Deriving Mission-Specific Radiation Test Fluence Targets from Orbital Survivability Requirements

Joran Heldens
EDGX
Ghent, Belgium
joran.heldens@edgx.space

***Abstract*—Standard radiation hardness assurance methodologies prescribe fixed heavy-ion test fluences, typically 10^7 ions/cm^2, to screen for Destructive Single Event Effects (DSEE). This paper examines the statistical and physical foundations behind this number. We show that a radiation test is a sequence of Bernoulli trials whose aggregate behaviour is described by the Poisson distribution, establishing a direct relationship between test fluence, device cross section, and statistical confidence. By inverting this relationship through orbital event rate analysis, we derive the maximum permissible cross section for a given mission and translate it into a target fluence. For a 550 km sun-synchronous orbit over 7 years, the analysis recovers the standard 10^7 for the most conservative case, an unknown device without redundancy, providing a first-principles justification for the established value. The analysis further reveals that system-level redundancy is a powerful and underutilized lever: a single redundant device reduces the required test fluence by approximately one order of magnitude, independent of device characteristics. At stringent survival targets such as 99.997%, redundancy transforms an effectively untestable per-device requirement (MTTF $>$ 200,000 years) into a practical one (MTTF $\sim$ 1,200 years with one spare). These results do not replace existing standards but offer a complementary perspective that connects test specifications to the missions they serve, and provide a quantitative basis for incorporating redundancy into radiation test planning.**

***Index Terms*—single event effects, radiation hardness assurance, test fluence, IRPP, Poisson statistics, COTS, survivability**

I. Introduction

The radiation hardness assurance process for spacecraft electronics has historically been anchored to standardized test procedures. For Destructive Single Event Effects (DSEE), including Single Event Latchup (SEL), Single Event Gate Rupture (SEGR), and Single Event Burnout (SEB), standards such as ESCC 25100 [4] and JEDEC JESD57 prescribe test fluences of 10^7 ions/cm^2 at each LET step. But where does this number come from, and when is it appropriate?

The prescribed fluence of 10^7 ensures that a cross section as small as $\sim 3\times10^{-7}$ cm^2 would produce at least one observable event with 95% probability. This provides high confidence for very sensitive devices. However, the number carries an implicit assumption: that cross sections this small are mission-relevant. For many orbits, devices with much larger cross sections already satisfy mission reliability requirements, making 10^7 unnecessarily conservative.

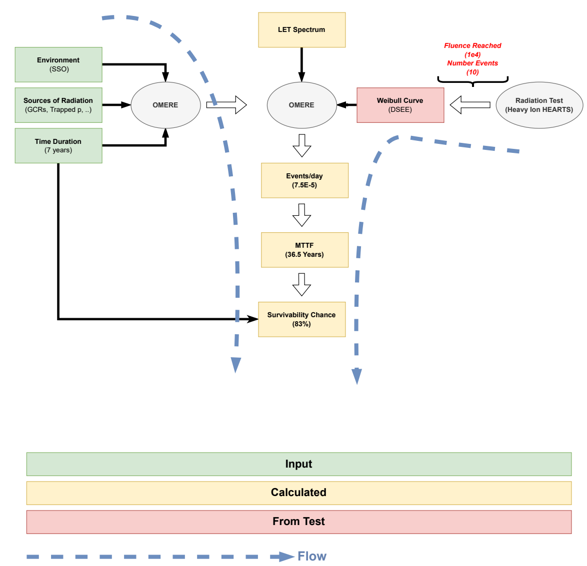

Fig. 1. Standard radiation processing flow.

For commercial-off-the-shelf (COTS) parts, increasingly prevalent in the NewSpace sector, the economic pressure is acute. Beam time at heavy-ion facilities is expensive and scarce. Testing a large candidate list to 10^7 each is prohibitive for small satellite programs. Yet screening to lower fluences without statistical justification provides false confidence.

This paper addresses the gap by developing a methodology that derives the required test fluence from first principles, connecting mission survivability requirements to test specifications through a traceable mathematical chain.

II. Standard vs. Inverse Approach

A. Standard Flow

The conventional DSEE qualification follows Figure 1. The mission environment is modeled to produce a LET spectrum. Device characterization from a heavy-ion campaign yields a Weibull cross-section curve, combined with the spectrum via IRPP rate calculation [1], [2] to compute an event rate, MTTF, and survivability. The test fluence and the survivability are decoupled: fluence is set by standard, and survivability computed after the fact.

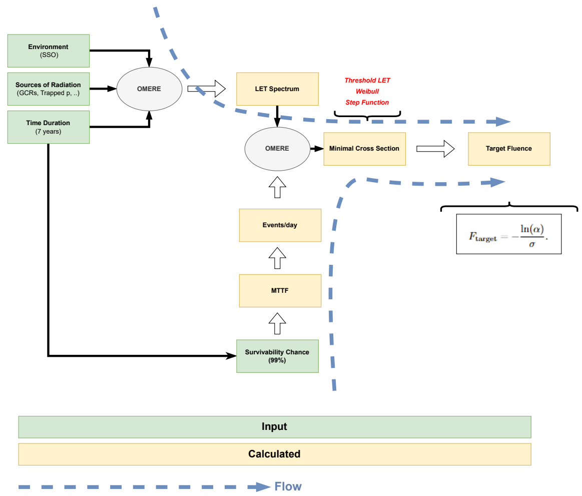

Fig. 2. Inverse approach: deriving target fluence from survivability.

B. Inverse Flow

We propose inverting the flow (Figure 2). The survivability requirement becomes an *input*; the target fluence is a *derived output*. The key intermediate is the maximum permissible cross section $\sigma_{\max}$.

III. Statistical Basis

A. A Radiation Test as Bernoulli Trials

A radiation test can be understood as a sequence of independent random experiments. Each incident ion either causes a destructive event or it does not. The unknown, true underlying probability of causing an event is p, the device's cross section. The number of successes (observed events) is k, and the number of trials is n, the applied fluence. The event count follows the binomial distribution:

$$P(k \mid n, p) = \binom{n}{k} p^k (1-p)^{n-k} \tag{1}$$

Since the cross section p is extremely small (10^{-7} to 10^{-3} cm^2) while the fluence n is very large (10^5 to 10^7), the binomial is mathematically equivalent to the Poisson distribution with expected value

$$\lambda = p \times n = \sigma \times F \tag{2}$$

giving

$$P(k \mid \lambda) = \frac{\lambda^k e^{-\lambda}}{k!} \tag{3}$$

B. Zero-Event Confidence Criterion

For DSEE screening ($k = 0$):

$$P(k = 0) = e^{-\sigma F} \tag{4}$$

1) Test context: We require that a device with true cross section $\geq \sigma_{\max}$ is detected with confidence $1 - \alpha_{\text{test}}$:

$$e^{-\sigma_{\max} F} = \alpha_{\text{test}} \implies \boxed{F_{\text{target}} = \frac{-\ln(\alpha_{\text{test}})}{\sigma_{\max}}} \tag{5}$$

TABLE I

CROSS-SECTION UPPER BOUNDS FOR $F = 1.10 \times 10^5$ IONS/CM^2, ZERO EVENTS.

CL	α	Multiplier	σ_{upper} [cm^2]
63%	$1/e$	$1/F$	9.09×10^{-6}
95%	0.05	$3.0/F$	2.73×10^{-5}
99%	0.01	$4.6/F$	4.18×10^{-5}

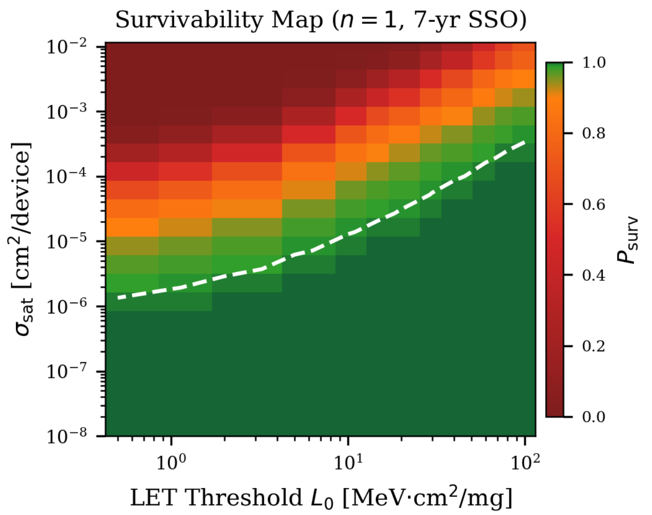

Fig. 3. Survivability map: Poisson zero-event probability in the L_0 vs. σ_{sat} plane for a 7-year SSO ($n = 1$). Dashed contour marks 99% survival.

2) Mission context: A device with rate R survives with probability e^{-RT}. A survival requirement P_{surv} constrains:

$$R_{\max} \cdot T = -\ln(1 - P_{\text{surv}}) = \lambda_{\max} \tag{6}$$

Both contexts share the same Poisson framework.

C. The Rule of Three

At 95% confidence ($\alpha = 0.05$): $F_{\text{target}} \approx 3/\sigma_{\max}$. This is the classical "rule of three": $3N$ trials with zero failures gives 95% confidence that the rate is below $1/N$. Table I shows the sensitivity to confidence level.

IV. From Mission to Fluence

A. Connecting Orbit to Cross Section

The event rate $R(\sigma_{\text{sat}})$ is computed from the orbital LET environment and the Weibull cross section using the IRPP method [1], [2]. Since R is monotonically increasing in σ_{sat}, there exists a unique $\sigma_{\max}$ satisfying $R(\sigma_{\max}) \cdot T = \lambda_{\max}$. The target fluence follows from (5).

Fig. 3 shows the resulting survivability map for a 550 km SSO (7 yr, $n = 1$).

B. Effect of Redundancy

For n independently redundant devices:

$$P_{\text{sys}} = 1 - (1 - P_{\text{single}})^n \tag{7}$$

TABLE II
PER-DEVICE REQUIREMENTS FOR 99% SYSTEM SURVIVAL, 7-YR.

n	P_{dev}	$\lambda_{\max}$	$R_{\max}$ [/day]	Relaxation
1	99.00%	0.010	3.93×10^{-6}	–
2	90.00%	0.105	4.12×10^{-5}	10.5×
3	78.46%	0.243	9.49×10^{-5}	24.1×

TABLE III
PER-DEVICE REQUIREMENTS FOR 99.997% SYSTEM SURVIVAL, 7-YR.

n	P_{dev}	$R_{\max}$ [/day]	MTTF [yr]	Relaxation
1	99.997%	1.17×10^{-8}	233,000	–
2	99.45%	2.15×10^{-6}	1,275	184×
3	96.89%	1.24×10^{-5}	222	1,053×

Tables II and III quantify the per-device relaxation.

Fig. 4 shows how redundancy expands the safe region at the 99.997% level.

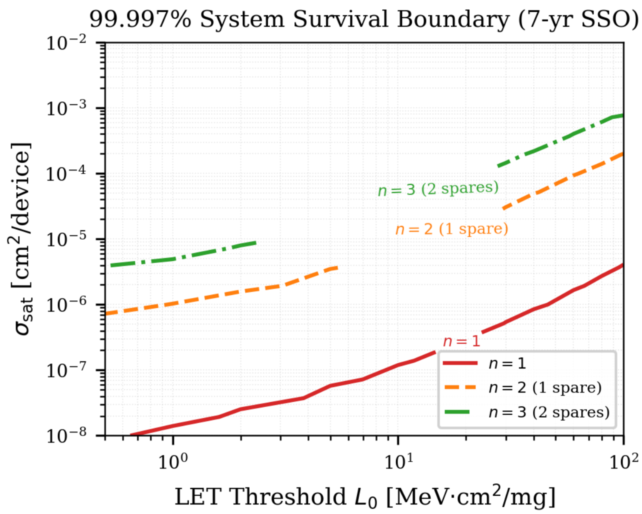

Fig. 4. 99.997% system survival boundary for $n = 1, 2, 3$. Redundancy dramatically expands the safe region in Weibull parameter space.

V. WORKED EXAMPLE: 550 KM SSO

We consider a 7-year mission at 550 km, 98° (ISO 15390 GCR solar min, AP-8 Min, ESP solar, 1 g/cm^2 Al). Our IRPP implementation was verified against OMERE 5.9 to within 9%.

Fig. 5 shows the required test fluence as a function of L_0 for three survival targets, each with $n = 1, 2, 3$. This is the central result: it directly answers the question *"for my orbit, my survival target, and my redundancy level, how much fluence do I need at each LET threshold?"*

Table IV gives numerical values for the 99% case without redundancy.

TABLE IV
DERIVED FLUENCE FOR 99% SURVIVAL, 95% CL, $n = 1$.

L_0	$\sigma_{\max}$ [cm^2]	F_{target} [ions/cm^2]	vs. std
1	3×10^{-7}	10^7	1×
10	5×10^{-6}	6×10^5	17×
30	5×10^{-5}	6×10^4	170×
100	5×10^{-4}	6×10^3	1700×

Fig. 6 compares how the safe region in Weibull space scales across survival targets and redundancy levels.

VI. THE CIRCULARITY PROBLEM

A. The Problem

Table IV shows that the required fluence depends strongly on L_0, which is precisely the quantity one intends to measure. For an unknown COTS device the most conservative assumption ($L_0 \to 0$) recovers the standard 10^7. However, the framework remains practically useful in two ways.

B. Redundancy Requires No Prior Knowledge

The redundancy relaxation is independent of device parameters. Adding one spare reduces the fluence by $\sim$10× regardless of L_0 or σ_{sat}. This is a system design decision, not a device characterization question. For a non-redundant unknown device the fluence is 10^7; with one spare it is $\sim 10^6$. No device knowledge needed.

C. Post-Test Fluence Justification

Even when the full fluence cannot be reduced before the first test, the framework enables rigorous *early stopping*. If an initial run establishes that no destructive events occur up to a certain LET, the device's L_0 is bounded from below. The methodology then provides a justified fluence target for subsequent runs, transforming it from a pre-test planning tool into a real-time decision tool during the campaign.

VII. DISCUSSION

The standard 10^7 corresponds to $\sigma = 3 \times 10^{-7}$ cm^2 at 95% confidence. Our analysis shows this is the correct target when the device is untested, the system has no redundancy, and high survival is required in LEO. Any relaxation reduces the justified fluence: system redundancy gives $\sim$10× per spare with no device knowledge needed; known L_0 gives 10 to 1,000×; these compound.

The most striking result is at high survival targets. At 99.997% without redundancy, the per-device MTTF must exceed 233,000 years, which is effectively untestable. With one spare this drops to 1,275 years; with two spares, 222 years. Redundancy transforms an impossible requirement into a routine one. Yet current practice typically derives test fluence from the per-device requirement without accounting for system redundancy.

Limitations include: (1) redundancy assumes independent failures (physically separated devices); (2) the rate model introduces $\sim$10% uncertainty; (3) GCR rates vary $\sim$2× over the solar cycle.

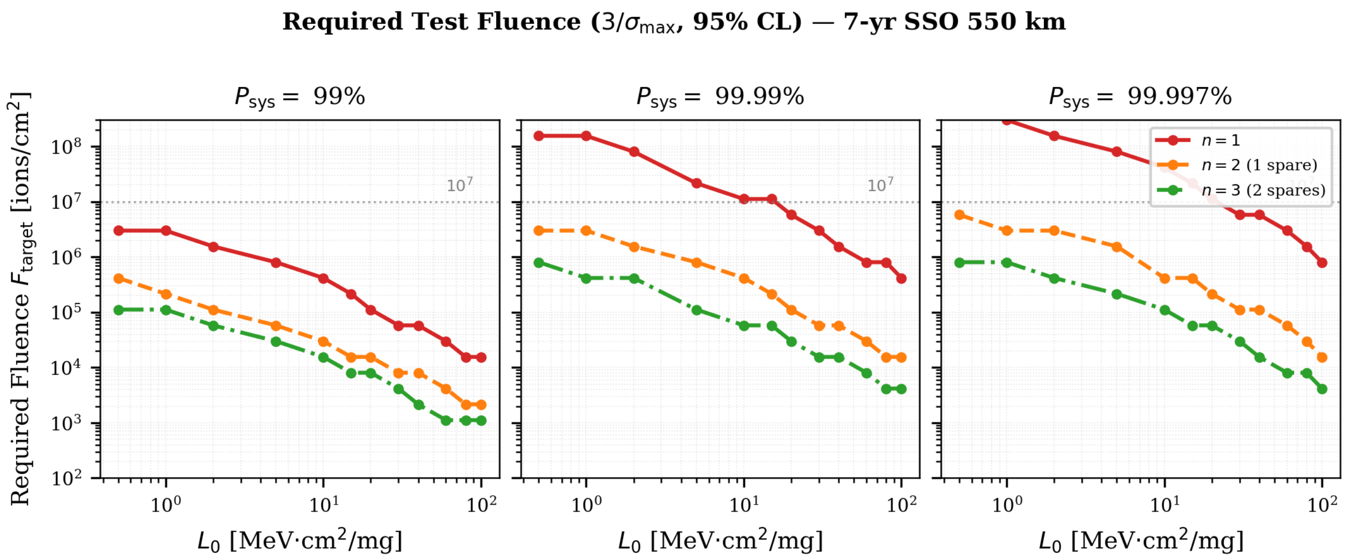

Fig. 5. Required test fluence ($3/\sigma_{\max}$, 95% CL) vs. LET threshold for three survival targets and three redundancy levels. The gray dashed line marks the standard 10^7. Redundancy shifts each curve downward by $\sim 10\times$ per added device.

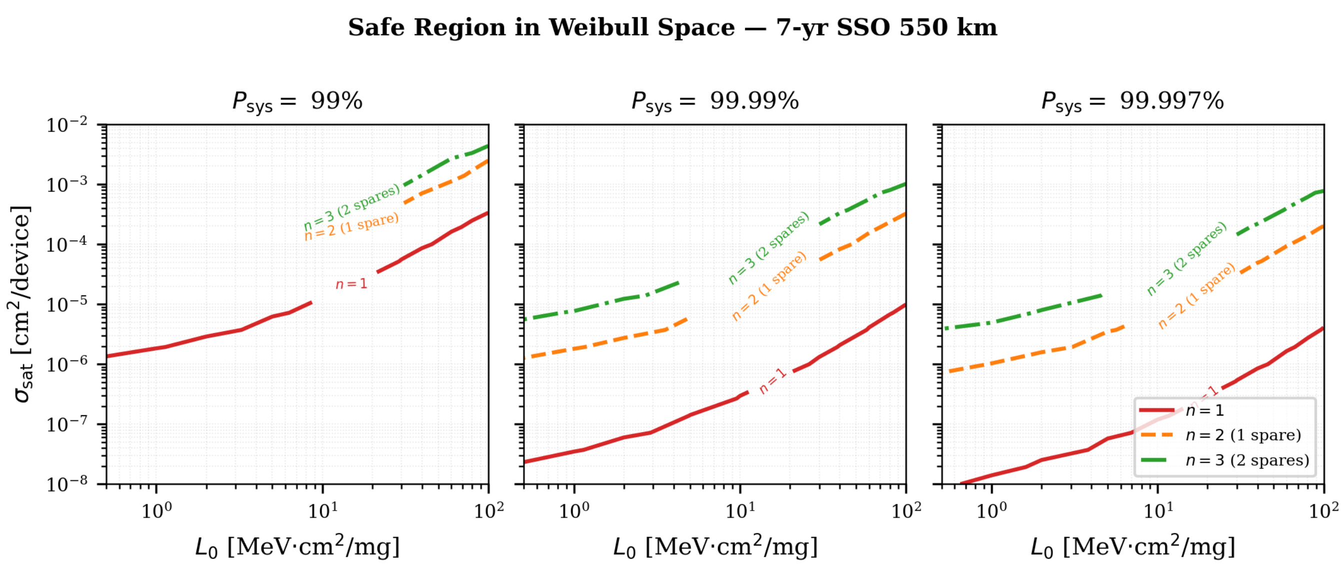

Fig. 6. Comparison of safe-region boundaries for 99%, 99.99%, and 99.997% system survival. Redundancy ($n = 2$, $n = 3$) shifts the boundary by orders of magnitude. Note that 99.997% with $n = 3$ is comparable to 99% with $n = 1$.

VIII. Conclusion

We have presented a methodology that connects radiation test fluence to mission survivability through the statistics of Bernoulli trials and their Poisson limit. The central equation, $F_{\text{target}} = -\ln(\alpha)/\sigma_{\max}$, reveals fluence as a derived quantity. For a 7-year 550 km SSO:

1) For an unknown device without redundancy, the framework recovers the standard 10^7, providing a first-principles justification for the established value.
2) System-level redundancy reduces the required fluence by $\sim 10\times$ per additional device, without any prior device knowledge.
3) Once initial characterization bounds L_0, justified early stopping can reduce beam time by one to three orders of magnitude.
4) At stringent targets (99.997%), redundancy transforms an untestable requirement (MTTF $>$ 200,000 yr) into a practical one (MTTF $\sim$ 1,200 yr with one spare).

This work does not claim that existing standards are wrong. They were designed to be universally applicable, and universality has clear value. Rather, we propose that the community consider an additional, complementary path: one where the test is designed *for* the mission, not independently of it. The methodology and tools presented here are intended as a starting point for that discussion, not a final word. We welcome scrutiny, counterexamples, and refinement from the radiation effects community. If the framework withstands that challenge, it may offer a practical route to reducing test costs, particularly in the growing small satellite sector, without sacrificing the statistical rigor that radiation hardness assurance demands.

References

[1] E. Petersen, *Single Event Effects in Aerospace*, Wiley, 2011.

[2] J. C. Pickel and J. T. Blandford, "Cosmic-ray-induced errors in MOS devices," *IEEE Trans. Nucl. Sci.*, vol. NS-27, no. 2, pp. 1006–1015, 1980.

[3] E. Petersen, "Single event analysis and prediction," *IEEE NSREC Short Course*, 1983.

[4] ESCC Basic Specification No. 25100, "Single Event Effects Test Method and Guidelines," ESA, 2014.