Enhancing Autonomous Satellite Communication Systems with Weather-Aware Scheduling and Reconfiguration

Aaron Smith
Cognitive Signal Processing Branch
NASA Glenn Research Center
Cleveland, OH
aaron.smith@nasa.gov

Elmer Weston Brown
Electrical, Computer, and Systems Engineering
Case Western Reserve University
Cleveland, OH
ewb4@alumni.cwru.edu

Adam Gannon
Cognitive Signal Processing Branch
NASA Glenn Research Center
Cleveland, OH
adam.gannon@nasa.gov

Francis Merat
Electrical, Computer, and Systems Engineering
Case Western Reserve University
Cleveland, OH
flm@case.edu

***Abstract*—NASA currently provides communication support to over 100 satellite missions, and the agency is driving developments in Ka-band communications and network management automation to support additional future missions. At Ka-band frequencies, rain can degrade a communication link by more than 10 dB, which may be mitigated by agile scheduling and data rate control. We present a weather forecasting module for Ka-band communications that is intended to be used in an autonomous network management service that employs machine-to-machine scheduling systems for dynamic user access opportunities. Our forecasting module (NIMBUS) runs on AWS Cloud, consumes the freely and publicly available NOAA MRMS precipitation rate dataset (1km x 1km x 2-min), produces 30-minute Nowcasts using the pySTEPS algorithm, and publishes high-level ground station specific link quality predictions. We evaluate two potential NIMBUS outputs, a binary classifier that predicts rain attenuation greater than 3 dB and a rain attenuation estimator, and we backtest these outputs using one year of power measurement data collected from observations of the geostationary ANIK F2 satellite's Ka-band beacon. We report, with a 30-minute lead time, a binary classifier accuracy of 84% and an estimator RMSE of 1.67 dB. Additionally, we discuss how the NIMBUS module could be incorporated into a user-initiated service framework to enable weather-aware scheduling and reconfiguration. Integrating NIMBUS into a user-initiated service framework provides an alternative to static link budget padding based on a statistical long-term rain exceedance rate attenuation, and leverages advances in software defined radio, autonomous machine-to-machine scheduling, and precipitation Nowcasting.**

***Index Terms*—Satellite communication, scheduling, automation, precipitation nowcasting**

I. Introduction

NASA provides communication support to many satellite missions, from planning to operations. Generally, the spacecraft needs to transmit science data from on-board sensors back to Earth. Ground stations often employ large high gain antenna that point to and track spacecraft as they arc across the sky. These highly directional systems help increase the signal to noise ratio and thus increase the potential data rate that can be achieved over that link, while minimizing the size, weight, and power of the spacecraft's terminal. Therefor, ground stations generally communicate with only one satellite at a time and assets must be scheduled in advance. Scheduling can be a complex non-linear optimization problem, which is why human operators are often involved. Similarly, when a schedule fails to satisfy the service needs of a satellite, humans are tasked with identifying a remedy, which can backlog network operators if it happens too frequently. To circumvent this, communication services are generously padded with excess capacity.

Ka-band communication systems are becoming more common due to spectral congestion in the lower frequency bands and the higher data rates that can be achieved. However, Ka-band communication systems are also susceptible to rain attenuation, which can exceed 10 dB. To mitigate this, ground station site selection involves measuring, or modeling, the long-term weather statistics of prospective ground sites. For an established ground site, rain is generally accounted for by applying a static rain attenuation to a communication link budget. Setting a low exceedance rate reduces the number of rescheduling events that will be required during a year, but it also results in a reduced data rate. To meet science data volume

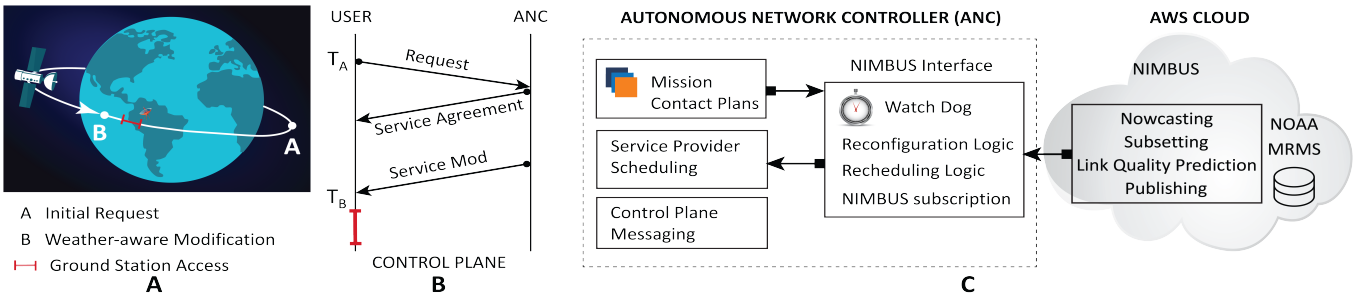

Fig. 1. (A) A user spacecraft requests service while at orbit position "A". When the spacecraft reaches position "B", the autonomous network controller retrieves high-level link quality predictions from the NIMBUS module and uses the expected link quality to reconfigure the data rate of the upcoming contact and schedules additional services as required. (B) A sequence diagram showing the control plane messages that are passed prior to the ground station contact. (C) A system diagram of an autonomous network controller that is subscribed to the NIMBUS module, which runs independently on AWS Cloud.

requirements, the mission may need to resize the satellite's communication terminal, or schedule additional ground station time.

An alternative method, in a perfect world, would be to communicate at the highest possible data rate during good weather, and then reduce the communication data rate only when the link is degraded. Additionally, an ideal system would only schedule as much time as is needed on the ground station, thus maximizing the efficiency of the ground asset. However, such a system would need to accurately predict the future channel capacity, assuming a constant coding and modulation scheme is used duration a ground station pass. Underestimating channel capacity is similar to setting a low exceedance rate and overestimating channel capacity creates rescheduling tasks.

There are several factors that make weather-aware scheduling and reconfiguration relevant today.

- Ka-band communications is becoming increasingly prominent [1], and it is prone to rain attenuation [2].
- Advances in Ka-band software defined radios provide spacecraft with an ability to rapidly alter their communication protocols and configurations [3], [4].
- NASA's user-initiated services (UIS) concept of operations [5]–[7] would support autonomous scheduling and reconfiguration and UIS has been included in NASA's LunaNet Specification [8].
- Weather related datasets are widely available [9] and precipitation Nowcasting has improved significantly [10]–[12].

In the following sections, we describe a weather mitigation system for Ka-band communications that predicts future rain rates and estimates rain attenuation. We provide an analysis of our algorithm by backtesting against historical precipitation data that is paired with Ka-band received power measurements over the same time period. In our discussions section, we use an example scenario to describe a concept of operations for integrating our algorithm into an autonomous scheduling agent.

II. Methodology

The NOAA-Informed Meteorological Broadcast Updater for Satellites (NIMBUS) module (see Fig. 1) was designed to provide high-level link quality predictions to an autonomous network scheduling agent that could rapidly update schedules and radio configurations without the need for human intervention. In this section, we describe the systems used to collect radar-based precipitation rate data, which would be used operationally, and the Ka-band power measurement data that was used to validate the algorithm. Following this, we describe the NIMBUS algorithm and our vision for an AWS Cloud-based implementation.

A. Data Acquisition

1) NOAA's MRMS: National Oceanic and Atmospheric Administration (NOAA)'s Multi-Radar Multi-Sensor System (MRMS) system integrates the data streams of multiple sources, including radars, satellites, and numerical weather prediction models. In this work, we use the radar-only precipitation rate [mm/hr] estimate. The precipitation rate values represent a surface area of 1km x 1km, and are updated every two minutes. In the operational tables [13], this is labeled "PrecipRate", which is category 6, parameter 1. These estimates are provided by stitching together the outputs of more than 158 WSR-88D radars, which are positioned across the continental United States. In future work, we expect to use the precipitation flag as well, which describes the precipitation type, i.e. convective, stratiform, tropical, hail, snow. The output files are packaged in the GRIB2 file format, with each file containing more than 24 million coordinate points.

2) Power Measurement and Attenuation: In 2009, NASA installed an interferometer near White Sands New Mexico to study the site's Ka-band tropospheric phase stability and rain attenuation [14]. The interferometer was designed to observe an unmodulated beacon signal broadcast from the geostationary ANIK F2 satellite. In this work, we used the complex I/Q data that was captured during 2015 to produce a received power time series. In Table I we provide the relevant link information.

Our received power time series contains a periodic variation of approximately 0.35 dB, with a 24-hour period. We expect

TABLE I
KA-BAND LINK MEASUREMENT PARAMETERS

Satellite	Name	ANIK F2
	Beacon Freq.	20.2 GHz
	Polarization	Linear
	Orbital Lon.	111.1°W
	Elevation	51.8°
	Azmuth	188.3°
Ground Site	Latitude	32.5423°N
	Longitude	106.6139°W
	Altitude	1.469 km

that this is due to the day/night cycle and variations in pointing error as the satellite has a small inclination angle and eccentricity. At Ka-band frequencies, rain attenuation is the dominant source of atmospheric attenuation. Therefore, we approximate rain attenuation by subtracting the received power level from a reference power level. The reference level was computed using a rolling median filter with a window length of seven days, which smoothed out day/night variation. The seven day period and the median operation allowed the reference to remain stable during periods of rain. This method of approximating rain attenuation ignores other potential contributions, such as polarization offset–which should be constant in this case, and other atmospheric attenuation sources, such as clouds and gas.

B. Data Sanitization and Alignment

In Fig. 2, we show the rain attenuation, after removing missing and/or erroneous data. For example, in October of 2015, severe rain caused damage to the ground station receiver's low-noise amplifier, this caused the digital receiver (DRX) to recorded abnormally low power measurements. The issue was resolved by the 2nd week of November. Additionally, there were brief periods of outage throughout the year, which needed to be removed from the dataset. A largely manual inspection of the data was used to isolate and remove outage periods.

The MRMS dataset from 2015 was collected from the Iowa State University repository. The grid coordinate points are defined using latitude and longitude. The definitions were shifted starting Jan 21 2015, and due to this, we only used the data from January 21st to the end of the year. Additionally, there are a number of missing files for this year (with each file representing a 2-min dataset). There were also several instances of -1 "missing" and -3 "no coverage". In our evaluation, we rejected all time steps that were invalid due to corrupt power measurements in the DRX or invalid due to issues with the MRMS data. The DRX data is collected at 1-sec intervals, while the MRMS data is collected at 2-min intervals. We resolved this by taking the mean power measurement value over 120 seconds.

C. NIMBUS

NIMBUS is fed the precipitation rate "PrecipRate" values from the MRMS GRIB2 dataset. Each pixel in this data represents a 1km x 1km area. Initial testing and evaluation was done using the historical data archived by Iowa State

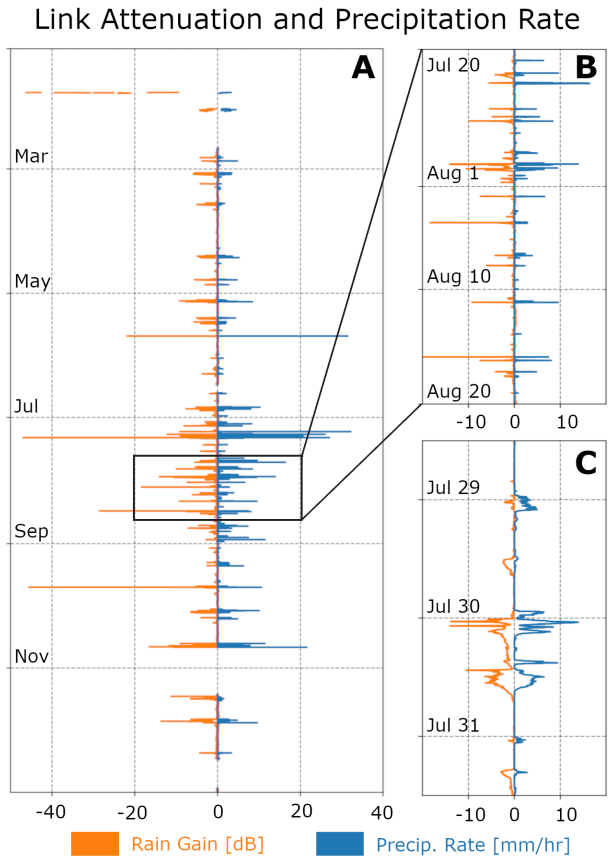

Fig. 2. (A) Link attenuation (rain gain) [dB] is overlaid on the same x-axis as the precipitation rate [mm/hr] to highlight the correlation between the two signals. The precipitation rate shown is the average taken over a 21km x 21km grid and the attenuation data was averaged over a 2-min period. Periods with missing and/or erroneous data have been removed. The data is for all of 2015. (B) This figure focuses on data between July 20 and August 20, 2015, which is a particularly rainy season at White Sands. (C) This figure focuses on four days worth of data at the end of July, 2015.

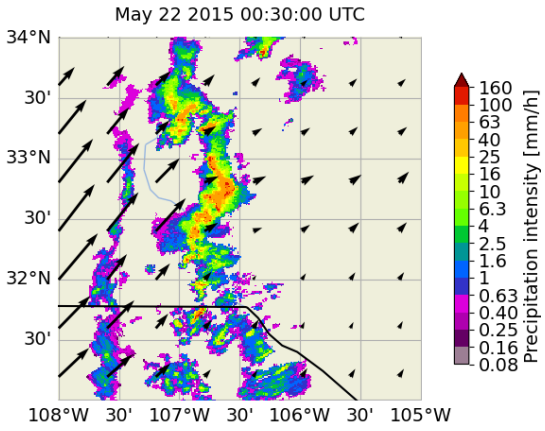

Fig. 3. This is an output from the pySTEPS Nowcast algorithm. Here, we have computed a motion field using the Lucas-Kanade optical flow algorithm. The field is visualized as black vectors over a 300km square area. The motion field is used to produce an advection-based Nowcast of the precipitation rate data.

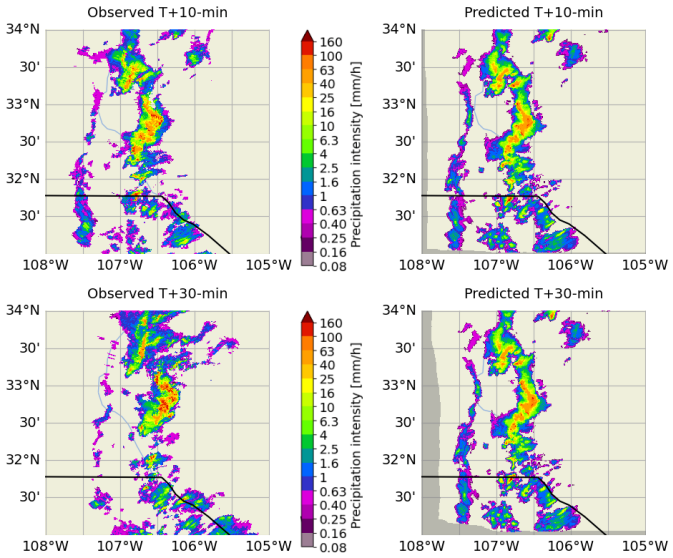

Fig. 4. The top row shows the actual 10-min evolution of precipitation (starting from Fig. 3), and the 10-min Nowcast output by the pySTEPS algorithm. The bottom row compares the 30-min evolution of precipitation with the 30-min Nowcast.

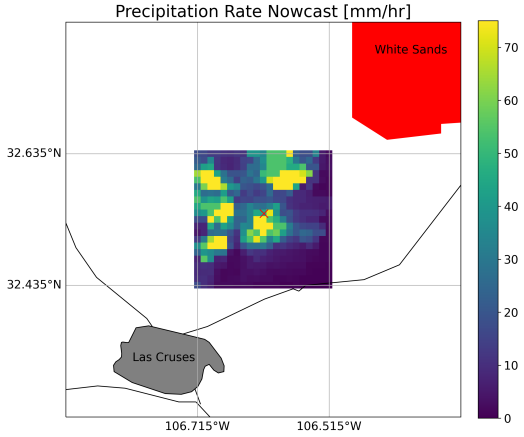

Fig. 5. After Nowcasting, a 21km square grid of precipitation rate values are extracted at each ground site of interest. This figure visualizes the subsetting operation of a Nowcast and overlays the grid on a map for scale. The 21x21km grid of rain data is updated every two minutes.

University at https://mtarchive.geol.iastate.edu/. We produce a Nowcast using pySTEPS [11], we compute a motion field via the Lucas-Kanade optical flow algorithm (see Fig. 3), and then extrapolate the most recent observation using a Lagrangian persistence method. This assumes that air parcels are constant and that change is due to advection. In Fig. 4, we visualize a Nowcast at two different lead times (10-min, and 30-min), projecting forward from May 22, 2015 00:30:00 UTC. For our single ground station example, we used a 300km square grid around our site of interest. The Nowcast area must be large enough to capture relevant precipitation data and motion vectors. In the future, we assume that we will be interested in sites across Continental United States (CONUS), and thus would run a single Nowcast on the entire CONUS dataset, before subsetting to select out small (21km square) grids for each individual ground station location. After computing a Nowcast of precipitation rates, we subset a 21km square grid (21x21 pixels) around each ground station site of interest– shown in Fig. 5. We then collapse the 2D matrix of Nowcast precipitation rates into a single scalar using the mean of the values. In our NIMBUS results section, we also analyze using a max function on this data. Finally, we use this value to generate two outputs 1) we use a binary classifier to predict whether or not the link will experience >3dB rain attenuation and 2) we convert the average rain rate to a scalar approximation of rain attenuation.

Our binary classifier is a supervised learning algorithm that categorizes the precipitate rate data as belonging to the class "0: the precipitation rate will not disturb the link by more than 3 dB" or "1: the precipitation rate will disturb the link by more than 3 dB". We provide an analysis in the results section based on the 2015 data. To estimate rain attenuation directly from the precipitation rate, we use ITU-R P.530-18 Section 2.4.1 "Long Term Statistics of Rain Attenuation", where we choose $R_{0.01}$ to be equal to the instantaneous precipitation rate.

D. AWS Cloud

NOAA's MRMS data products are now available on the Registry of Open Data on AWS and on AWS Data Exchange. Conveniently, the data architecture has leveraged SNS notifications, which provide subscribers with information about new objects when they are placed in an S3 bucket. In the case of MRMS, this occurs every two minutes. Hosting processes and data storage in the cloud provides scalability and transfers the burden of server management to the commercial entity. In this case, because NOAA's MRMS data products are already placed in an S3 bucket, they are automatically reachable from AWS compute options, including Elastic Cloud Compute (EC2), Lambda, and Elastic Kubernetes Service (EKS). If a weather prediction system, such as NIMBUS, is eventually scaled up to support a large fraction of NASA's mission portfolio, then leveraging a serverless cloud compute architecture may be practical. In this section, we describe the process of deploying NIMBUS in an AWS Cloud environment.

In Fig. 6, we provide a diagram of the NIMBUS module implemented in AWS Cloud. There are largely three main processes that take place. 1) NOAA deposits a set of MRMS data into an S3 bucket. The MRMS data contains many categories that are not used in this algorithm, but the action of filling the S3 bucket triggers an SNS topic notification that is passed to our Virtual Private Cloud (VPC), with allows us to act on the new data. 2) Our first Lambda compute process, which exists in our own VPC, is subscribed to the MRMS topic and thus can take action on the new precipitation rate data. As described in a previous section, we used the latest precipitation rate frames to compute a precipitation Nowcast for CONUS and we then place the predicted rain data into an S3 bucket. Fortunately, AWS Cloud also provides a documented path to

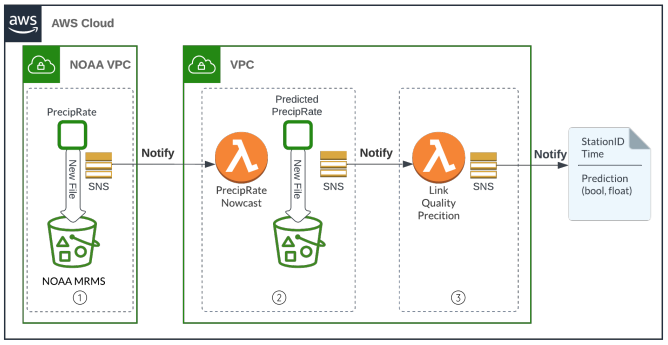

Fig. 6. A diagram of the NIMBUS MRMS data handling process in AWS Cloud. 1) Every two minutes, NOAA deposits a GRIB2 file containing the precipitation rate for CONUS into an S3 bucket, which triggers an SNS topic notification. 2) The four most recent precipitation rate GRIB2 files are processed by pySTEPS to produce a 30-min Nowcast. The Nowcast is placed in an S3 bucket, which triggers an SNS notification. 3) For each ground station, a 21km grid is subset from the precipitation rate Nowcast. The average of the 21km grid used as input by the binary classifier and the attenuation estimation function. The results for each ground site are transmitted to the automated scheduler via an SNS topic notification.

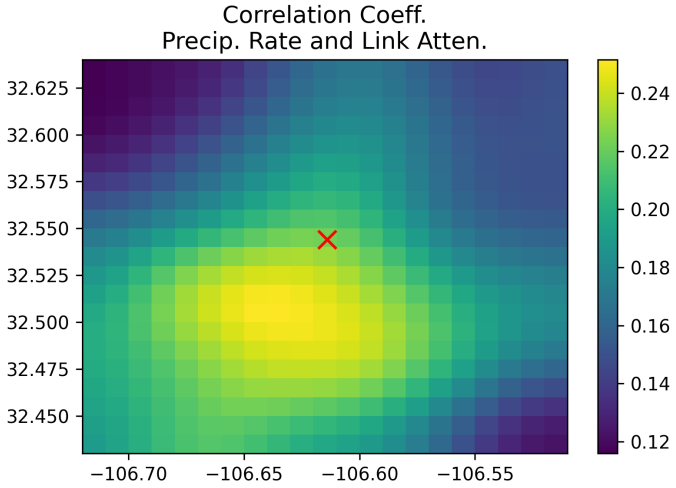

Fig. 7. The rain attenuation time series for 2015 was correlated with the MRMS precipitation rate of each pixel in a 21km square grid near the ground station at White Sands, which is denoted by the red "X".

enhance the base Lambda image with the GRIB2 processor program `wgrib2`[15] and the `pywgrib2_s.py`[16] Python interface. 3) The Link Quality Prediction Lambda is triggered by an SNS topic notification. This function is provided with a list of relevant ground station locations, so that it can subset out the 21km square areas that are used to provide localized link predictions. For each of these sites, an SNS topic notification is used to send the results to a subscriber, with each result containing a site identifier, the prediction time, and the prediction results.

III. Results

We evaluate the NIMBUS algorithm using power measurements collected during 2015 and the corresponding NOAA MRMS precipitation rate data for the same year. The DRX was located near White Sands NM, where it observed the 20.2 GHz continuous wave (CW) beacon transmitted by the Anik F2 satellite. We analyze the precipitation Nowcasts and the link quality predictions independently, and then analyze the end-to-end performance by predicting link quality based on Nowcast data.

A. Data Correlation

In Fig. 7, we visualize the correlation coefficient, for the rain attenuation and the rain rate data, for each pixel in a 21km square grid around the ground site near White Sands. The highest correlation coefficient is 0.238. This correlation was computed without any Nowcasting, or lead time. Given a mean rain height of 4.74km above mean sea level for this location, the peak correlation is further south and west than expected. However, this does highlight the spacial dependency of the rain data and future systems may want to consider the satellite trajectory during a pass.

B. Precipitation Nowcasting

To analyze quality of the Nowcast predictions, we use the Fractions Skill Score (FSS). This score measures the number of grid points that exceed a threshold error, which was set to be 1 mm/hr. Here, we use a grid size of 21km, to match the subset grids sizes that were used during the prediction step. In Fig. 8 we report the FSS scores for every two minute time step over a 24 hour period, for three different lead times. There are approximately 720 sample results for each lead time. In this work, we computed Nowcasts over a 300km square area around White Sands.

C. Link Quality Prediction

Our first prediction method was a binary classifier that predicted class labels based on the average rain rate over a 21km square grid at the ground site. The labels were designed to predict whether or not the link would be degraded by more than 3 dB, based on the 2015 data previously described. The integrated area under the receiver operating characteristic curve was 92%. For a false positive rate less than 5%, we choose a 0.414 mm/hr threshold on rain rate, which results in a true positive rate of 82.5%. The recall at this level is 76% and the precision was 31%. The total accuracy over one full year was 97%. 7.85% of the 2-min frames, over one year, contained some amount of rain in the 21km subset. Our second prediction method was an attenuation estimator. We generate instantaneous point estimates of attenuation from rain rate measurements using the long-term statistical attenuation methods in the ITU recommendations, by replacing the $R_{0.01}$ rain rate with the instantaneous measurement. These methods use the specific attenuation power-law model $\gamma_R = kR^{\alpha}$ described in ITU-R P.838-3. We compute the path length $d = 4.17$km using the mean annual rain height method described in ITU-R P.839-4. We evaluate $\gamma_R d$ as well as $\gamma_R d_{eff}$, where d_{eff} is computed using ITU-R P.530-18. Additionally, we compute

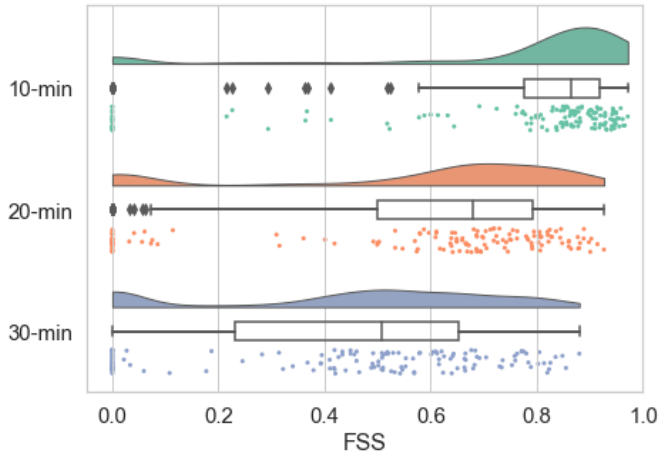

Fig. 8. The Fractions Skill Score was measured for Nowcast lead times of 10-min, 20-min, and 30-min. The FSS was computed based on a 21km grid. Nowcasts used the most recent four frames and were computed every two minutes for a 24-hour period.

the attenuation in a similar manner, but with the equations in ITU-R P.618-13 Section 2.2.1.1.

Both the binary classification and the attenuation estimation methods use a single scalar rain rate value to produce a prediction. We collapse the 21km square grid of values into a single scalar by averaging. Alternatively, we investigated taking the maximum rain rate, but this produced poorer performance in all cases. Additionally, we experimented with only using rain rates from grid coordinates with a high a correlation coefficient (see Fig. 7), but this also generally degraded the results.

The three different methods used to estimate rain attenuation produced similar results, with the method described in ITU-R P.530-18 being slightly better performing. The other methods were generally within one tenth of a dB in both mean absolute error (MAE) and root mean squared error (RMSE) error.

Using $\gamma_R d_{eff}$, we compute an MAE of 0.919 dB and an RMSE of 1.517 dB. This calculation was performed using one full year of data, but only for frames that contained some amount of rain. As previously stated, there was some amount of rain in 7.85% of the 2-min frames. If we included all frames, the MAE and RMSE would be significantly lower. There were several frames with abnormally large errors ($>$ 30 dB), which could have been due to equipment error. To mitigate this, we computed the MAE and RMSE over frames with errors that were in the 99% quantile, thus rejecting 1% as outliers.

D. Nowcasting + Prediction

Here, we analyze the end-to-end results for a particularly rainy day, July 20, 2015 (see Fig. 2C). First, we computed 30-min Nowcasts for every 2-min interval over the 24-hour period. Then, we used the Nowcast frames to produce a binary classification that predicts link degradation and an estimate of rain attenuation, as previously described. The binary classifier achieved an overall accuracy of 84%, a 91% recall, and 56% precision. The rain attenuation estimates over the 24-hour period had an MAE of 1.17 dB and an RMSE of 1.67 dB.

IV. Discussion

In this section we illustrate a possible system architecture with NIMBUS integrated into an operational system. As shown in Fig. 1, NIMBUS output is used to select a radio configuration which maximizes performance as well as schedule additional services as needed.

A. Basic Concept of Operations

We imagine a scenario in which a network management center performs the function of allocating Ka-band ground station passes to several Earth-orbiting spacecraft. We assume each spacecraft is compatible with several commercial providers of ground station service. While providers are capable of accepting requests for service with minimal lead time (15 minutes or less) based on availability, we assume nominal ground station contacts are planned out in a forecast schedule days to weeks in advance. At this time horizon, no real-time weather information is taken into account during scheduling and rain attenuation is expected to be minimal. Spacecraft radios and the corresponding modems at ground station sites can operate at several data rate modes (e.g. low, medium, and high rate) but maintain a constant rate during a pass. Typical of near-Earth missions, spacecraft transmitters remain off between passes but their receivers are always listening for commands.

An automated monitoring process at the network management center uses NIMBUS to evaluate each scheduled pass as it comes within the 30-minute Nowcast window. Contacts for which the classifier predicts negligible ($<$ 3dB) disturbance due to weather are left to execute as-is. If significant ($\geq$ 3dB) link disruption due to weather is predicted the monitoring process can take several actions. First, commands are sent to the spacecraft and ground station modem to lower data rate. Increased link margin at lower rate will increase likelihood some data is successfully transferred over the pass, as opposed to the link failing entirely. This will result in less data transferred over the pass than originally anticipated. Second, for high-priority data transfers a backup pass is immediately scheduled in the near future with another ground station. The duration of the backup pass is calculated to account for the expected remaining data volume after the original pass is completed at the lower rate. Using automation this backup can be scheduled without incurring the delay of a human mission operator reacting to data loss during a degraded pass. Given sufficient ground station availability, this backup pass can take place even during the same orbit.

B. Enhanced Operations

Additional options exist if the spacecraft is capable of varying its link rate during a contact such as the adaptive coding and modulation feature of the 2nd Generation Digital Video Broadcasting - Satellite (DVB-S2). Given a fixed symbol rate, spacecraft radios using this standard vary their data rate by adjusting modulation and coding on a frame-by-frame basis [17]. Link margin requirements at the lowest versus highest rates allow an 18.4 dB adjustment to handle weather, fading,

and path loss differences. However, operation at a middle rate within this range is more typical giving less adjustment downwards to handle weather. We also see in Fig. 2 that rain attenuation can exceed even 18.4 dB. In this case, the monitor system could use NIMBUS predictions to adjust symbol rate in advance of the pass to maintain connectivity.

A reasonable question to consider is whether or not the monitor system should cancel existing contacts if they are predicted to be heavily impaired. In a per-minute cost model, this would potentially save money. However, a service provider may seek to discourage last-minute cancellations since it would have difficulty finding a replacement customer. They may implement a mechanism to prevent near-term cancellations or discount the refund provided (if any) back to the customer if cancellations are made with too little notice. Instead we assume the monitor system will seek to keep existing contacts and use them to transfer the maximum amount of data possible, scheduling additional passes when necessary to make up the difference.

Future refinement of NIMBUS could produce impairment predictions more granular than the current binary classifier. Predictions of exact level of link degradation would help fine-tune the monitoring system's reaction. In this work, we do not explore changing transmit powers, but optimizing over both power and data rate may have benefits in certain power constrained platforms. Finally, other sources of forecasting information in addition to NIMBUS could provide a future monitor system with more comprehensive information for decision making.

V. Conclusion

We presented the results of our NIMBUS algorithm, which predicts rain impairment in satellite communication links, and rain attenuation estimates. We used power measurement data collected by a digital receiver, that was observing the Ka-band CW beacon transmitted from the geostationary ANIK F2 satellite, to backtest our algorithm. When operating on a 30-min lead time, our binary classifier achieved an overall accuracy of 84% and the attenuation estimation RMSE was 1.67 dB. In our discussion section, we described a concept of operations for satellite communications that would use NIMBUS to respond dynamically to rain events, as opposed to adding a flat margin to a static link budget.

References

[1] K. McCarthy, F. Stocklin, B. Geldzahler, D. Friedman, and P. Celeste, "Nasa's evolution to ka-band space communications for near-earth spacecraft," in *SpaceOps 2010 Conference Delivering on the Dream Hosted by NASA Marshall Space Flight Center and Organized by AIAA*, 2010, p. 2176.

[2] D. Chakraborty, F. Davarian, and W. Stutzman, "The ka-band propagation measurements campaign at jpl," *IEEE Antennas and Propagation Magazine*, vol. 35, no. 1, pp. 7–13, 1993. DOI: 10.1109/74.210825.

[3] C. B. Haskins and W. P. Millard, "Multi-band software defined radio for spaceborne communications, navigation, radio science, and sensors," in *2010 IEEE Aerospace Conference*, IEEE, 2010, pp. 1–9.

[4] M. Piasecki, J. Downey, N. Pham, *et al.*, "Development and demonstration of a wideband rf user terminal for roaming between ka-band relay satellite networks," 2021.

[5] D. J. Mortensen, C. Roberts, and R. Reinhart, "Automated spacecraft communications service demonstration using nasa's scan testbed," in *2018 SpaceOps Conference*, 2018, p. 2583.

[6] D. J. Israel, C. J. Roberts, R. M. Morgenstern, J. L. Gao, and W. S. Tai, "Space mobile network concepts for missions beyond low earth orbit," *Space Operations: Inspiring Humankind's Future*, pp. 25–41, 2019.

[7] A. Gannon, S. Paulus, C. Gemelas, and L. Vincent, "Spacecraft-initiated scheduling of commercial communications services," in *39th International Communications Satellite Systems Conference (ICSSC)*, Stresa, Italy, Oct. 2022.

[8] J. Esper, "Draft lunanet interoperability specification," National Aeronautics and Space Administration, Tech. Rep., 2022.

[9] "Noaa multi-radar/multi-sensor system (mrms)." (2022), [Online]. Available: https://registry.opendata.aws/noaa-mrms-pds (visited on 08/11/2022).

[10] S. Agrawal, L. Barrington, C. Bromberg, J. Burge, C. Gazen, and J. Hickey, "Machine learning for precipitation nowcasting from radar images," *arXiv preprint arXiv:1912.12132*, 2019.

[11] S. Pulkkinen, D. Nerini, A. A. Pérez Hortal, *et al.*, "Pysteps: An open-source python library for probabilistic precipitation nowcasting (v1. 0)," *Geoscientific Model Development*, vol. 12, no. 10, pp. 4185–4219, 2019.

[12] S. Ravuri, K. Lenc, M. Willson, *et al.*, "Skilful precipitation nowcasting using deep generative models of radar," *Nature*, vol. 597, no. 7878, pp. 672–677, 2021.

[13] "Mrms: Operational mrms grib2 tables." (2022), [Online]. Available: https://www.nssl.noaa.gov/projects/mrms/operational/tables.php (visited on 03/21/2023).

[14] R. J. Acosta, M. Zemba, J. Morse, and J. Nessel, "Two years of simultaneous k (sub a)-band measurements: Goldstone, ca; white sands, nm; and guam, usa," in *18th Ka and Broadband Communications Conference*, 2012.

[15] "Wgrib2: Wgrib for grib-2." (2023), [Online]. Available: https://www.cpc.ncep.noaa.gov/products/wesley/wgrib2/ (visited on 03/14/2023).

[16] "Pywgrib2_s: Home page." (2023), [Online]. Available: https://www.cpc.ncep.noaa.gov/products/wesley/wgrib2/pywgrib2_s.html (visited on 03/14/2023).

[17] "Digital Video Broadcasting (DVB); Second generation framing structure, channel coding and modulation systems for Broadcasting, Interactive Services, News Gathering and other broadband satellite applications (DVB S2)," Standard ETSI EN 302 307 v1.2.1, 2009.