IOP Conf. Series: Materials Science and Engineering **1344** (2026) 012016 doi:10.1088/1757-899X/1344/1/012016

British (Ball) Aerospace Cryogenic Thermal Margins

David Glaister[1] and Carlita Gorham[2]

[1,2] Cryogenics Solution Center, BAE Systems, Broomfield, CO, USA

[1]E-mail: david.glaister@baesystems.us
[2]E-mail: carlita.gorham@baesystems.us

Abstract. Thermal margins are critical to the development, the production, the test, and ultimately the performance of thermal subsystems. The subset of thermal subsystems that are cryogenic are even more reliant on thermal margins for success. Thermal uncertainty margins cover the inherent uncertainty, due to both design and analysis limitations, of predicting performance in space. They are not equivalent to sensitivity analyses that focus on the impact on performance of individual design and analysis parameters and assumptions. Thermal margins cover all the areas of uncertainty in the design. These include areas that are nearly impossible to eliminate uncertainty such as: thermal interstitial joints (workmanship influenced), orbit environments (transient), surface properties (vary with individual surfaces), thermophysical properties (vary with individual material lots), MLI (empirical ranges), etc. These areas often have empirically derived ranges. At BAE Systems (formerly Ball Aerospace) cryogenic margin is based on the MIL-STD philosophy, which this paper will cover in more detail. It is critical to not only establish, but track margins, throughout a program, from pre-proposal to on-orbit. Examples will be provided that show the tracking of margin on several actual BAE programs.

Acronyms and Nomenclature

ATP	=	Authority to Proceed
BAE	=	British Aerospace
CDR	=	Critical Design Review
FPA	=	Focal Plane Assembly
GSFC	=	Goddard Space Flight Center
K	=	Kelvin
M	=	Thermal uncertainty or heat load margin (%)
MLI	=	Multi-Layer Insulation
NASA	=	National Aeronautics Space Administration
NTE	=	Not To Exceed
PDR	=	Preliminary Design Review
$Q_{capacity}$	=	Cooling source capacity
Q_{margin}	=	Predicted heat load with margin
Q_{pred}	=	Predicted load without margin
Q_{rqmt}	=	Required cooling capacity without margin
PFPA	=	Power dissipation of FPA (Focal Plane Assembly)
TVAC	=	Thermal Vacuum test

1. Margin Philosophy and Foundation

Thermal margin needs to be appropriate to the uncertainty risk. Typical stereotypes have the Thermal Engineer defaulting to "the more margin, the better" and the Program Manager response "we don't need any thermal margin as Thermal Engineers are always too conservative". Instead, there needs to be a balance to collaboratively establish margin commensurate with risk. Too little margin can result in real risk of late program or on orbit failure to meet temperature, performance, or requirements. Too much margin can, at a minimum, increase system mass and power. At its worst, too much margin can prevent a system from even being proposed as feasible.

The foundational documents that provided guidance on thermal margin uncertainty in cryogenic systems were MIL-HDBK-340 [1] and MIL-STD-1540D [2]. MIL-HDBK-340 was cancelled in 2017. MIL-STD-1540D was officially cancelled in 1999. Nonetheless, members of the practitioner community continued to use the document's guiding principles up through the release of its successor document, SMC-S-016 [3]. The original construct used for margin in these documents was similar to that implemented for room temperature applications whereby high and low margin limits are determined relative to a desired operational temperature. However, at cryogenic temperatures and with cryogenic cooling sources, it is typically more effective to use heat load as opposed to temperature margin. For example, 50% heat load margin at either 45 K or 100 K has similar equivalent uncertainty risk, but 8 K of temperature margin at 45 K has a significantly higher uncertainty risk than 8 K margin at 100 K. Also, cryogenic cooling sources have discreet interfaces for the cooling, which enable application of the heat load margin at those interfaces. Gradated temperature margins as defined in SMC-S-016 (shown in Table 1) have a gradual decrease in high/low margin temperature during the approach to 40 K and below. In addition, for passively controlled components, a margin of ±11°C is to be applied to worst-case analytical predictions for passively controlled components. The suggested control authority margin for actively controlled components is a minimum of 25%.

Table 1. Passive thermal uncertainty margins for hardware at low temperature.

Predicted Temperature [K]	Uncertainty Margin [K]
>203	11
203 ⇔ 186	10
185 ⇔168	9
167 ⇔150	8
149 ⇔132	7
131 ⇔114	6
113 ⇔ 96	5
95 ⇔ 78	4
77 ⇔ 60	3
59 ⇔ 42	2
<42	1

IOP Conf. Series: Materials Science and Engineering **1344** (2026) 012016 doi:10.1088/1757-899X/1344/1/012016

2. Cryogenic Margin Definition and Application

The following establishes the definition and specific application of cryogenic margin, both of which have had been debated extensively. The key for both is based on the original Martin Donabedian cryogenic MIL-STD margins [4] which utilized a database of cryogenic programs, which this paper's author compiled as an early career engineer at The Aerospace Corporation. That database is consistent with the following definition and application.

The definition is as follows in Equation 1 (there has been some confusion on this due to a typographical error in the first edition of the Spacecraft Thermal Control Handbook where the denominator mistakenly has Qmargin instead of Qpred):

M: Thermal uncertainty or heat load margin (%)
Q_{margin}: Predicted heat load with margin
Q_{pred}: Predicted load without margin

$$M = 100 * \frac{Q_{margin} - Q_{pred}}{Q_{pred}} \quad (1)$$

Another method to express this is shown in Equation 2:

M: Thermal uncertainty or heat load margin (%)
$Q_{capacity}$: Cooling source capacity
Q_{rqmt}: Required cooling capacity without margin

$$M = 100 * \frac{Q_{capacity} - Q_{rqmt}}{Q_{rqmt}} \quad (2)$$

The margins and heat loads are all applied at the interface to the cooling source, as opposed to the location of each heat load. For a cryocooler, this is at the interface to the cold tip of the cooler. For a cryoradiator, at the interface to the radiator; for example, between a conduction bar and the radiator or between a transport heat pipe flange and the radiator. This is consistent with the historical data which was the basis for the MIL-STD.

Q_{pred} is the realistic worst-case prediction of the heat load as opposed to unrealistic. Realistic includes the use of assumptions (i.e.-properties) that can occur and/or are within the range of empirically derived values. It does not mean unrealistic beyond worst-case assumptions that, for example, could be the result of sensitivity analyses where properties are varied by a percentage.

3. Recommended Cryogenic Margin

As margins cover uncertainty, the recommended levels are a function of the design and hardware maturity, which in turn is dependent on the program maturity and milestones. Thermal uncertainty should decrease as the program advances. This can be due to requirements becoming more mature and stable, the thermal design becoming more defined, and/or the completion of component, assembly, or system tests and model correlations to verify the performance.

Table 2 provides a listing of the margins used for contemporary cryogenic programs throughout the Project Life Cycle at BAE Systems and NASA Goddard Space Flight Center. Per the MIL-STD, at BAE the margin applied at ATP or during Pre-Phase A activities is 50%. Following the Authority to Proceed (ATP) this value is held through Phase A. Beginning with PDR or Phase B, the standard margin applied is gradually reduced at each of the subsequent Life Cycle Phases until it reaches a value of 25% at launch (i.e., Phase E).

It is important to note that the margin cannot be reduced until the milestone is completed (i.e.-after PDR completion, not leading into the PDR). Using the original MIL-STD cryogenic margin database, these margins correspond to a bounded (less than) 98% probability or 2 sigma standard deviation range. In other words, if applied, at these levels, 98% of on orbit performance should be within the worst-case prediction plus this margin.

NASA GSFC has adopted a more conservative approach in Pre-Phase A whereby 100% margin is applied. This value is held through the completion of Phase A. Upon entering Phase B, the margin requirement becomes 80%. Subsequent reductions are applied at each of the remaining phases of the Life Cycle, ending with a margin of 33% at the time of Launch. For passively controlled systems, heat load margin is applied to the projected "flight" minimum operating temperature.

Table 2. Example cryogenic margins spanning the project life cycle.

Life Cycle Phase: NASA/MIL	MIL-STD[a] and BAE Systems	NASA GSFC[b]
Pre-Phase A/ATP	50	100
Phase A/SDR	50	100
Phase B/PDR	45	80
Phase C/CDR	35	50
Phase D/Qualification	30	40
Phase E/TVAC Correlation	25	33

[a]Established from MIL-STD-1540 [2]
[b]Taken from NASA/GSFC GEVs Cryogenic Section [5]

4. Tailoring of Cryogenic Margin

As margin covers uncertainty, it can and should be updated and modified when the uncertainty changes. The generalized update for program milestone was provided in the previous section. However, there is benefit to modifications or tailoring at a lower level. For example, some individual heat loads can have larger or smaller uncertainties. For cryogenic systems, detector electrical heat loads can be measured with good accuracy. Thus, a common practice at BAE is to reduce the applied margin at ATP for detector dissipation to 20% against prior, similar measured detector electrical dissipations. If an NTE (Not to Exceed) dissipation or quantity (e.g., thermal strap conductance) is specified in a deliverable or requirement flow down, then the margin for that component's performance can be reduced, and sometimes eliminated.

Ultimately, the best method for reducing uncertainty and cryogenic thermal margin is through test and model correlation. Following qualification/final TVAC testing, the generic approach margin can be reduced to its Life Cycle minimum (i.e., generally 25% - 33%) prior to flight. However, it is not uncommon to incorporate other tests at lower levels that can also reduce the uncertainty and margin. For example, a component such as a thermal strap or conductor bar can be tested at the component level. A thermal brassboard can be built and tested earlier in an instrument design (e.g., prior to PDR). It is important to revisit the margin requirements based on results from these lower-level tests, not only because the uncertainty has been reduced, but also to justify with project management and/or customers the value of performing early risk reduction tests. From a program or customer perspective, if the margin is not reduced following the completion of an early risk reduction test, then there may be limited benefit to performing the test at all.

On the reverse side, for resource/schedule/cost constrained programs, the margin should often be tailored to a higher level due to increased risk from less resources for testing and analyses. This may seem counterintuitive (that a cost constrained program needs more margin), but the key is to equate to thermal uncertainty and program success probability. If a cost constrained program still wants a 98% probability of success, but has less manpower or test resources, then the uncertainty increases and so does the margin need. However, if a program can have a lower probability of success, then the margin could be decreased. At BAE, it is typical to double the cryogenic margin requirements for a low resource program (for example: 100% at ATP). Additionally, the experience and expertise of the program team can also be taken into account. For example, a less experienced team should carry more margin.

5. BAE Mission Cryogenic Margin Examples

This section includes actual tracking data of cryogenic margins (defined in Equation 1 or 2) on six BAE (formerly Ball) programs with seven instruments. Each was at an instrument level and was within the last 8 years (2017 to 2025). Some key takeaways are:

1. These programs demonstrate the adequacy of the MIL-STD based BAE margin philosophy of essentially 50% at ATP, as all had post-delivery or on orbit performance meeting requirements with >25% margin
2. For the two programs with flight data, the delivery (post TVAC test and model correlation) to on-orbit margin either stayed the same or increased. In other words, the on-orbit conditions did not cause an additional loss in margin relative to ground testing.
3. For some of these programs, the margin significantly exceeds the requirements. This can be a sign of overdesign. However, for these examples with large margins, the additional margin wasn't due to overdesign, but instead related to leveraging heritage design features (to reduce cost) and mature technologies (reduce risk).
4. These examples illustrate that the margin becomes more stable as the design and requirements mature.

Margin data is focused on heat loads, however, for background only all these programs use a similar cooler and all of these systems are for a single active cooler per instrument.

5.1 Program A

For Program A, it was challenging to provide tracking information across the entire life cycle of the program. However, this program did have a TVAC Correlated Model/delivery cryogenic margin of 37% and, on orbit, has never seen less than 53% margin.

5.2 Program B

For Program B, the tracking history is shown in Figure 1. Program B had two independent cryogenic instruments, and the plot includes the margin for both. Instruments 1 and 2 operate at two different temperatures, resulting in the margin offset seen in the figure. Oscillations in margin around CDR are due to design changes that occur earlier in programs. This program had a significant assembly and integration anomaly that affected both instruments at the time of the final TVAC, with a corresponding loss in margin. However, there were mitigations enacted that compensated for some of the margin loss, resulting in meeting requirements at delivery and now, on-orbit.

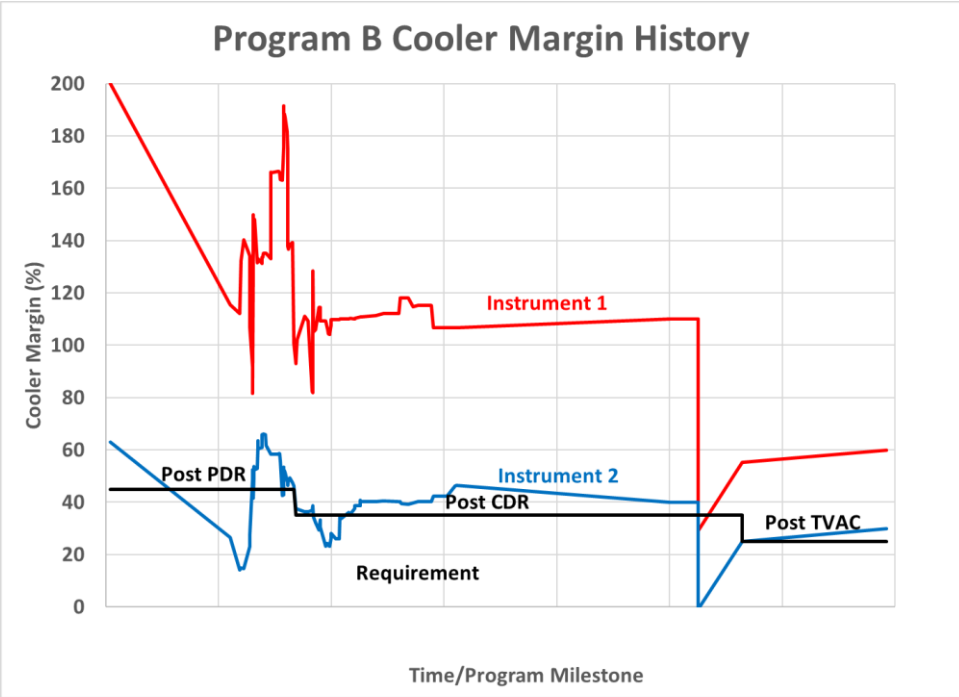

Figure 1. Cryogenic margin history for BAE Program B with two instruments.

5.3 Program C

For Program C, the tracking history is shown in Figure 2. This figure shows margin calculated and tracked both with and without the FPA dissipation included in the total heat loads. The FPA dissipation was a NTE value, thus, was contractually not required to be included (subtracted from both the total heat load and the cooler capacity) in the margin requirement. Early in Program C there was below requirement margin due to technical challenges inherited from a previous program. However, mitigations were developed after PDR that enabled meeting the margin requirements at CDR. Program C has now completed final TVAC, model correlation, and delivery; and met all margin requirements.

IOP Conf. Series: Materials Science and Engineering **1344** (2026) 012016
doi:10.1088/1757-899X/1344/1/012016

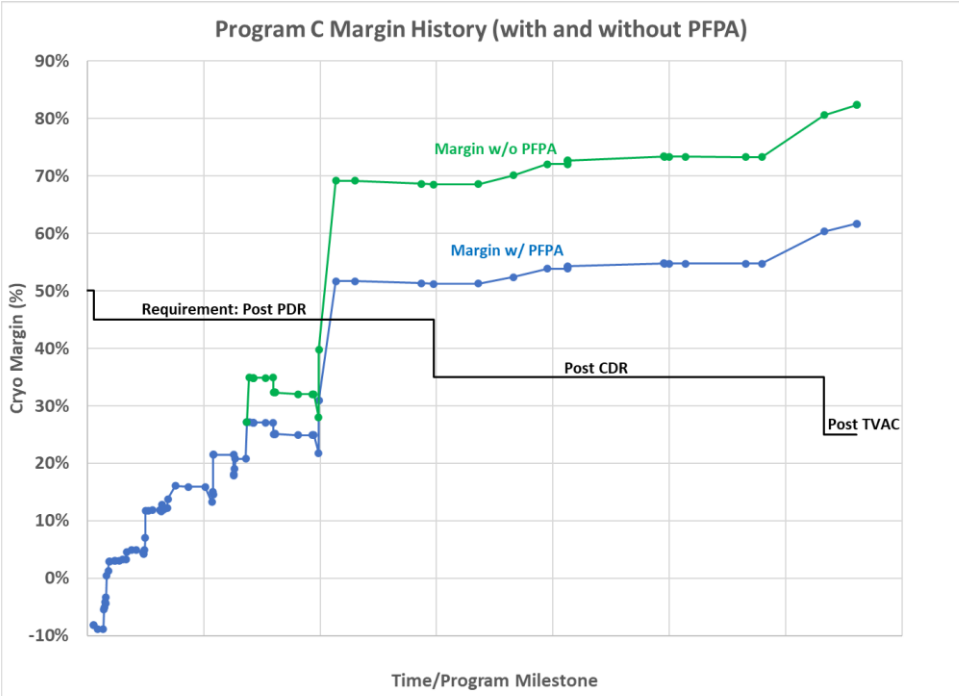

Figure 2. Cryogenic margin history for BAE Program C with and without the NTE FPA dissipation.

5.4 Program D

For Program D, the tracking history is shown in Figure 3. This figure shows margin calculated and tracked both with and without the FPA dissipation included in the total heat loads. The FPA dissipation was a NTE value, thus, was contractually not required to be included (subtracted from both the total heat load and the cooler capacity) in the margin requirement. Early in Program C there was a significant margin impact due to technical challenges inherited from a previous program. However, mitigations were developed after PDR that enabled margin recovery by CDR. Program D has now completed final TVAC, model correlation, and delivery; and met all margin requirements.

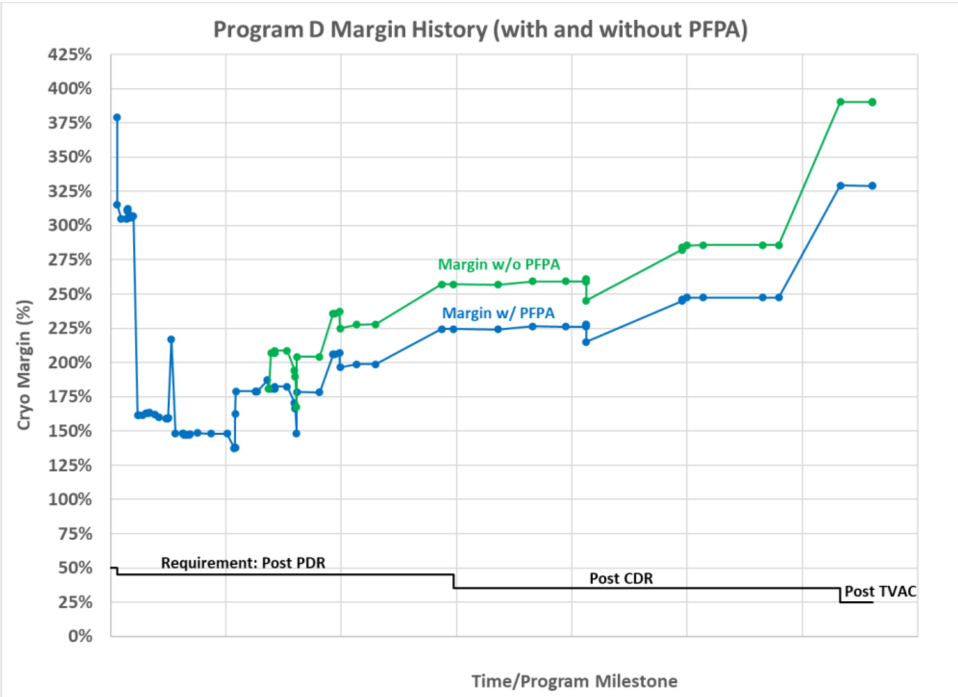

Figure 3. Cryogenic margin history for BAE Program D with and without the NTE FPA dissipation.

5.5 Program E

For Program E, the tracking history is shown in Figure 4. This figure shows margin calculated and tracked without including the FPA dissipation. The FPA dissipation was a NTE value, thus, was contractually not required to be included (subtracted from both the total heat load and the cooler capacity) in the margin requirement. Program E was leveraged off a previous program design and, thus, started with lower thermal uncertainty. To compensate, the margin requirements were contractually tailored for lower levels initially with the traditional 25% requirement at delivery/TVAC model correlation. Program E has now completed final TVAC, model correlation, and delivery; and met all margin requirements.

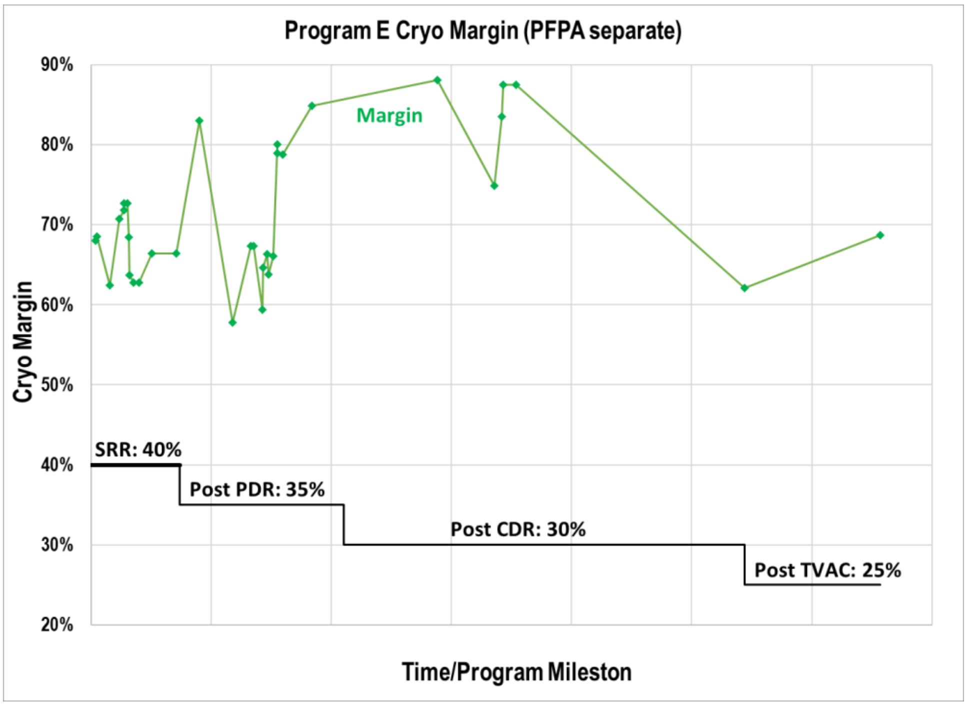

Figure 4. Cryogenic margin history for BAE Program E without the NTE FPA dissipation.

5.6 Program F

For Program F, the tracking history is shown in Figure 5. This figure shows margin calculated and tracked without including the FPA dissipation. The FPA dissipation was an NTE value, thus, was contractually not required to be included (subtracted from both the total heat load and the cooler capacity) in the margin requirement. Program F was the second unit of Program E and, thus, started with lower thermal uncertainty. Similarly to Program E, the margin requirements were contractually tailored for lower levels initially with the traditional 25% requirement at delivery/TVAC model correlation. Also, Program F had the identical thermal model and, thus, margin prior to its TVAC test. Program F has now completed final TVAC, model correlation, and delivery; and met all margin requirements.

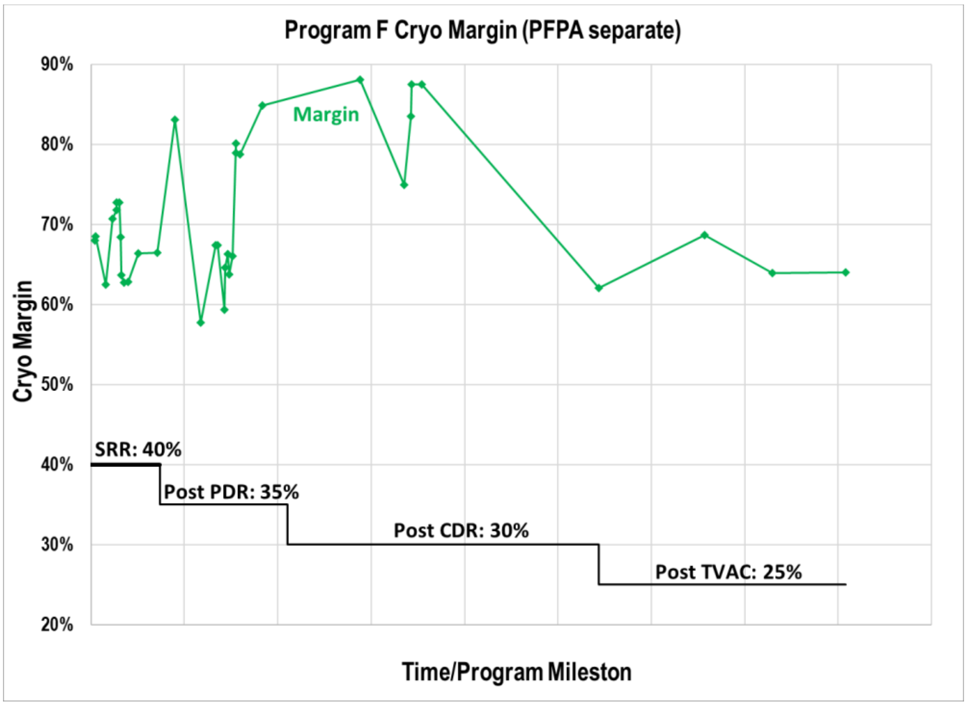

Figure 5. Cryogenic margin history for BAE Program F without the NTE FPA dissipation.

6. Conclusion

Thermal/cryogenic margin covers uncertainty and is not the same as a sensitivity analysis. Cryogenic margin appropriate to risk is optimum and critical to program success. The cryogenic margin foundation is key to understanding how to apply margin. As such, cryogenic margin should be applied at the cooling source interface and against realistic worst-case predictions. Margin definition divides the net cooling capacity margin (capacity minus required load) by the required heat load. Standard cryogenic margin recommendations are 50% at ATP and 25% at Delivery. Margin tailoring is important and includes tailoring for individual heat loads, for lower-level tests, and for lower resource/higher risk programs.

Six BAE programs and seven instruments with cryogenic margin tracking were presented. Those programs demonstrated the adequacy of the BAE margin philosophy as every program met at least 25% margin at delivery and/or on-orbit. Additionally, for every on-orbit program, cryogenic margin from delivery (post TVAC test and model correlation) to on-orbit margin either stayed the same or increased, illustrating the accuracy of BAE ground TVAC testing to simulate on-orbit conditions.

References

[1] Military Handbook, 1985, Application Guidelines for MIL-STD-1540B; Test Requirements for Space Vehicles, MIL-HDBK-340.

[2] Military and Government Specs and Standards (NPFC), 1999, Product Verification Requirements for Launch, Upper Stage, and Space Vehicles, MIL-STD-1540.

[3] AFSC Space and Missile Systems Center Standard: Test Requirements For Launch, Upper-Stage, and Space Vehicles, Aerospace Corporation, Report # SMC-S-016.

IOP Conf. Series: Materials Science and Engineering **1344** (2026) 012016 doi:10.1088/1757-899X/1344/1/012016

[4] Donabedian, M., 2002, Thermal Control Margins, Risk Estimation, and Lessons Learned, Spacecraft Thermal Control Handbook Vol. II: Cryogenics, edited by Gilmore, D., The Aerospace Press and The American Institute of Aeronautics and Astronautics, Inc., Chap. 19, pp. 469-480.

[5] General Environmental Verification Standard (GEVS) for GSFC Flight Programs and Projects", 2013, GSFC Technical Standards Program, NASA Goddard Space Flight Center, Greenbelt, MD., GSFC-STD-7000A.