SAND2016-3281C

Where Error to Expect When You Are Expecting a Bit Flip

James Elliott[1,2], Mark Hoemmen[1], Frank Mueller[2]

[1]Sandia National Laboratories

[2]North Carolina State University

NC STATE UNIVERSITY

Department of Computer Science

U.S. DEPARTMENT OF ENERGY

NNSA

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000. SAND NO 2015-2090 C

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

Motivation

- Titan/Jaguar case study †
  - Constant stream of single bit flips
  - Double bit flip every 24 hours
  - 20 faults per hour
    - Heartbeat fault every 3 minutes
    - 12 kernel panics in 3 days

†Al Geist, Monster in the Closet, 2011

Motivation

- Titan/Jaguar case study †
  - Constant stream of single bit flips
  - Double bit flip every 24 hours
  - 20 faults per hour
    - Heartbeat fault every 3 minutes
    - 12 kernel panics in 3 days

Root cause analysis difficult …

Is it even necessary for algorithm designers?

†Al Geist, Monster in the Closet, 2011

Myth of a Reliable Machine

- Computers are inherently **unreliable**
- We design software and hardware that work given some tolerance
  - Windows crashes…but not too often
  - Screens flicker…but faster than the eye can see
  - Video games draw frames…but can skip some
- Numerical Mathematics has adapted
  - IEEE spec mandates behavior, standardizes representation
  - Integrity of approximations can be proven given assumptions promised by floating point (e.g., Higham'93)

Motivation Part 2

- How to build resilient algorithms?
  - Fault model allows <u>very small</u> and **very large** errors
- How to assess an algorithm's robustness to errors?
  - Common approach "create" bit flips by modifying input or output data

- If bit flips are randomly injected, what types of errors are these creating inside the app?

This talk explores the relative errors that bit flips can introduce.

- This work is motivated by the desire to both build and assess application's resilience to soft errors

IEEE Representation

- Scientific Notation $v = a \times b^m$
- Analytic Representation $$v = (-1)^{sign}\left(1 + \sum_{i=0}^{n-1} b_i 2^{i-n}\right) \times 2^{e-bias}$$
- Implementation (binary64)
  - 1 sign bit
  - 52 bits in the mantissa
  - 11 bits in the exponent

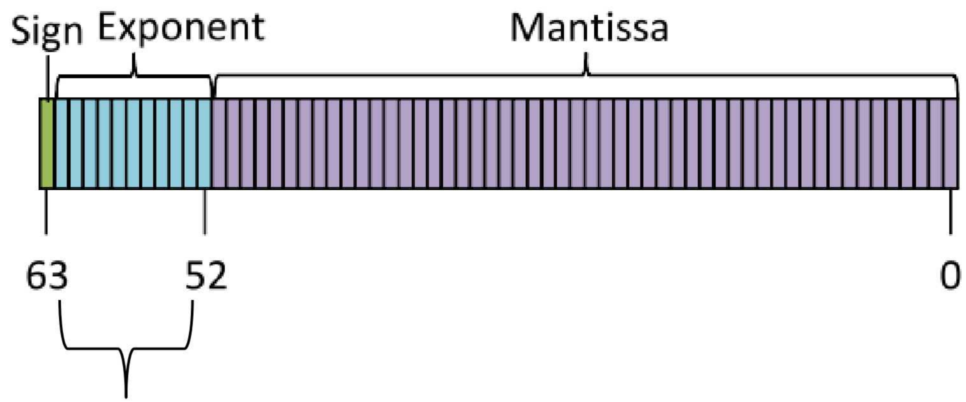

No explicit sign bit, bias of 1023 used

Example: $2^0 = 2^{1023-1023}$

IEEE Representation

- Analytic Representation

$$v = (-1)^{sign}\left(1 + \sum_{i=0}^{N-1} b_i 2^{i-N}\right) \times 2^{e-bias}$$

- Terminology

β	Trailing significand field
$1.\beta$	Complete significand $(1+\beta)$
$0.\xi$	Error from a bit flip in the significand
α	The exponent of the scalar $\left(2^{e-bias}\right)$
$bias$	Bias used for exponent storage
η	Error introduced from a corrupted biased exponent (2^{η})
N	Number of significand bits
M	Number of significand values (2^N)
ϵ	2^{-N}
Z	Number of exponent bits
K	Number of biased exponents

IEEE Representation

- Model sets
  - Significand: $1.\beta \in \{1 + 2^{-N}(M - i)\}$ for $i = 1, \dots, M$
  - Significand error: $0.\xi \in \left\{2^{-i}\right\}$ for $i = 1, \dots, N$
  - Exponent errors (2^{η}): $\eta \in \{\pm 2^0, \pm 2^1, \dots, \pm 2^{Z-1}\}$
    - $\eta^+ \in \{2^0, 2^1, \dots, 2^{Z-1}\}$ (corresponds to 0->1)
    - $\eta^- \in \{-2^0, -2^1, \dots, -2^{Z-1}\}$

IEEE Representation

- Model $v = (-1)^{sign} \alpha \times 1.\beta$
- Bit flip models
  - Significand: $\tilde{v} = \alpha \times 1.\tilde{\beta} = \alpha \times (1.\beta \pm 0.\xi)$
  - Exponent: $\tilde{v} = \tilde{\alpha} \times 1.\beta = \alpha \times 2^{\eta} \times 1.\beta$
    - If flip 0->1, then η is positive
    - If flip 1->0, then η is negative
  - Sign: $\tilde{v} = -v$

IEEE Representation

- Model $v = (-1)^{sign} \times \alpha \times 1.\beta$
- Bit flip models
  - Significand: $\tilde{v} = \alpha \times 1.\tilde{\beta} = \alpha \times (1.\beta \pm 0.\xi)$
  - Exponent: $\tilde{v} = \tilde{\alpha} \times 1.\beta = \alpha \times 2^{\eta} \times 1.\beta$
    - If flip 0->1, then η is positive
    - If flip 1->0, then η is negative
  - Sign: $\tilde{v} = -v$
- Model error measures

Bit Flip Location	Relative Error
Significand $(1.\tilde{\beta})$	$\frac{0.\xi}{1.\beta}$
Exponent $\tilde{\alpha}$	$\lvert 1 - 2^{\eta} \rvert$
Sign	2

Model Error Statistics

- What is the expected relative error for each classification given a bit flip?
  - Compute $\mathrm{E}\left[\frac{0.\xi}{1.\beta}\right]$, $\mathrm{E}[1 - 2^{\eta^-}]$, $\mathrm{E}[1 - 2^{\eta^+}]$, $\mathrm{E}[1 - 2^{\eta}]$
  - Introduce discrete R.V.s taking values from each set
    - X_ξ (significand errors), X_ω (inverse significand)
    - $X_\eta, X_{\eta^+}, X_{\eta^-}$

Model Error Statistics

- Expected relative error given a significand bit flip
  - Compute $\mathrm{E}\left[\frac{0.\xi}{1.\beta}\right] = \left[\mathrm{X}_{\xi}\mathrm{X}_{\omega}\right] = \mathrm{Cov}\left(\mathrm{X}_{\xi}, X_{\omega}\right) + \mathrm{E}\left[X_{\xi}\right]\mathrm{E}\left[X_{\omega}\right]$
  - Can show $\mathrm{Cov}\left(\mathrm{X}_{\xi}, X_{\omega}\right) = 0$
    - Should be zero! Significand error is independent of the significand
    - Example

$$1.5 = 2^{0} \times [1 + (1)2^{-1} + (0)2^{-2} + \cdots + (0)2^{-N}]$$

Flip most significand bit ($0.\xi = |-0.5| = 0.5$)

Absolute value forces errors to be independent (can prove rigorously)

  - $\mathrm{E}\left[\frac{0.\xi}{1.\beta}\right] = \mathrm{E}\left[X_{\xi}\right]\mathrm{E}\left[X_{\omega}\right]$

Model Error Statistics

- Expected relative error given a significand bit flip
  - $\mathrm{E}[X_\xi] = \frac{1}{N}\sum_{i=1}^{N} 2^{-i} = \frac{1}{N} - \frac{\epsilon}{N}$
  - $\mathrm{E}[X_\omega] = \frac{1}{M}\sum_{i=1}^{M} \frac{1}{1+2^{-N}(M-i)} = \ln(2) + \frac{\epsilon}{2}$
  - $\mathrm{E}\left[\frac{0.\xi}{1.\beta}\right] = \frac{\ln(2)}{N} + \frac{1-2\ln(2)}{2N}\epsilon - \frac{\epsilon^2}{2N}$

Model Error Statistics

- Expected relative error given a significand bit flip

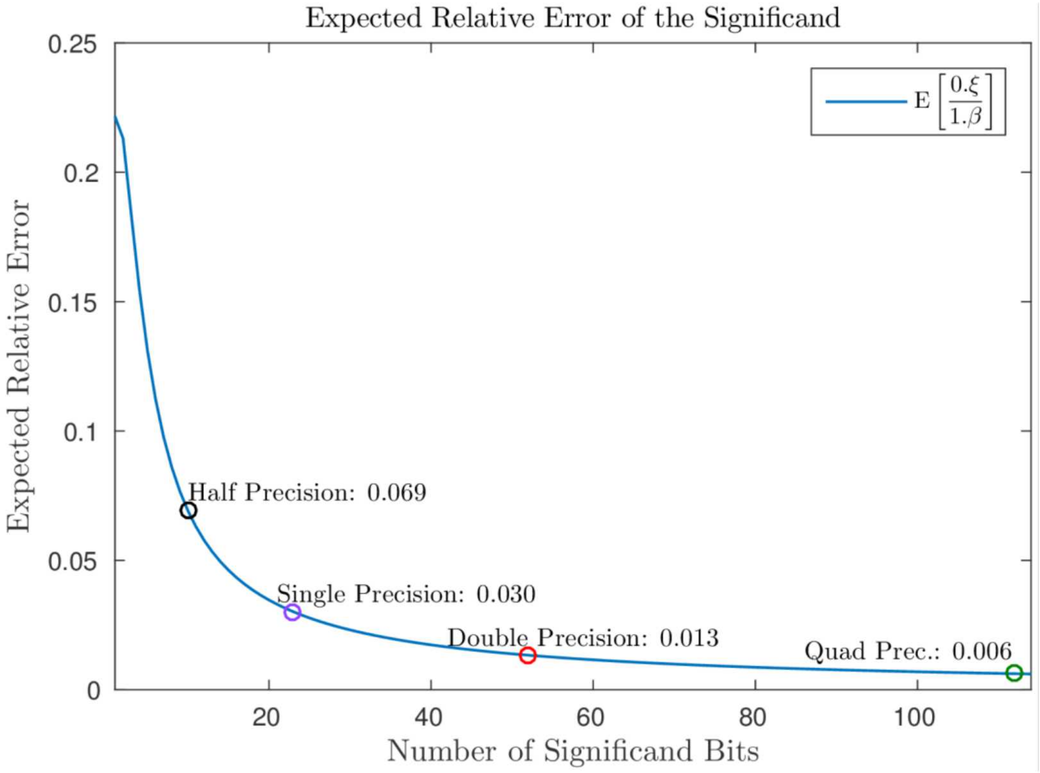

  - Relative error range $[\epsilon, 0.5]$

Model Error Statistics

- Expected relative error for an exponent bit flip?
  - All exponent flips: $\frac{2^{bias}}{z} < \mathrm{E}[1 - 2^{\eta}] < \frac{2^{bias+1}}{z}$
  - 0->1 exponent flips: $\frac{2^{bias+1}}{z} < \mathrm{E}\left[2^{\eta^{+}} - 1\right] < \frac{2^{bias+2} - 1}{z}$
  - 0->1 exponent flips: $\mathrm{E}[1 - 2^{\eta^{-}}] < 1$