Statistical Techniques for Analyzing Process or "Similarity" Data in TID Hardness Assurance

R. Ladbury

Abstract— We investigate techniques for estimating the contributions to TID hardness variability for families of linear bipolar technologies, determining how part-to-part and lot-to-lot variability change for different part types in the process.

Index Terms—radiation effects, reliability estimation, quality assurance

INTRODUCTION

Radiation Hardness Assurance (RHA) methodologies against Total Ionizing Dose (TID) degradation impose rigorous statistical treatments for data from a part's Radiation Lot Acceptance Test (RLAT)[1] and/or its historical performance.[2],[3],[4],[5] However, no similar methods exist to treat "similarity" data—that is, data for similar parts fabricated in the same process as the part under qualification. This lack is notable because interpreting similarity data is more difficult and potentially riskier than interpreting RLAT or historical data for the same part type.

Rigorous techniques for using similarity data are essential for establishing the quantity or quality of data needed to judge a fabrication process radiation tolerant. Moreover, by casting a broad net over the parts in a process, statistical treatments of process data are more likely than a small-sample RLAT to identify whether parts in a process may exhibit extraordinary lot-to-lot or part-to-part variability or other failure distribution pathologies. Statistical analysis of process data can also be very valuable for part selection and design by bounding likely parametric degradation for a part before we have the opportunity to test it. This last factor is a significant boon for parts susceptible to enhanced low-dose rate sensitivity (ELDRS) where waiting for test data could take months.

Here we develop methods to disentangle part-to-part, lot-to-lot and part-type-to-part-type variation. The methods are useful for qualification decisions as well as quality control and detecting process changes and other "out-of-family" behavior.

We begin by discussing the data used in the study and the challenges of developing a statistic providing a meaningful measure of degradation across multiple part types, each with its own performance specifications. We then develop analysis techniques and apply them to the different data sets.

I. DATA

All data are from public sources. Data in Table I are taken from the Goddard Space Flight Center (GSFC) Radhome archival database,[6] (test dates 1997-2008) and are for op amps fabricated in the Analog Devices Inc. (ADI) bipolar process (minimum feature size >2.5 μm). Data in Table II are from reports for low-dose-rate (LDR) tests of Linear Technologies Corp. (LTC) RH-series parts, and are available on LTC's website.[7] Table III contains a subset of parts from Table I—those parts where we have data for multiple wafer lots, allowing us to explore lot-to-lot and part-type-to-part-type as well as part-to-part variation. For each lot in Tables I and III, we determined mean failure dose (where the first parameter exceeded specifications) and standard deviation about the mean. For the parts in Table II, we calculated the mean percent change in input bias current: %ΔIb and its standard deviation.

TABLE I: ADI BIPOLAR OP AMP FAILURE DOSES—KRAD(SI)

Part #	# lots	Mean Failure Dose	Avg. Lot Stand. Dev.
OP177	1	2.5	
OP400	4	5.32	1.7
OP497	1	5.7	1.2
OP467	2	9.2	4.5
OP07	3	9.4	1.9
OP97	1	11.1	1.5
OP270	1	12	2.8
OP271	1	12.5	2
OP470	1	12.5	
OP11	1	12.6	2.3
OP77	3	16.4	2.6
OP200	1	22	1.7
OP27	3	29.5	4.3

TABLE II: ΔIb FOR LTC RH SERIES PARTS (@50 KRAD(SI))

Part #	Function	Mean ΔIb	St. Dev.
RH1499	Op Amp	5%	0.02
RH119	V. Comp.	45%	0.04
RH1011	V. Comp.	56%	0.03
RH108	Op Amp	56%	0.10
RH1014	Op Amp	126%	0.06
RH1078	Op Amp	191%	0.04
RH27	Op Amp	215%	0.76

TABLE III: ADI OP AMP FAILURE DOSES (KRAD(SI)) FOR MULTI-LOT SAMPLES

Part #	Lot Mean Fail. Level	Lot St. Dev.
OP27(1)	39.00	2.80
OP27(2)	21.00	5.00
OP27(3)	28.50	5.07
Mean	29.50	4.29
St. Dev.	9.04	1.29
OP07(1)	9.50	3.20
OP07(2)	11.00	1.30
OP07(3)	7.60	1.10
Mean	9.37	1.87
St. Dev.	1.70	1.16
OP467(1)	9.90	1.30
OP467(2)	8.50	7.70
Mean	9.20	4.50
OP77(1)	21.19	4.08
OP77(2)	17.50	2.20
OP77(3)	10.50	1.60
Mean	16.40	2.63
St. Dev.	5.43	1.29
OP400(1)	8.00	1.76
OP400(2)	2.50	
OP400(3)	4.69	3.04
OP400(4)	6.08	0.28
Mean	5.32	1.69
St. Dev.	2.32	1.39

For all datasets, we chose either a normal distribution (if the quantity being fit was distributed from -∞ to +∞) or a lognormal fit (for quantities with values from 0 to +∞—e .g.

failure dose). However, since fit parameters for the lognormal distribution can be found by fitting the logarithms of the data to the desired distribution, we will represent the resulting means with the symbol μ and the standard deviations by σ.

Because similarity data must consider TID degradation in many different part types, a first challenge is developing a meaningful criterion for comparing degradation across such different part types. For the parts in Tables I and III, we defined failure dose as that where the first parameter (usually input bias current, Ib) goes out of specification for the device.

Such a failure criterion will not work for parts in Table II, since none of the parts failed parametrically or functionally at the highest dose of 50 krad(Si). Here, we compare the parts' parametric degradation—with input bias current change (ΔIb) as a proxy for degradation. Because pre-rad specifications of Ib varied widely from part to part, we normalized input bias current changes to pre-rad values (to arrive at %ΔIb). Although the LTC site has only one lot of data for each part type, previous data show series performance to be quite stable from lot to lot. For instance, 38 lots of RH1014 op amps showed that mean lot ΔIb varying by less than 2-3x from 60-200 krad(Si).[5]

Exploratory data analysis[8] revealed neither systematic correlation between mean failure level and standard deviation about the mean, nor any systematic time dependence of TID results that could indicate a process change. As such, we determined distributions for mean failure doses and standard deviations independently, assuming all parts in a given process are representative of a single distribution.

II. INFERENCE WITH LIMITED DATA

Our goal is to quantify types of variation that affect TID response for parts fabricated in a given process. These include part-to-part variation within the flight lot and variation in both mean hardness and part-to-part variability from lot to lot. We can estimate part-to-part variation within a lot using RLAT data.[1] We can also bound it to a desired confidence level if lot-to-lot variation is well behaved and we have enough representative historical data.[2]-[5] Likewise, unless the flight parts are exceptional, we can bound lot-to-lot variation with a large enough dataset of data for similar parts. However, the dataset must include data for enough part types to be a representative sample for the process. That is, we must be able to assess variability of mean hardness and part-to-part variation across lots for several (>3) part types in the process.

Because many of the parts in Table I include data only for a single lot, it is not possible to disentangle lot-to-lot variation from the part-type-to-part-type contribution. Rather, the rank plot in Fig. 1 (upper) shows the probability (abscissa) that the mean failure dose of a random lot of a random part type drawn from the process will exceed a given failure dose. Assuming Weibull statistics (which give the best fit), with 90% confidence 90% of lots in the process will not exhibit first failure below 3.3 krad(Si), and a "typical" lot of a typical part will be hard to ~13 krad(Si). The lower plot in figure 1 is a

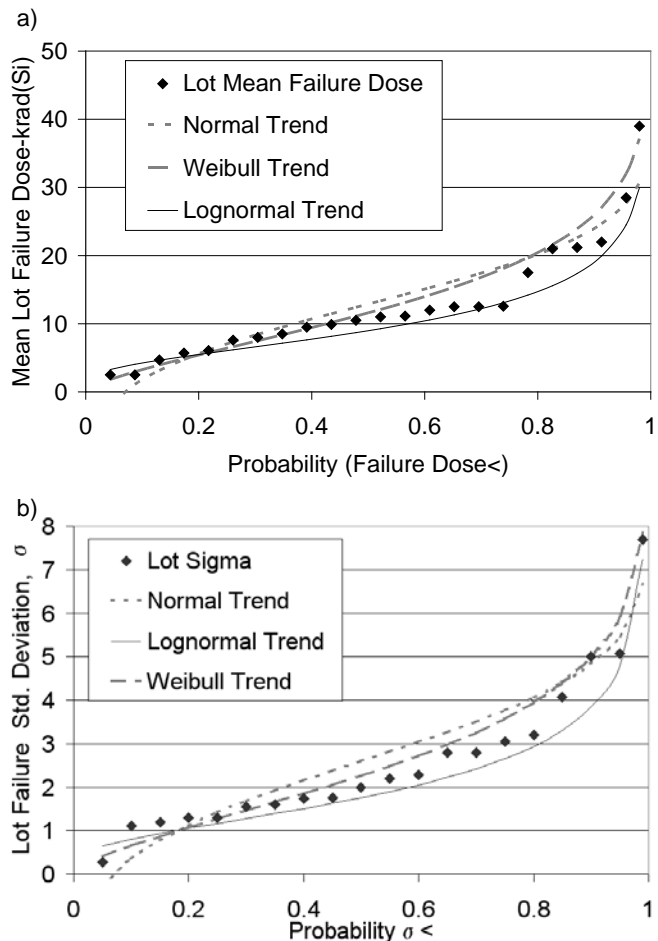similar rank plot for the standard deviation of failures in a lot.

Fig.1 Rank plots of (a) mean and (b)standard deviation of failures in a lot for OP series op amps in Table I. Plot a shows the probability that a random lot of a random part type remains within specification to a failure dose. Plot b shows the distribution of part-to-part standard deviations about those mean failure doses.

Likewise, for part types in Table II, we have data only for a single lot. However, while we have data for fewer part types, the RH series is a radiation hardened process, so we expect lot-to-lot and part-to-part variation to be moderate.

As such, the larger part-to-part variation exhibited by the RH27—sufficient to distort the lognormal fit to the other parts (Fig. 2a, b)—is surprising. The part also exhibited the highest overall %ΔIb, although this was less out of family (Fig. 2c). A query to LTC[9] revealed that the RH27 uses the same design as the commercial OP27 with no additional design hardening. As such, it likely represents a worst case for parts fabricated in the RH process. In contrast, the RH1499 exhibits very little proportional increase in Ib, as well as excellent part-to-part consistency because it was specifically designed to capitalize on the hardness of the RH process. This part probably approaches the limit of what can be achieved with process and design hardening.
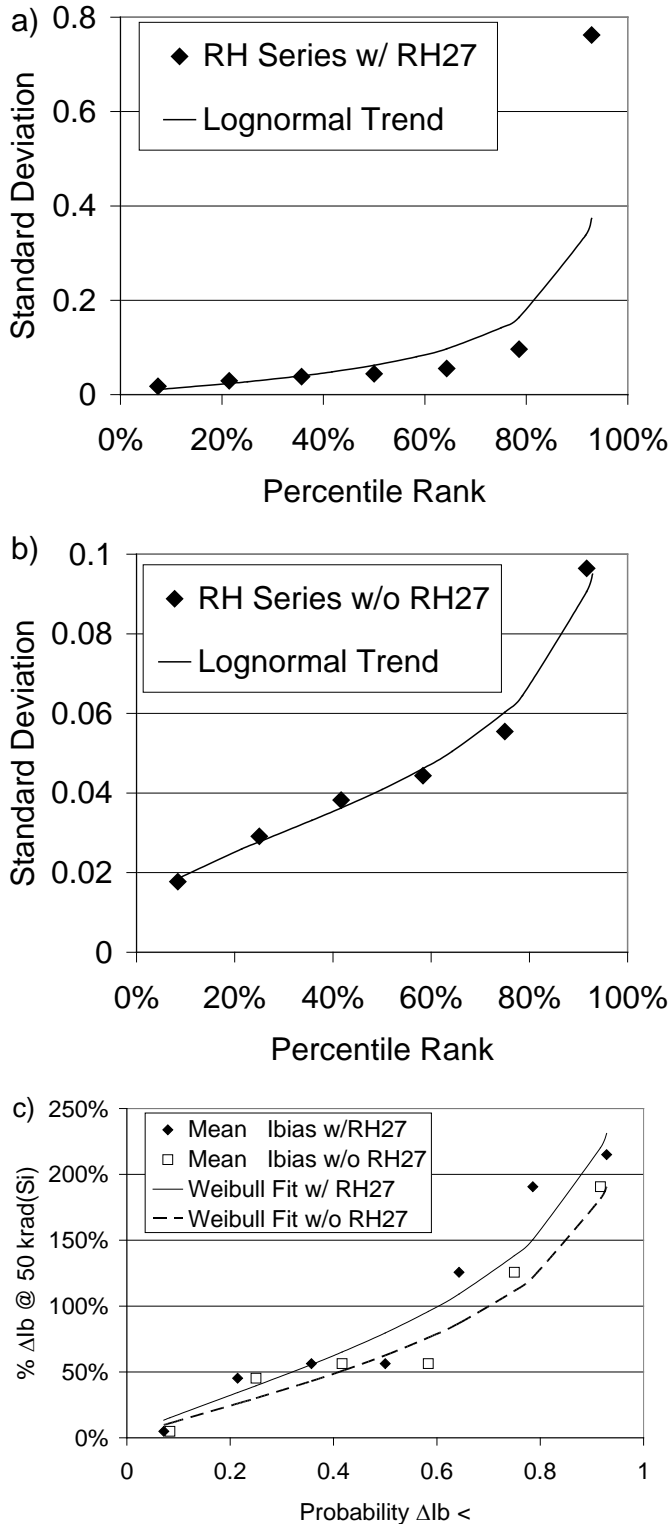
is suspended before failure is seen) provide poor constraints on the failure distribution, since it can be fit equally well by a narrow distribution centered just above the highest test levels or a broad distribution centered far above the failure levels. These two extremes give very different answers as we look at behavior in their tails. Even so, a fit of this data to a lognormal distribution (Fig. 3) shows that >99% of parts in the XFCB process will survive 45 krad(Si) with 90% confidence unless the failure distribution is extremely broad (e.g. even worse than a random failure distribution).

TABLE IV: **SUSPENSION DATA FOR ADI XFCB PARTS**

Part #	Type	Dose krad(Si) (Dose Rate)	# Parts
AD6640	ADC	>100 (50 rad(Si)/s)	5
AD6640	ADC	>1000	5
AD9042	ADC	>1000	5
AD8184	MUX	>103 (0.06 rad(Si)/s)	10
AD8063	Op Amp	>103 (0.06 rad(Si)/s)	10
AD8001	Op Amp	>100 (1.5 rad(Si)/s)	5
AD8001	Op Amp	>1000 (150 rad(Si)/s)	5
AD8009	Op Amp	>1000	5

μ \ σ	0.1	0.3	0.5	0.7	0.9
5.0					
5.2	136				
5.4	166	93			
5.6	203	114			
5.8	248	139	78		
6.0	302	170	96	54	
6.2	369	208	117	66	
6.4	451	254	143	80	45

Fig.3 Estimated lower-limit doses where 99% of XFCB parts would pass based on lognormal fits of Table IV's data. Parameters (μ,σ) within the 90% confidence interval are in unshaded cells. The analysis shows that the 99/90 lower limit for the process is 45 krad(Si) unless the failure distribution is exceptionally broad (σ>0.9), which is very unlikely given prior experience.

III. DISENTANGLING VARIABILITY

The data for multi-lot samples of ADI op amps in Table III allow us to estimate the contributions of part-to-part, lot-to-lot and part-type-to-part-type variation to the overall process variability. However, because our dataset is still small, we want a process that is robust and relatively insensitive to fluctuations so that it evolves smoothly as we add more data. Our basic philosophy can be summarized as: "When in doubt, average." Averages tend to be robust against minor sampling fluctuations and other random errors. As a measure of variability, we use the standard deviation (normal or lognormal), which, as the square root of the variance, or second moment, is also robust.

Fig.2 Rank plots of mean and standard deviation for %ΔIb. a) The RH27's standard deviation is an outlier, distorting the fit to the other part types (b). c) A Weibull form fits the mean increase in Ib, with the RH27 distorting the fit only slightly.

Even minimally restrictive data can constrain failure distributions. For the data in Table IV on ADI's eXtra-Fast Complementary Bipolar (XFCB) process,[10] we know only that all parts performed within specifications at the highest dose (column III). Such suspension data (that is where the test

We describe the method in figure 4. For each lot of each part type, we start with the mean failure dose and the standard deviation of failures about that mean. Then we fit these means and standard deviations (both positive definite quantities) to their own lognormal distributions for each part type. This results in 4 parameters as in reference 5:

1) μ_μ—the expected mean failure dose for an average lot of a given part type

2) μ_σ—the expected part-to-part variability (standard deviation) in an average lot for a given part type

3) σ_μ—lot-to-lot standard deviation of expected mean failure level for a given part type

4) σ_σ—lot-to-lot standard deviation of part-to-part variability for a given part type.

Then we fit each of these 4 parameters to appropriate distributions across part types, resulting in likelihoods for normal or lognormal parameters (one mean and one standard deviation) for each of the above variables (1-4):

1) $P(\mu_\mu, \mu_{\mu\mu}, \sigma_{\mu\mu})$ describes the distribution of μ_μ, across part types as a lognormal distribution with lognormal mean $\mu_{\mu\mu}$, and lognormal standard deviation $\sigma_{\mu\mu}$.

2) $P(\mu_\sigma, \mu_{\mu\sigma}, \sigma_{\mu\sigma})$ describes the distribution of μ_σ across parts as a lognormal distribution with lognormal mean $\mu_{\mu\sigma}$, and lognormal standard deviation $\sigma_{\mu\sigma}$.

3) $P(\sigma_\mu, \mu_{\sigma\mu}, \sigma_{\sigma\mu})$ describes the distribution of σ_μ, across parts as a lognormal distribution with lognormal mean $\mu_{\sigma\mu}$, and lognormal standard deviation $\sigma_{\sigma\mu}$.

4) $P(\sigma_\sigma, \mu_{\sigma\sigma}, \sigma_{\sigma\sigma})$ describes the distribution of σ_σ, across parts as a lognormal distribution with lognormal mean $\mu_{\sigma\sigma}$, and lognormal standard deviation $\sigma_{\sigma\sigma}$.

The parametric combination with the highest likelihood is the maximum likelihood distribution describing how the given variable is distributed across part types in the process under study. However, because our dataset is small, it is unlikely to produce a sharp maximum in the likelihood. As such, we use a likelihood-weighted model averaging approach similar to that in reference 11 and used later in 5, with weights given by:

$$w(\mu_{xy}, \sigma_{xy}) = \frac{L(\{\mathbf{z}\}, \mu_{xy}, \sigma_{xy})}{\int L(\{\mathbf{z}\}, \mu_{xy}, \sigma_{xy}) d\mu_{xy} d\sigma_{xy}} \quad (1)$$

where z represents the data.

Thus, the probability distribution for expected hardness is

$$P(\mu_\mu) = \int P(\mu_\mu, \mu_{\mu\mu}, \sigma_{\mu\mu}) w(\mu_{\mu\mu}, \sigma_{\mu\mu}) d\mu_{\mu\mu} d\sigma_{\mu\mu} \quad (2)$$

This method takes the contribution to $P(\mu_\mu)$ for each parametric combination $(\mu_{\mu\mu}, \sigma_{\mu\mu})$ weighted according to the likelihood of that combination based on our data. In many cases such model-averaging approaches can outperform maximum likelihood, because we preserve the contribution of every parametric combination, merely changing their weights as we add data, rather than selecting different parameters with each new addition.[11] As such, the averaged probability distribution $P(\mu_\mu)$ tends to evolve more continuously. We calculate similar averaged distributions over the other parameters μ_σ, σ_μ and σ_σ.

These distributions allow us to bound expected failure

dose, part-to-part and lot-to-lot variability to any desired confidence level for parts in the process. While the limited dataset we have here likely results in a broad distribution, this merely means that the bounds we estimate are likely conservative. The bounds will improve as we add data for more parts to constrain the model. Moreover, we can use the distributions to spot parts that may perform out of family for the process.
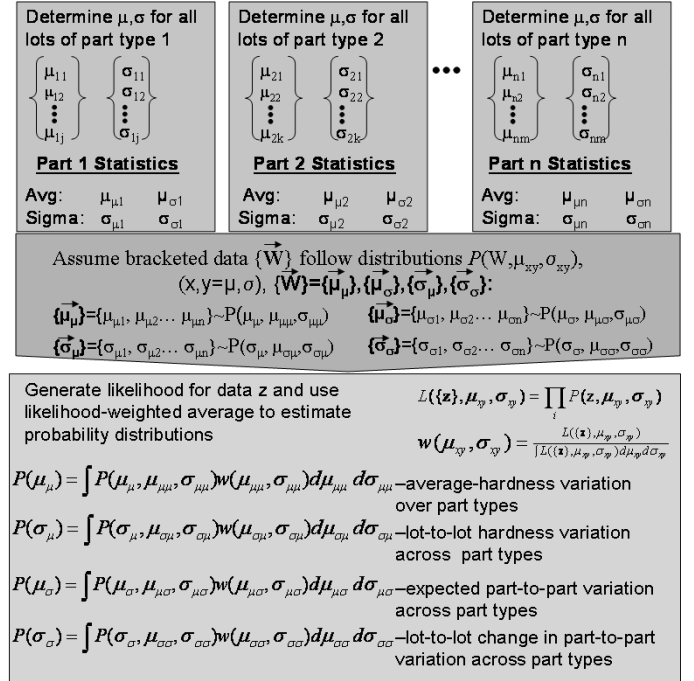

Fig. 4 (Top) Fitting the mean failure dose and standard deviation (μ_{ij}, σ_{ij}) for each lot j of each part type i to lognormal distributions yields 4 parameters for each part-type $(\mu_{\mu i}, \sigma_{\mu i}, \mu_{\sigma i}, \sigma_{\sigma i})$.

(middle) Fitting these four parameters across part types to suitable distributions yields 4 distributions that describe process variability in terms of 4 distributions, each with 2 parameters: $P(\mu_\mu, \mu_{\mu\mu}, \sigma_{\mu\mu})$, $P(\sigma_\mu, \mu_{\sigma\mu}, \sigma_{\sigma\mu})$, $P(\mu_\sigma, \mu_{\mu\sigma}, \sigma_{\mu\sigma})$ and $P(\sigma_\sigma, \mu_{\sigma\sigma}, \sigma_{\sigma\sigma})$.

(Bottom) Because our dataset is small, rather than taking the single parametric combination that maximizes likelihood, we perform a weighted average over all parametric combinations using likelihood weights for the fit parameters. The resulting distributions describe part-to-part, lot-to-lot variation over the range of similar part types in the process.

Applying this method to the data in Table III, the OP400 appears to exhibit greater variability than other parts--especially in its part-to-part standard deviation—from one lot to the next. As we did for the RH27, we initially perform the procedure outlined above excluding the OP400 data. Figures 5-8 summarize the various contributions to variability within the process (omitting the OP400). These curves show that while the OP400 is somewhat soft and variable in its lot mean hardness for the process (86% WC in expected mean hardness and 91% WC in lot-to-lot variation of mean hardness), lot-to-lot changes in part-to-part standard deviation are at the 95% WC level for the process. This indicates that the part could be out of family for the process. Certainly, inclusion of the OP400 results in much broader distributions except for that of mean hardness.
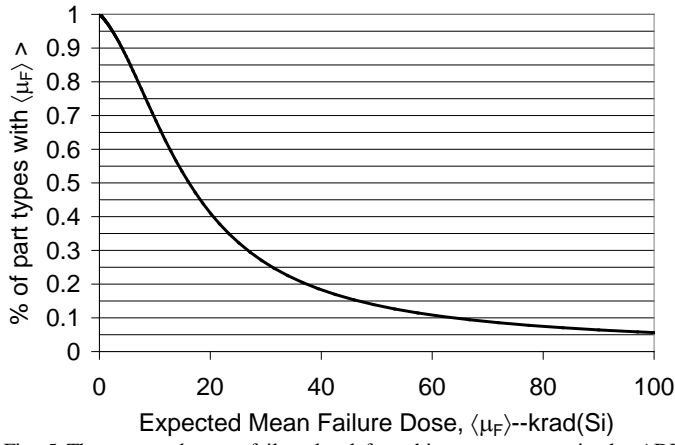
Fig. 5 The expected mean failure level for arbitrary part types in the ADI bipolar (>2.5 μm feature size) process.

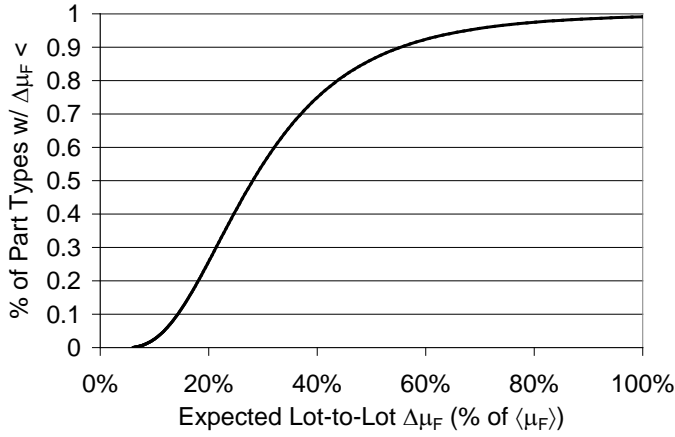Fig. 7 The distribution of expected part-to-part standard deviations of failure doses for arbitrary part types in the ADI bipolar (>2.5 μm) process.

Fig.6 The expected lot-to-lot variation of mean failure dose for arbitrary part types in ADI's bipolar (>2.5 μm feature size) process.

Fig. 8 The distribution of variation of part-to-part standard deviations of from lot to lot for arbitrary part types in the ADI bipolar (>2.5 μm) process.

Another part we would like to investigate is the OP484 quad op amp, which has been seen to exhibit significant lot-to-lot and part-to-part variability.[4],[5]   For example in one sample of 9 lots, the hardest lot had a mean failure level about 9× harder than the softest lot. This would place this part roughly at 97% worst-case for the process based on figure 6, and it places at about the 90th percentile for variability in standard deviation (based on figure 8).

Although we cannot say definitively based on our small sample size that the OP400 and OP484 are out of family for the ADI bipolar (>2.5 μm) family, these parts exhibit much more lot-to-lot variability that other parts we have investigated in the family. It is important to be able to identify such out-of-family parts produced in a process, because they preclude qualification of parts in the process based solely on similarity data—unless we understand why the parts are out of family. That we could spot the OP400 and OP484 even with the limited dataset available illustrates the model's robustness.
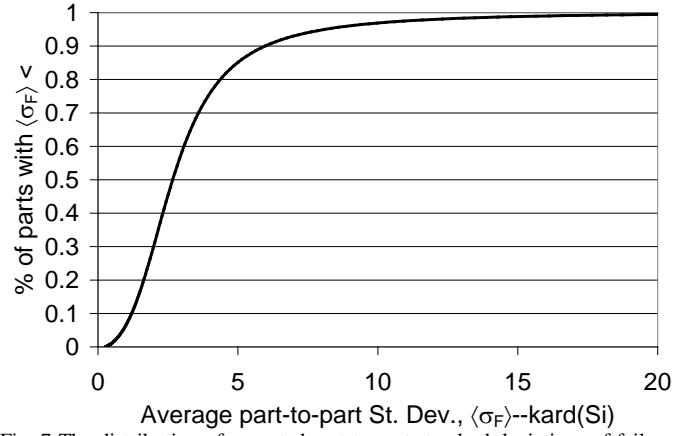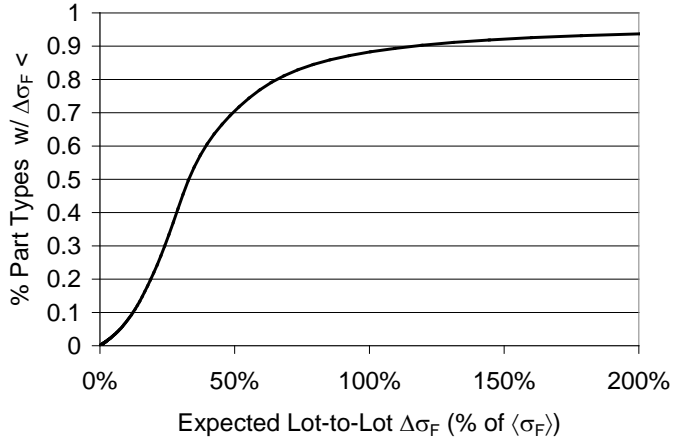
IV. INFERENCE WITHOUT MULTI-LOT DATA

The lack of multi-lot data for the LTC parts in Table II precludes disentangling lot-to-lot and part-type-to-part-type contributions to variability for the process. However, we can look at performance of a "typical" lot for a "typical" part type in the RH process. Figure 9 shows the distribution of lot mean ΔIb after 50 krad(Si) of low-dose-rate irradiation. The curves were generated using a fitting and averaging procedure similar to that used in section IV. Based on these curves, the RH27 does indeed appear out of family, falling at the 81st worst case (WC) percentile for mean increased bias current $\langle \% \Delta Ib \rangle$ and at the 99.8th WC percentile for part-to-part standard deviation σ_{ptp}.   Even so, the degradation and variability are mild compared to the commercial processes like the ADI process investigated in section IV.
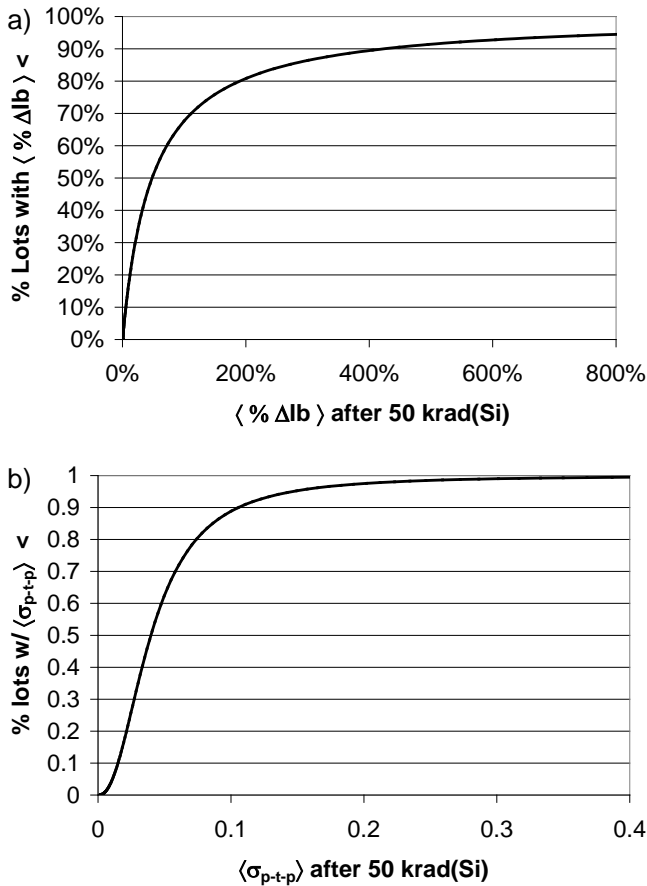
Fig. 9 Mean (a) and standard deviation (b) of percent increase in Ib $\langle \% \Delta Ib \rangle$ across part types and lots for LTC's RH process. 90% of parts would exhibit less than a 4x increase in Ib after 50krad(Si). Part-to part variation is small for the sample surveyed, with 90% of parts showing a part-to-part standard deviation σ_{ptp} on %ΔIb less than 0.11.

While the above plots do not separate the different contributions to variability, we noted above that the lot-to-lot consistency of the RH1014 quad op amp suggests that lot-to-lot variation should be mild for the RH process. As an example, mean ΔIb at 100 krad(Si) (standard dose rates) varied by only a factor of 2.6 over 38 lots.[4],[5] We assumed that ΔIb was lognormally distributed and found that this lot-to-lot variability equates to about half as much variability as we saw for the parts in Table II. If the RH1014 is typical for the process, this suggests that lot-to-lot variability and variation of response over part types are roughly commensurate in the LTC RH process. While this represents only a first guess that could change as we add data for more part types to refine our model, it demonstrates that even limited data can provide a useful working hypothesis to guide part selection while LTC continues to add to its list of parts tested at ELDRS rates.

V. CONCLUSIONS AND RECOMMENDATIONS

We have developed a statistical method that allows use of similarity data to bound likely radiation behavior of other parts in the process for which we do not yet have data or to supplement limited data for a part fabricated in the same process. The method is sufficiently robust that it can provide meaningful constraints on radiation behavior even with small, imperfect data sets drawn from public sources. Although the methods developed do not necessarily require multi-lot samples of data for a given part, we can begin to disentangle the contributions to variability for the process if we have at least three part types with data for three lots. The multi-lot analysis provides estimates for expected mean performance and part-to-part variability, as well as how these quantities vary from lot to lot and part type to part type.

Even if we lack data for multiple lots, we can still draw useful conclusions about how an "average lot" will perform— especially for radiation hardened part families like the LTC RH series of parts. Indeed, if the process is sufficiently hard, even suspension data like that for the ADI XFCB process in Table IV can place useful constraints on possible failure distributions.

Although similarity data has traditionally been used mainly for qualification decisions, the methods developed here have several other uses. Perhaps most important, they indicate when similarity data are inappropriate for qualification. Thus, the "out of family" performance of the OP484 and OP400 relative to other parts in the ADI bipolar (>2.5 µm) process is problematic because it is not understood, while the RH27's greater degradation relative to other parts in the RH process can be understood based on its reliance exclusively on process hardening. Also, because the method of necessity looks at all available data for parts in a process, it is useful for spotting process changes that may affect other parts in the process as well. Because the method provides separate estimates of part-to-part, lot-to-lot and part-type variation, it will likely also provide useful input for physics-based modeling for TID and for process and circuit hardening efforts.

Although our goal in this study has been to develop a method that is sufficiently robust to yield useful results even with small, imperfect datasets, the model's utility will improve significantly as we increase the dataset size or quality. In particular, if data are drawn from a long time series of lot qualification efforts with consistent test procedures and conditions, much greater precision is possible. Moreover, while here we have concentrated on a single parameter (Ibias), one can apply it across the board to all parameters or to different definitions of failure (e.g. functional).

The robustness of the method should also allow it to be used with data from different sources, provided that test conditions are known to be consistent. As a result, it increases the value of data sharing and of implementing consistent test methods that facilitate data interpretation as well as data sharing.

REFERENCES

[1] A. Namenson, "Lot uniformity and small sample sizes in hardness assurance," IEEE Trans. Nucl. Sci., vol. 35, no. 6, pp. 1506–1511, Dec. 1988
[2] "Ionizing dose and neutron hardness assurance guidelines for microcircuits and semiconductor devices", 2009. MIL-HDBK-814
[3] R. L. Pease, "Microelectronic Piece Part Radiation Hardness Assurance for Space Systems, Atlanta, GA, July, 19. 2004 NSREC Short Course.

[4] R. Ladbury, R. and J. L. Gorelick, "Statistical methods for large flight lots and ultra-high reliability applications," IEEE Trans. Nucl. Sci., Vol. 52, No. 6,pp. 2630 – 2637, 2005.

[5] R. Ladbury et al., "Statistical Model Selection for TID Hardness Assurance," IEEE Trans. Nucl Sci., vol. 56, no.6, pp. 3354-3360. Dec. 2009.

[6] Radhome Searchable database, http://radhome.gsfc.nasa.gov/radhome/RadDataBase/RadDataBase.html ,12/28/2009.

[7] Linear Technology Corp, ELDRS data: http://cds.linear.com/docs/Information%20Card/LTC%20Space%20Port folio.pdf , 11/23/2009.

[8] *Explor. Data Analysis, NIST/SEMATECH e-Handbook of Statistical Methods*http://www.itl.nist.gov/div898/handbook/eda/eda.htm, 12/28/2009.

[9] Rob Davies, Personal Communication.

[10] S. Burns, Personal Communication.

[11] K. P. Burnham and D. R. Anderson, Model Selection and Multimodel Inference: A Practical Information-Theoretical Approach, 2nd ed., New York: Springer-Verlag, 2002.