaker to its east. This is due to the fact that the planetary boundary layer is stable east of the sea breeze, as the air near the surface is substantially cooler than the air above it. This is particularly evident over the water, where the vertical velocities are very small.




Figure 12. Winds at 400 ft with vertical motion shaded for a sea breeze over Boston on July 29, 2020.

Historically, low stratus clouds and fog have been extremely difficult to forecast in the San Francisco Bay area leading to significant challenges for air traffic management at San Francisco International Airport (SFO). The aerial coverage and height of these clouds has implications for the rate at which takeoffs and landings can occur at SFO. To provide better guidance to air traffic managers, the SFO Marine Stratus Forecast System was developed in the 2000s by MIT Lincoln Laboratory to provide short-term forecasts of cloudiness at different waypoints along typical landing paths [1]. This tool has been extensively used since its deployment, but hardware and instrumentation used in the current system are degrading and the system will need to be overhauled in the near future. New technology that has been developed since this system was initially commissioned will be explored to enhance the capabilities and increase confidence in the forecast going forward. One forecast solution could be real-time high resolution modeling, similarly to the capability developed here. To investigate this capability, the high-resolution WRF was also run over a domain centered on SFO on a day (July 31, 2020) where marine stratus was apparent in the morning and gradually burned off throughout the day. A snapshot of the simulated cloud cover, along with HRRR forecasted cloud cover and a satellite image, are shown in Figure 13. Compared to the HRRR forecast, the high-resolution WRF can capture fine-scale details such as the absence of clouds over the Santa Cruz Mountains south of the San Francisco Bay, likely due to the fact that the mountains are higher than the marine layer where the clouds are embedded. Conversely, the HRRR depicts a broad gradient of

cloudiness over the ocean and clear skies inland, with the cloud cover decreasing roughly linearly between the two regions. While the high-resolution WRF generally over-predicts the cloud cover slightly, it can predict localized clear areas not resolved by the HRRR. With additional tuning of model parametrizations and observations assimilated into the WRF for a more accurate representation of the initial conditions, the high-resolution WRF model could be considered as a component of an updated Marine Stratus Forecast System.




Figure 13. Modeled cloud cover from the HRRR, high-resolution WRF, and observed visible satellite over San Francisco on the morning of July 31, 2020. The WRF and HRRR forecasts are 4-hr forecasts initialized at 12:00 UTC.

Another potential application area for high-resolution weather forecasting is supporting wildfire suppression operations. Strong winds, which can be highly localized due to orographic effects or convective outflows, coupled with low relative humidity, are the primary drivers for rapid wildfire growth and spread. Accurate forecasts are crucial for prediction of conditions in which wildfires can start, such as high winds that may knock down power lines, as well as assessing the propagation speed and direction that wildfire fronts spread. Accurate prediction of wind shifts is especially important, when the predominant wind direction changes and fans the flames in a different direction and can endanger firefighters. To demonstrate the potential of leveraging these high-resolution simulation capabilities for wildfire response purposes, the model was also run centered on Paradise, CA on November 8, 2018. On this morning around the time of the snapshot shown in Figure 14, the Camp Fire was started underneath power transmission lines by hot katabatic winds flowing down the mountain slope and rapidly pushed the wildfire into the town of Paradise. These winds are depicted in the simulation where they are seen to be highly localized. The strong

northeasterly winds in the center of the domain, locally exceeding 40 knots, are the katabatic winds that quickly spread the wildfire. These results clearly show that this modeling capability can be used to inform decision making through mitigation of fire initiation when the local conditions are ripe for spread, as well as responding to active wildfires in effective and safe ways. In the future, the open-source WRF-Fire module could be incorporated into this capability if it were to be developed further. This module simulates fires themselves, including a prediction of the spread, and also how the weather itself is altered by the massive heat release from burning vegetation.




Figure 14. Wind forecast centered on Paradise, CA on November 8, 2018. The forecast is valid at 13:15 UTC, which is 75 minutes after it was initialized at 12:00 UTC.

This page intentionally left blank.

3. METEOROLOGICAL FLIGHT TESTING

Flight test were to be conducted using the platform described in Section 1 to collect meteorological samples, which in return would be assimilated into the high-resolution forecast system described in Section 2. The ultimate goal was to demonstrate that including additional meteorological measurements within the atmospheric boundary layer in region where in situ measurements are lacking would provide improved forecast of wind, temperature, humidity, and other properties in regions fundamentally downwind of the sample locations, which could ultimately lead to more accurate forecasts of storms and other high-impact weather. The original intent of the program was to fabricate as many as 10 meteorological sensing platforms, obtain or operate in one or more COAs in the Northeast US where the platforms could operate simultaneously while collecting data up to 10,000 feet. A number of issues thwarted this goal, some of which are cited above.

Instead, flight tests were planned to be conducted using the four MTD-2 platforms from the small Unmanned Aircraft System (sUAS) Testbed where we would conduct our tests at the same time another program was using the assets. Modification of these platforms for our intended purposes would not affect the operation of the platforms for the other tests. Fortunately, the design of the flight tests themselves in terms of simultaneous operations, altitudes, and routes were easily developed to satisfy both missions without any compromises to achieving either program's mission goals.

The coordinated flight tests were planned to be conducted at Fort Devens May 26-29, 2020: see Figure 15. In the shakedown prior to this week, significant data communication issues between the platforms and the grounds stations continued to plague the systems. Command and control communications were not consistent amongst the platforms and none showed the expected performance demonstrated in previous flight tests. At the end of the shakedowns, communications were considered performant enough to proceed with the flight testing with better communications expected in the open field environment, as compared to the electronically noisy environment inside the laboratory.



Figure 15. Flight testing site at Fort Devens, Devens, MA.

The data communication woes exhibited in the laboratory unfortunately continued to occur in the field. Flight trials over the first several days in a number of cases led to cases of loss of control (LOC) of some of the platforms. On several occasions, these LOCs led to significant damage to a platform. An example of an outcome from a complete LOC situation is presented in Figure 16.



Figure 16. Negative effect of loss of control due to communications failure. Complete separation of motors from wing mounts and cracked fuselage. Creative overnight reconstruction produced a flyable platform.

Fortunately, in many cases, including the aftermath shown in Figure 16, many platforms could be returned to service when enough adhesive and spare parts were available. On some occasions, that was not the case and components from an inoperable platform needed to be salvaged to put another platform into service. Fortunately, all members of both teams were skilled at field repairs, as is evident in Figure 17.



Figure 17. Field replacement of platform motors and power control units. Customized airframe modification was necessary for installation of parts not originally intended for unit.

Because of the continued data communication issues occurring in the field, and the attrition of platforms, no germane weather data were collected until the last day of flight testing. All the original coordinated flight plans were appended at that point. For the last intended flight of the campaign (see Figure 18), the flight plan that took the platform to the peak altitude for the restricted airspace of 3995 feet was employed. Fortunately, there were no communications issues with this flight and a valid set of water vapor and temperature data were collected along the ascent and descent. The flight path remained within a $\frac{1}{4}$ square nautical mile region and the total flight duration was 11 minutes.



Figure 18. Successful meteorological data collection flight. platform proceeded to 3995 ft, which was the maximum allowable flight altitude for the restricted airspace R-4102.

The processed data from the flight are presented in Figure 19. Note the hysteresis present in the profile data. The missing data in the traces near the lower pressure levels (higher altitudes) are occurrences where the platform passed through clouds and the water vapor sensor became saturated. That particular sensor cannot function at 100% humidity as water condenses on the sensor and gives false readings.




Figure 19. Meteorological data, water vapor and static temperature, as recorded at varying pressures (altitudes) along the platform's ascent and descent.

Even though the quantity and range of data collected is significantly less than what was anticipated, the one collection made still provided enough information to make it suitable for later data assimilation (DA) into the high resolution forecasting system.

4. ASSIMILATION OF DRONE MEASUREMENTS INTO WRF

As discussed earlier, the WRF system is composed of a number of subsystems that each perform different processing, culminating in running the WRF model itself to produce numerical weather prediction forecasts. Observations are typically assimilated into the model through WRFDA, in which the observations are compared with the initial conditions for quality control checks to ensure the data is at least realistic. Following this quality control check, the initial conditions of the model are updated to better resemble the observed atmospheric state. The effect of the meteorological observations is over a volume, with the greatest effect being where the observation was taken and smoothly decreasing to the edge of the volume. The horizontal size and vertical depth of the volume are determined based on length scales of the observation type (i.e., wind speed, wind direction, temperature, humidity) being assimilated. These length scales are calculated from previous WRF runs over approximately the prior month in the domain of interest, thus they are climatologically representative of the typical scales in that region and time of year. After the initial conditions are updated through this DA, the boundary conditions of all nested domains are also updated and are used to run the model. While DA can be performed in three- or four-dimensions (including time), only three-dimensional DA was used here since the drone data were all collected during a short time window (≈ 10 minutes) and three-dimensional is much less computationally intensive, thus quicker to run. If observations taken over a longer period of time (several hours) were ingested, four-dimensional DA would be more appropriate to use.

Before the drone meteorological observations could be ingested into WRFDA, the data needed to be converted into a format that WRFDA could consume. The American Standard Code for Information Interchange ASCII-based LITTLE_R was the format chosen, as it was the easiest to work with and compatible with other community-developed software. The LITTLE_R format contains information on the latitude, longitude, observation type, elevation, date, time, atmospheric pressure, and relevant meteorological quantities along with quality control flag fields. Most of these fields are simply filled in using the drone observations, and translated into LITTLE_R format using a Python script developed in this program. Because drone observations are new and rarely assimilated into WRF (which is why this activity is novel and valuable), using LITTLE_R there was no way to indicate that the data were collected onboard a drone. Instead, the observations were treated as an Automated Meteorological Data Relay (AMDAR) data type, which is an observation reported from a commercial aircraft and is the most similar type supported. The data type determines how the observations are used in WRFDA, including the typical error characteristics. In the future, as drone meteorological measurements become more commonly collected, it will become necessary to create a data type that is particular to drones and supported in WRFDA to correctly ingest these data into the model.

Once the meteorological drone observations were converted into LITTLE_R format, the data were assimilated into WRFDA. During the test flights on May 29, the data collected to a height of up to 1217-m between 15:45-15:53 UTC was used for the DA demonstration. Using the three-dimensional variation (3DVAR) analysis scheme in WRFDA, these observations of temperature and humidity were assimilated

into a forecast initialized at 16:00 UTC using the HRRR analysis at that time as the initial background field. The DA is only conducted over the coarse 2.25 km resolution outer domain. (Note, it was too computationally expensive, given time constraints in this program, to run 90-m simulations over a forecasted month, which is necessary to build background error statistics that can improve the accuracy of running WRFDA at such high resolutions).

The results of this DA on the temperature and humidity profiles centered on the flight profile are shown in Figure 20. Here, the drone observations are shown from both the ascent and decent of the aircraft, as well as the profiles before and after the DA. Considerable spread exists between the ascent and decent portions of the flight. While a small part of this spread may be due to the atmospheric state evolving during the flight, most of this difference can be largely explained by the hysteresis of the sensor given its response time of several seconds. As the aircraft quickly passes through different altitudes, it takes time for the sensor to respond and sense these changes. Because temperature decreases and humidity increases with height in the convective daytime boundary layer, this leads to a warm dry bias during the ascent and a cool moist bias during the decent. By assimilating all of these observations during the ascent and decent portions into WRF, the bias is expected to average out in the resulting profile. This can be seen in the resulting post-DA profiles, which mostly lie between the ascent and decent drone observations. Assimilating these observations with WRFDA generally resulted in local temperature and humidity profiles more closely resembling the data compared the pre-DA initial conditions, with generally a warmer and drier profile throughout the boundary layer. Assuming that the measurements themselves are accurate and free from significant bias, this result is expected and shows that the DA is making the correct local changes to the initial conditions to better represent the true atmospheric state.




Figure 20. Comparison of the temperature and humidity profiles over Fort Devens, MA after assimilating meteorological drone observations using WRFDA.

As mentioned earlier, the DA does not only affect the initial conditions where the observations are taken, but instead over a wider volume that is related to typical length scales in the domain. The volumetric effect of the drone observations assimilated into the coarse 2.25 km resolution domain is shown in Figure 21, where the temperature differences in initial conditions before and after the DA are visualized. Despite the drone observations only being collected within a cylindrical volume less than 1.3 km high and 1 km wide, the observations influence a wide wedding-cake shaped volume almost 200 km in radius horizontally and 5 km vertically. This is due to the large length scales of temperature and humidity in the area, particularly in the horizontal. If the observations were to be assimilated into the high-resolution nested domain in the future, the effect is expected to be much more localized. Still, the most significant impact of the DA is where the profiles were conducted.




Figure 21. Temperature difference of the model initial conditions of before and after assimilating the meteorological drone observations using WRFDA.

After the initial conditions were updated through WRFDA, the high-resolution model was run to produce an 18-hour forecast. To evaluate the effect of the DA on the forecast, the temperature and humidity forecasts were compared with local observations in the area at varying time horizons. While aircraft-based observations were considered for validation, significant gaps in the availability of the data and sparseness, partially due to reduced air traffic during the COVID-19 pandemic, made such a comparison to produce statistically robust results difficult. Instead, the model output was validated at several Automated Surface Observing System (ASOS) sites in the area.

In Figure 22, the temperature and dew point measured at the closest ASOS site to Fort Devens, located at Fitchburg Municipal Airport (KFIT), are compared to forecasts at that location from HRRR, WRF with

DA, and WRF without DA. This comparison shows mixed results in forecast accuracy from the high-resolution WRF alone compared to the HRRR, as well as from the DA into WRF. In comparing the HRRR and WRF (no DA) forecasts, it is clear that the HRRR more accurately predicts the temperature over the course of the day but the HRRR has a higher bias in the dew point throughout the day than the WRF without DA. Considering the additional DA of the drone measurements, the high-bias in the WRF forecast of temperature becomes even larger while the dew point forecast becomes lower and more accurate. This is likely due the fact that the drone observations led to warmer and drier initial conditions near the surface after DA, as was shown in Figure 20, resulting in a warmer and drier forecast. Thus the high-resolution WRF forecasted near-surface moisture more accurately than the HRRR or WRF without DA, but the temperature forecast was degraded. Within Figure 22, the two WRF forecasts become more similar with time as the volume of air influenced by the DA gradually advects out of the region.

Due to time and resource constraints of this program, there was insufficient time to properly calibrate the temperature and humidity sensors that were instrumented on the aircraft. It is suspected that the high-temperature bias in the model forecast is a direct result of a bias in the HTY 271 sensor itself, as is common with this type of instrumentation. Before future data collections, any sensors should be properly calibrated to ensure the measurements are accurate and will not lead to a degradation of a forecast when assimilated into a model. But these results successfully demonstrate the feasibility of assimilating novel UAS-based atmospheric measurements into a high resolution research NWP, which is a major accomplishment in and of itself.




Figure 22. ASOS temperature and dew point one minute observations (OMO) taken at KFIT compared to forecasts from HRRR, WRF, and WRF run with data assimilated from drone observations (DO) at the same location.

This page intentionally left blank.

5. CONCLUSION

The program was a success in a number of particular avenues. It furthered the knowledge and experience of Division 4 personnel in operating small drones in BVLOS operations. This knowledge and experience can be leveraged in future programs requiring skills and expertise in such domains. The actions resulting in the reinterpretation of the DoD memorandum limiting the acquisition and use of drones considerably improves the flexibility to develop future UAS platforms and will make such development more expeditious across the Laboratory.

Of particular benefit resulting from this program was the establishment of a ready capability to produce high-resolution numerical weather forecasts over any region of interest on demand and demonstrating the feasibility of ingesting UAS-derived atmospheric measurements into the model. There are a multitude of programs across the Laboratory that could leverage such a capability, particularly in Division 4. While there is plenty of future work that could and should be applied toward improving the capability, its current incarnation, and the practical experience gained by developing will be applicable and beneficial to many future programs, including emerging areas such as Advanced Air Mobility. Efforts should be made to modify the Division's existing sUAS inventory to correct the deficiencies in magnetic heading sensing. This will support future sUAS meteorological research and other programs utilizing these platforms that depend on accurate pointing such as surveillance, search and rescue, and wild fire detection and monitoring.

This page intentionally left blank.

GLOSSARY

3DVAR	Three-Dimensional Variation
AAM	Advanced Air Mobility
ACAS Xu	Alert and Collision Avoidance System X unmanned
AMDAR	Automated Meteorological Data Relay
ASOS	Automated Surface Observing System
BVLOS	Beyond Visual Line Of Sight
COA	Certificate of Authorization
COVID	Corona Virus Disease
DA	Data Assimilation
DMPar	Distributed-Memory Parallelism
DO	Drone Observations
DoD	Department of Defense
DTED	Digital Terrain Elevation Data
FAA	Federal Aviation Administration
GFS	Global Forecasting System
HRRR	High-Resolution Rapid Refresh
KFIT	Fitchburg Municipal Airport
LLSC	Lincoln Laboratory Supercomputing Center
LOC	Loss Of Control
MIT	Massachusetts Institute of Technology
MODIS	Moderate Resolution Imaging Spectroradiometer
NCAR	National Center for Atmospheric Research
NCEP	National Centers for Environmental Prediction
NLCD	National Land Cover Database
NOAA	National Oceanic and Atmospheric Administration
NWP	Numerical Weather Prediction
OMO	One Minute Observations
SFO	San Francisco International Airport
SPI	Serial Peripheral Interface
sUAS	small Unmanned Aircraft System
UAS	Unmanned Aerial System
WPS	WRF Preprocessing System
WRF	Weather Research and Forecasting
WRFDA	WRF Data Assimilation

This page intentionally left blank.

REFERENCES

1. Reynolds, D.W., Clark, D.A., Wilson, F.W. and Cook, L., 2012. Forecast-based decision support for San Francisco International Airport: A NextGen prototype system that improves operations during summer stratus season. *Bulletin of the American Meteorological Society*, 93(10), pp.1503-1518.