Advanced Thermal Model Correlation Using Reduced-Order Models

Derek W. Hengeveld[1] and Jacob Moulton[2]
Redwire, Albuquerque, NM, 87106

Thermal model correlation uses test results to better estimate uncertain inputs. During correlation, selected inputs are modified (based on engineering expertise and intuition) traditionally in an iterative process. This becomes computationally expensive since the thermal model needs to be run for each iteration. Several advanced thermal model correlation methods have been developed to minimize this expense. These commonly involve gradient based searches that guide input parameter selection to more optimal points. The approach outlined in this paper leverages reduced-order models to reduce the computational expense associated with iterative methods. Reduced-order models provide computationally efficient surrogates of high-fidelity thermal models and are often built for a singular reason: reducing development cycle times and costs. By leveraging their speed, reduced-order models provide a new method for thermal model correlation. The typical computational expense inherent to traditional thermal model correlation is significantly reduced and thousands to millions of iterations can be achieved in seconds to minutes. The result is a collection of input factor combinations that meet correlation requirements. Reduced-order model thermal model correlation was applied to a nominal 6U CubeSat. Results were obtained for 50,000, 500,000, and 1,000,000 runs and distributions were plotted. The results showed a new statistical method of exploring correlation parameters.

Nomenclature

o	=	output response index
t	=	test index
ACS	=	Attitude Control System
BOL	=	Beginning of Life
CBERS	=	China-Brazilian Earth Resources Satellite
EOL	=	End of Life
GFOP	=	Group Factor of Performance
MMS	=	Magnetospheric MultiScale
PCB	=	Printed Circuit Board
RMS	=	Root Mean Square
ROM	=	Reduced-Order Model
RSS	=	Root Sum Square
TMC	=	Thermal Model Correlation
TVac	=	Thermal Vacuum

I. Introduction

UNCORRELATED thermal models are often based on an assortment of uncertain input parameters. Examples include interface conductance based on handbook values and estimated insulation ε* values. Because of this, it is hard to know which combination of parameter values accurately predict reality. To overcome this, thermal model correlation uses test results to better estimate and validate these uncertain inputs. This process commonly involves institutional knowledge and an iterative method that quickly becomes time-consuming and costly.

[1] Senior Engineer, 2309 Renard Place SE, Suite 101, Albuquerque, NM 87106.
[2] Staff Engineer, 2309 Renard Place SE, Suite 101, Albuquerque, NM 87106.

To overcome these challenges, reduced-order models (ROMs) were used. ROMs provide computationally efficient surrogates of high-fidelity thermal models (e.g., Thermal Desktop® models). An approach for creating these surrogates using efficient sampling and data fitting was developed and successfully applied to a broad range of spacecraft applications. This approach provides numerous benefits including computational speed. Leveraging this speed, ROMs can be used to calibrate thermal model parameters to experimental test data using an automated, repeatable, and simple methodology. This investigation will explore the use of ROMs for thermal model correlation. Details of the methodology behind this process will be provided. An example will illustrate how this method can be used in practice and highlight the computational efficiency and capabilities of this approach.

II. Thermal Model Correlation

Thermal models are developed to predict hardware performance (e.g., temperatures) across a broad range of operating conditions (e.g., hot- and cold-orbits). These models help ensure in-situ performance meets requirements. However, predictions can suffer from input parameter uncertainties; therefore, correlation is necessary. Thermal model correlation is an activity that helps improve a thermal model to better predict in-situ performance. In the current state of the art, a series of thermal tests (e.g., TVac) are completed to gather experimental data at prescribed temperature levels; this data is then compared to thermal model predictions using a suitable criterion (e.g., temperature difference) [1]. More specifically, the correlation process consists of finding a thermal model input parameter set that generates the best fit to the test data within a level of goodness deemed appropriate to satisfy mission requirements [2].

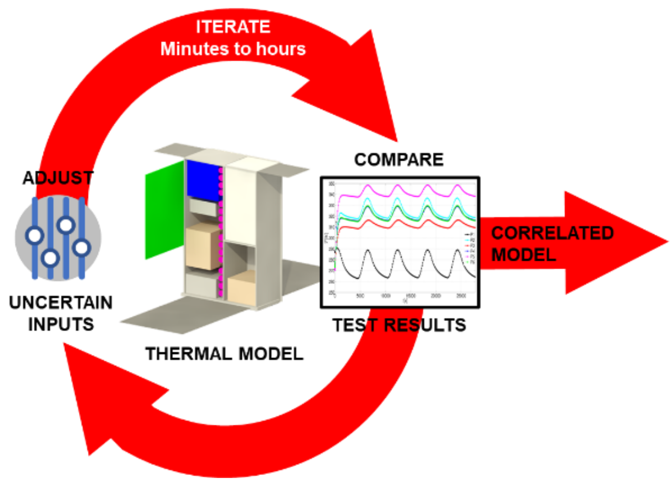

Figure 1. Traditional Thermal Model Correlation. *Traditional thermal model correlation is an iterative process based on engineering expertise and intuition.*

The literature provides several, real-world examples of thermal model correlation including the China-Brazilian Earth Resources Satellite (CBERS) [3] and Magnetospheric MultiScale (MMS) [4]. These works highlight the complexity of model correlation. First, thermal model correlation must be applicable to a broad range of input parameters. Examples include [2-5]:

- contact/interface conductance/resistance
- material (e.g., composite) conductivity
- equipment heat dissipations
- boundary conditions
- surface optical properties
- as-built insulation performance
- convection coefficients
- pressure losses
- capacitance at sensor locations
- harness conductance

Next, correlation must be able to handle many output responses over multiple tests. In these selected works, output temperature responses varied from 53 up to 468 thermocouple locations. For the MMS mission, results were obtained over 6 thermal balance tests [4]. Common amongst thermal model correlation is the criterion used for evaluation. Thermal model correlation typically reduces the difference between model and test temperatures. For CBERS, analysts reduced this difference to within 5°C. For MMS, six parameters were iteratively adjusted until the root mean squared (RMS) difference between model output and test data was minimized (RMS of less than 2°C). In fact, optimized correlation reduced the RMS to 1.4°C (from 7°C pre-correlation) [4].

Traditionally, selected parameters for correlation are modified (based on engineering expertise and intuition) in an iterative process, re-calculating the thermal model at each iteration. This becomes computationally expensive since the thermal model needs to be run for each iteration. For large and complex models, iterative correlation of a thermal model can take months using available processes and methods [2]. To overcome these challenges, several advanced thermal model correlation methods have been developed. Frey et al. [2] focused their correlation efforts on linearizing steady-state systems and applying an optimization algorithm to adjust linear and radiative conductors until the solution to the temperature vector compares favorably to test temperatures. Cataldo et al. developed a model-based correlation method for complex, large-scale systems [1]. They use the differences between test measurements and model predictions to derive incremental corrections to model parameters through a Taylor expansion truncated to the first order. The Jacobian matrix is calculated by perturbing model parameters one at a time and re-running the model. This method was applied to the James Webb Space Telescope [1, 6]. Anglada et. al. combined genetic algorithms with in-house thermal analysis software [7]. They applied their correlation tools to both steady-state and transient analyses. Finally, C&R Technologies provides features within Thermal Desktop® for calibrating both steady-state and transient models [8]. Calibration setup includes defining uncertainties and acceptable ranges for input factors of interest, adding measures if needed, and defining one or more data logger compare routines in their logic manager. Next, a case is setup with calls to SOLVER and Dynamic SINDA with the goal to minimize object set from the data loggers. In the next section, details of a ROM thermal model correlation approach is described.

III. Reduced-Order Modeling Thermal Model Correlation

ROMs provide computationally efficient surrogates of high-fidelity thermal models (e.g., Thermal Desktop® models). ROMs are often built for a singular reason: reducing development cycle times and costs; however, methods for building ROMs vary considerably. Methods include those that reduce the dimension of the underlying high-fidelity model. Examples include: projection-based methods [9, 10] and nodal reduction through lumped parameter methods [11]. Alternatively, ROMs can be built by interrogating the high-fidelity model, generating training data, and then creating a metamodel or statistical emulator by interpolating the observed data [12, 13]. An approach for creating a statistical emulator using efficient sampling and Gaussian process data fitting was developed and successfully applied to a broad range of spacecraft applications [14-18]. The developed approach provides numerous benefits, the primary being computational speed. Leveraging their speed, ROMs can be a useful tool for thermal model correlation.

The thermal model correlation method in this paper essentially takes the traditional thermal model correlation approach and leverages the speed of ROMs to iterate thousand to millions of times in seconds (Figure 2). The ROM Thermal Model Correlation (ROM-TMC) approach can be applied to most any thermal model assuming you can appropriately convert it into a ROM and requires just a few simple steps. This approach can also be used to correlate against a collection of test results (e.g., hot-soak, cold-soak, etc.).

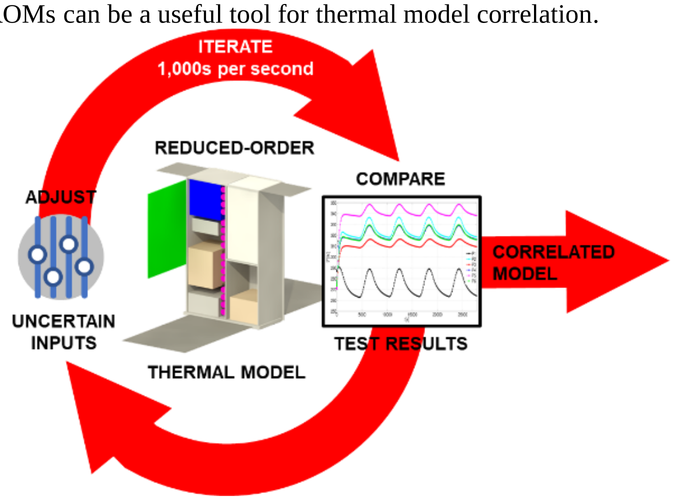

Figure 2. ROM-TMC. *Leveraging the speed of a ROM, thousands to millions of iterations can be completed in seconds.*

First, ROM inputs are identified as either fixed or correlation factors. Fixed input factors are those that remain static during a given test (e.g., hot soak heater power) and are therefore given a test-specific value during correlation. Correlation input factors are those that are uncertain, and we want to correlate. These input factors are given a range for us to correlate within across all test results. Next, ROM output responses are selected that will be correlated to. For each ROM output response and test case, both measured values and margins are allowed. Margins can be used to help filter out solutions.

Once inputs and outputs are setup, the iterative correlation process can begin. For each iteration, a random set of input factors is selected and quickly processed using the ROM. ROM outputs are provided and compared against measured values for all test cases. A performance metric is calculated by taking the root sum square (RSS) of the difference between measured and ROM values. A Group Factor of Performance (GFOP) metric for the trial is calculated by summing the RSS values (root sum square between ROM and Measured values) across all output responses and tests (Equation (1)). Currently, the GFOP equally weights the contributions from all outputs and tests. For example, cold- and hot-tests contribute the same to GFOP although hot-tests typically provide richer test information.

$$GFOP = \sum_{t}^{Test} \sum_{o}^{Output} \sqrt{\left(ROM_{t,o} - Measured_{t,o}\right)^2} \tag{1}$$

Leveraging the speed of ROMs, this process can be repeated very quickly. In fact, thousands to millions of iterations can be completed in just a few seconds using typical laptop/desktop processing power. This method then provides not just one but a collection of viable solutions, each with a unique GFOP. The solutions with the lowest GFOP could be considered the 'best' although the collection of solutions provide opportunity for more advanced correlation analysis and better understanding.

The advantage of the ROM-TMC approach is its ability to systematically identify viable solutions that meet correlation requirements. Further, ROM-TMC provides not just a single, but a collection of viable solutions. Additionally, this method takes very little user intervention; both ROM creation and model correlation are relatively 'hands-off'. The most significant disadvantage of ROM-TMC is the computational expense of ROM creation. As model complexity grows, so does this expense and in some cases could make this method too costly. In many cases though, ROM-TMC can provide time savings; however, the biggest value comes from having a collection of viable solutions which provides a lot more insight, confidence, and reduced risk.

IV. Case Study

A nominal Thermal Desktop® model was used to evaluate the use of ROMs for thermal model correlation. This 1795 node generic 6U CubeSat consists of fourteen submodels including: antenna, attitude control system, avionics, battery, 1U payload, 2U payload, propulsion, deployable radiator (100 cm x 239.4 cm), radios, solar array, tabbed structure, top plate of the structure, structural rails, and a test submodel. The Thermal Desktop® model included a TVac chamber to capture two test conditions, hot and cold (with different temperatures of the mounting flange, thermal shroud, and chamber door being fixed per test condition). This Thermal Desktop® model was converted into a ROM using the methods outlined in [14, 15, 17].

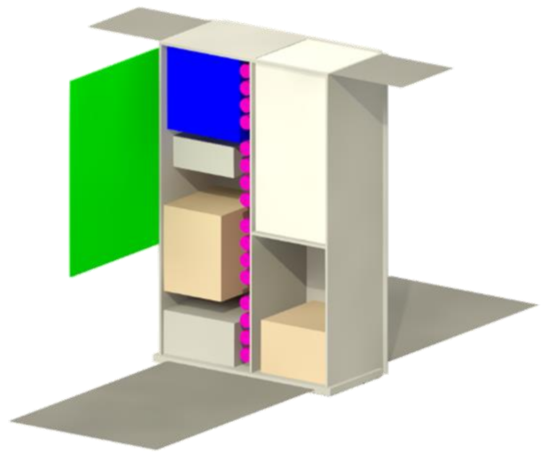

Figure 3. Thermal Desktop® Model. *A ROM was developed based on a 1795 node generic 6U CubeSat model.*

Testing included 13 thermocouple locations (3 on the TVac chamber and test fixture, and 10 placed on various locations of the CubeSat that were used for model correlation purposes). Thermal tests included a 3-hour hot and cold soak, where the average soak temperatures for both test conditions were used for model correlation. Eight input parameters were selected to be varied during this model correlation work, including: conductance values between components and the structure, PCB board thermal conductivity, and component optical properties. Model parameters adjusted during this model correlation effort, and therefore include in the ROM, were:

- Conductance Values: Conductance values connecting different components and subsystems within a CubeSat are usually estimated based on handbook values. As a result, five symbols were created that control the conductance value between five different CubeSat components and the CubeSat structure, and can all be varied as input factors during model correlation. It is anticipated that the conductance value affiliated with both the Attitude Control System (ACS) and the Payload_2U component will be larger than the other three conductance values, as approximately twice as many bolts were used to connect these two components to the CubeSat structure.
- Thermophysical Properties: Thermophysical properties assigned in a thermal model are typically estimated depending on handbook values, vendor material certificates, and/or well-accepted modeling standards (e.g., property files). However, the material conductivity of PCB boards is not as certain and can vary more widely depending on the amount of copper included. As a result, a symbol that controls the conductivity value of the PCB board components in the CubeSat was included and is one input parameter that can be varied during model correlation. A large value for PCB conductivity is anticipated, as PCB boards with high levels of copper were purposefully used in assembling the 6U CubeSat.
- Optical Properties: Optical properties assigned in a thermal model are typically estimated depending on a vendor's material or surface treatment certificate. However, these certified values are typically assigned a range that includes beginning of life (BOL) and end of life (EOL) expected values. As a result, two sensitivity factor symbols were created that get added (or subtracted if the sensitivity factor value is negative) to the optical properties nominal values during the TVac test, which allows for slight variations in the estimated absorptivity and emissivity of each component to be considered.

A complete summary of thermal model correlation input parameters with their ranges is shown in Table 1. Output responses include nodes from the following submodels: AVIONICS, BATTERY, PAYLOAD_1U, PAYLOAD_2U, PROPULSION, RADIATOR_DEPLOYABLE, SOLAR_ARRAY, STRUCTURE_TAB, and STRUCTURE_WALLS. Steady-state test results for both cold- and hot-soaks is summarized in Table 2.

Table 1. Summary of thermal model correlation input parameters

Symbol Value	*Description*	*Nominal Value*	*Low value*	*High value*
Cond_ACS_to_Struc	Conductance between components and the structure [W/K]	2.1	0.05	10.0
Cond_Avionics_to_Struc		1.1	0.05	10.0
Cond_Payload1U_to_Struc		1.1	0.05	10.0
Cond_Payload2U_to_Struc		2.1	0.05	10.0
Cond_RadiatorHinge		2.5	0.05	10.0
PCB_k	PCB thermal conductivity [W/m-K]	25	0.50	60.0
SensFac_Abs	Optical property offset [---]	0	-0.07	0.12
SensFac_Emiss		0	-0.01	0.12

Table 2. Summary of thermal model correlation test parameters

Test [---]	*Avionics* [°C]	*Battery* [°C]	*Payload 1U* [°C]	*Payload 2U* [°C]	*Propulsion* [°C]	*Radiator End* [°C]	*Radiator Base* [°C]	*Solar Array* [°C]	*Bus Top* [°C]	*Bus Center* [°C]
Cold	-13.8	-12.7	-12.4	-13.1	-14.3	-16.2	-17.1	-23.5	-23.8	-14.4
Hot	46.2	47.6	46.7	46.3	45.5	40.4	40.5	28.8	28.5	46.4

It is important to note that although only two test conditions, eight input factors, and ten output responses (i.e., temperature measurement locations) were included in this example; this approach for model correlation can support dozens of test conditions, model parameters, and temperature measurement locations in an efficient and effective manner.

The ROM-TMC method was applied for 50,000, 500,000, and 1,000,000 runs. The computational expense to evaluate the runs was ~30 seconds per 50,000 runs for a Dell XPS 15 9570 and Intel i7-8750H CPU @ 2.20GHz. Therefore, it only took 30 seconds, 5 minutes, and 10 minutes to identify 36, 377, and 758 viable solutions for 50,000, 500,000, and 1,000,000 runs, respectively. On average, it took 1.2 seconds to identify a single viable solution. The first three results (sorted by Group Factor of Performance) are shown in Table 3 and Table 4 for 50,000 and 1,000,000 runs. Note these correlation test parameters apply to both hot- and cold-tests.

Table 3. Summary of thermal model correlation test parameters (50,000 runs)

Cond ACS to Struc [W/°C]	*Cond Avionics to Struc [W/°C]*	*Cond Payload1U to Struc [W/°C]*	*Cond Payload2U to Struc [W/°C]*	*Cond Radiator Hinge [W/°C]*	*PCB k [W/m-°C]*	*SensFac Abs [---]*	*SensFac Emiss [---]*	*Group Factor of Performance [---]*
5.6	9.5	5.2	4.5	0.4	56.6	-0.1	0.1	44.0
9.1	9.1	6.4	8.5	0.7	47.5	0.0	0.1	45.5
7.6	9.8	6.6	4.8	0.6	36.6	0.0	0.1	50.0

Table 4. Summary of thermal model correlation test parameters (1,000,000 runs)

Cond ACS to Struc [W/°C]	*Cond Avionics to Struc [W/°C]*	*Cond Payload1U to Struc [W/°C]*	*Cond Payload2U to Struc [W/°C]*	*Cond Radiator Hinge [W/°C]*	*PCB k [W/m-°C]*	*SensFac Abs [---]*	*SensFac Emiss [---]*	*Group Factor of Performance [---]*
7.2	10.0	6.9	4.1	0.4	51.4	0.0	0.1	39.6
8.7	9.2	7.0	6.1	0.5	39.6	0.0	0.1	41.9
9.0	9.5	7.3	5.9	0.5	52.5	0.0	0.1	42.7

Each viable solution is summarized for all tests (i.e., cold- and hot-cases) and demonstrates that several solutions can be quickly obtained using the ROM-TMC approach. In addition, each solution is a unique combination of input factors and not just a slight perturbation of other solutions. For example, the difference between 'PCB k' in the solutions varies considerably. Also shown in these tables is that increasing the run number from 50,000 to 1,000,000 improved the GFOP by ~10%. This trend in GFOP was further evaluated by examining its distribution for increasing number of runs (Figure 4). Each histogram in this figure is normalized such that the total shaded area is equal to 1 for 20 bins. For 50,000 runs, there is not a clear GFOP distribution; it appears that the ROM-TMC approach is almost equally likely to identify solutions over a broad range of performance. In addition, solutions were found from ~45 to 80 GFOP. However, as the run numbers increase it becomes clear that this method is most likely to identify solutions with a GFOP near 60, but this approach is also able to identify a solution nearly 33% better than average (i.e., GFOP = 40).

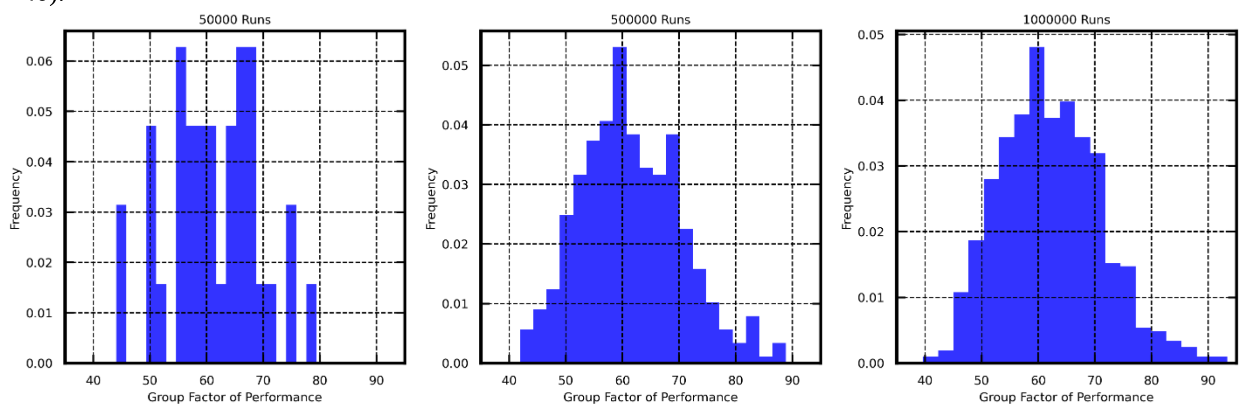

50,000 runs | 500,000 runs | 1,000,000 runs

Figure 4. Group Factor of Performance frequency distributions for various runs. *Group Factor of Performance values that met correlation requirements were plotted on a histogram for 50,000, 500,000, and 1,000,000 runs.*

These plots demonstrate that there are many solutions that meet our correlation requirements; however, there are solutions that work much better than others. The ROM-TMC is well-suited to finding solutions to this global optimization problem and can evaluate 1,000s to millions of input factor combinations due to ROM speed. This method is not limited to initialization conditions, which can limit the utility of gradient-based optimization methods and might only find local optimal solutions.

Exploring these trends further, we can look at how these distributions look for specific input factors. The frequency of 'Cond ACS to Struc' input factor values over all solutions was plotted for increasing run numbers (Figure 5). This figure shows a similar trend to Figure 4; for low run numbers, there is not a clear distribution of solutions. However, as the run number increases, it becomes clear that the input factor setting is more likely to be a larger number than small.

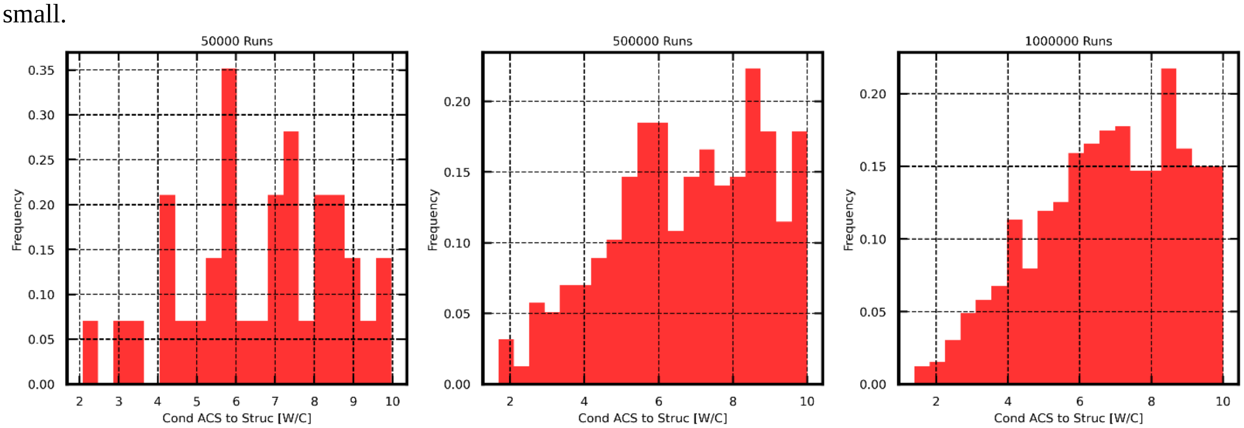

50,000 runs | 500,000 runs | 1,000,000 runs

Figure 5. Input factor frequency distributions for various runs. *Input factor values that met correlation requirements were plotted on a histogram for 50,000, 500,000, and 1,000,000 runs.*

Using 1,000,000 runs, these frequency plots were applied to all input factors to show well-developed distributions (Figure 6). These input factor distributions show several interesting trends. For example, the ‘SensFac Abs’ input factor indicates that all values are equally likely candidates for a solution (i.e., the distribution appears uniform). Other distributions, such as ‘PCB k’ indicate that solutions are more likely to be in a smaller range than that evaluated. This same trend is even more apparent for ‘Cond Avionics to Struc’ where solutions trend between 8.0 and 10.0 W/C. This type of evaluation is a unique feature of the ROM-TMC approach.

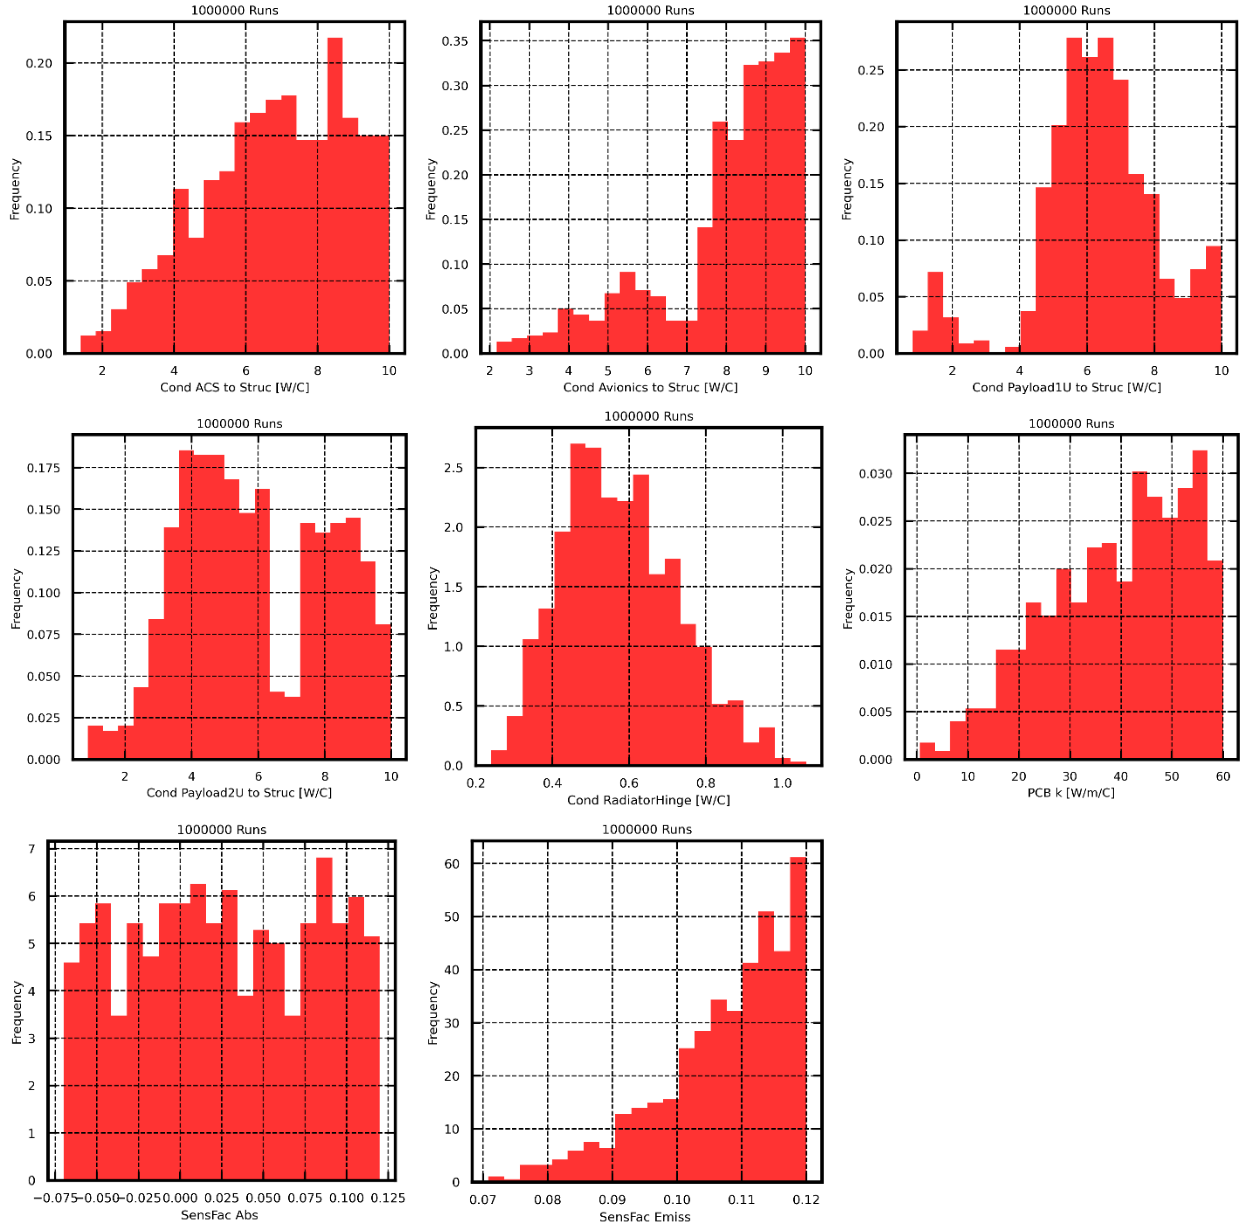

Figure 6. Input factor frequency distributions for 1,000,000 runs. *Input factor values that met correlation requirements were plotted on a histogram for 1,000,000 runs.*

Output response residual (i.e., ROM minus test result) frequency plots (all from -3 to 3°C with 20 bins) were created using 1,000,000 runs (Figure 7). As with the input factors, these show well-developed distributions and interesting trends. For example, the ‘RADIATOR_DEPOLOYABLE.70’ temperature difference output response shows that both negative and positive residuals are possible and a residual of 0°C is also possible. However, other outputs clearly show that 0°C residuals are not possible (i.e., SOLAR_ARRAY.2 and STRUCTURE_TAB.35). We could leverage this information to eliminate these outputs from our GFOP to help better focus correlation on those outputs we have more control over. As before, this type of evaluation is a unique feature of the ROM-TMC approach.

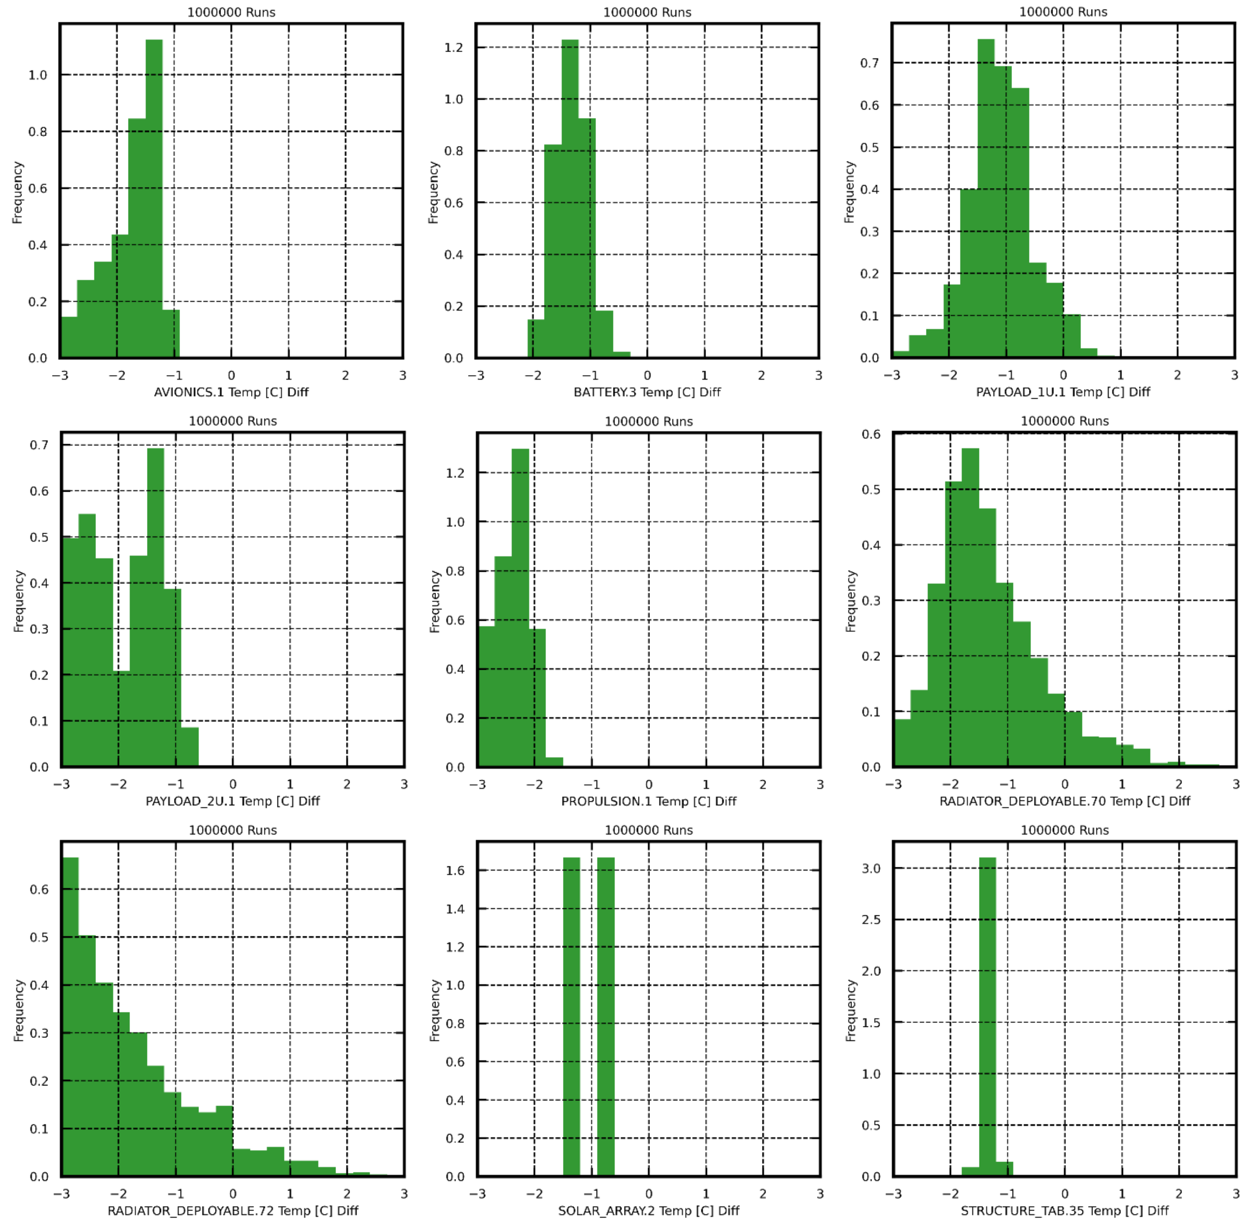

Figure 7. Output response frequency distributions for 1,000,000 runs. *Output response residual values that met correlation requirements were plotted on a histogram for 1,000,000 runs.*

V. Conclusions and Future Development

Thermal model correlation uses test results to better estimate uncertain inputs. Selected inputs for correlation are modified, based on engineering expertise and intuition, in an iterative process. This becomes a computationally expensive process since the full thermal math model needs to be run for each iteration. Several advanced thermal model correlation methods have been developed to help avoid the manual, iterative processes inherent to traditional methods. These typically involve gradient based searches that guide input parameter selection to more optimal points. The approach outlined in this paper leverages reduced-order models (ROMs) to reduce the computational expense associated with iterative methods. ROMs provide computationally efficient surrogates of high-fidelity thermal models and are often built for a singular reason: reducing development cycle times and costs. By leveraging their speed, ROMs provide a new method for thermal model correlation. The typical computational expense inherent to traditional thermal model correlation is significantly reduced and thousand to millions of iterations can be achieved in seconds to minutes. The result is a collection of input factor combinations that meet correlation requirements. ROM-TMC was applied to a nominal 6U CubeSat. Results were obtained for 50,000, 500,000, and 1,000,000 runs and distributions were plotted. The results showed a new statistical method of exploring correlation parameters.

Acknowledgments

This material is based upon work supported by Small Business Innovative Research projects with NASA and the Air Force Research Laboratory.

References

1. Cataldo, G., et al., *Model-based thermal system design optimization for the James Webb Space Telescope.* Journal of Astronomical Telescopes, Instruments, and Systems, 2017. **3**(4): p. 044002.
2. Frey, B., et al. *Development of an automated thermal model correlation method and tool*. 2015. 45th International Conference on Environmental Systems.
3. Leite, R., et al. *CBERS 3&4 TM Thermal Balance Test Results and Satellite Thermal Mathematical Model Correlation*. in *40th International Conference on Environmental Systems*. 2010.
4. Zara, R. and R.J. D'Antonio. *MMS Optical Bench Assembly (OBA) Thermal Design Validation and Thermal Model Correlation*. in *43rd International Conference on Environmental Systems*. 2013.
5. Myers, J.L., et al. *Hall Thruster Thermal Modeling and Test Data Correlation*. in *52nd AIAA/SAE/ASEE Joint Propulsion Conference*. 2016.
6. Cataldo, G. *Integrated modeling under uncertainty for the James Webb Space Telescope*. in *Modeling, Systems Engineering, and Project Management for Astronomy VIII*. 2018. International Society for Optics and Photonics.
7. Anglada, E. and I. Garmendia, *Correlation of thermal mathematical models for thermal control of space vehicles by means of genetic algorithms.* Acta Astronautica, 2015. **108**: p. 1-17.
8. Technologies, C.R. *Calibrating Models to Test Data Automatically*. 2021; Available from: https://www.crtech.com/videos/calibrating-models-test-data-automatically.
9. Qian, J., et al., *Projection-based reduced-order modeling for spacecraft thermal analysis.* 2015. **52**(3): p. 978-989.
10. Wang, Y., et al. *A projection-based model order reduction simulation tool for spacecraft thermal analysis*. in *Thermal and Fluids Analysis Workshop*. 2011.
11. Fernández-Rico, G., et al., *Quasi-autonomous thermal model reduction for steady-state problems in space systems.* Applied Thermal Engineering, 2016. **105**: p. 456-466.
12. Ranjan, P., R. Haynes, and R. Karsten, *A computationally stable approach to Gaussian process interpolation of deterministic computer simulation data.* Technometrics, 2011. **53**(4): p. 366-378.
13. Martin, J.D. and T.W. Simpson, *Use of kriging models to approximate deterministic computer models.* AIAA Journal, 2005. **43**(4): p. 853-863.
14. Hengeveld, D. and J. Moulton. *Automatic creation of reduced-order models using Thermal Desktop®*. in *48th International Conference on Environmental Systems*. 2018. 48th International Conference on Environmental Systems.
15. Hengeveld, D.W. *Reduced-Order Modeling for Rapid Thermal Analysis and Evaluation of Spacecraft*. in *46th AIAA Thermophysics Conference*. 2016.
16. Hengeveld, D.W. *Reduced-Order Modeling for Rapid Thermal Analysis and Evaluation of Spacecraft*. in *2016 Thermal & Fluids Analysis Workshop (TFAWS)*. 2016. Moffet Field, CA.
17. Hengeveld, D.W. and A. Biskner. *Enhanced data exploration through Reduced-Order Models*. in *47th International Conference on Environmental Systems*. 2017. Charleston, South Carolina.
18. Hengeveld, D.W. and J.A. Moulton. *Automatic creation of reduced-order models using Thermal Desktop®*. in *2018 Thermal & Fluids Analysis Workshop (TFAWS)*. 2018. Galveston, TX.