



Analyses de terminaison des calculs flottants

Termination Analysis of Floating-Point Computations

Fonenantsoa MAURICA
Thèse de doctorat

Soutenue le 8 décembre 2017 devant le jury composé de :

Amir BEN-AMRAM	Examineur
Sylvie BOLDO	Rapporteuse
Christian DELHOMME	Examineur
Laure GONNORD	Examinatrice
Salvador LUCAS	Rapporteur
Frédéric MESNARD	Directeur
Marianne MORILLON	Examinatrice
Etienne PAYET	Co-directeur

Version courte pour RAIM 2018

ImageMagick




Fig 1: ImageMagick v7.0.4-9 non-terminate

percent_hue=-1E25

Care about termination.

Why does such bug occur?


“95% of folks out there are completely clueless about FP” Gosling.

The best for these folks would be an [easy-to-use](#) (eg [fully automatic](#)) tool that prevents bugs from FPs.

FP Numbers

Similar to the “scientific notation”:

$$\hat{x} \in \mathbb{F} : \hat{x} = \overbrace{(-1)^s}^{\text{thus } \pm 0} \underbrace{\hat{m}}_{\substack{\text{finite expansion} \\ \text{thus rounding}}} \beta^{\underbrace{e}_{\substack{\text{finite thus} \\ \hat{x}_{\min} \leq \hat{x} \leq \hat{x}_{\max}}}}$$



*difference between two consecutive FP numbers

Fig 2: A toy FP by `myfloat`. Co $x = 100\pi = 314.1592\dots$ corre

$$\hat{x} = 310 \text{ if rounding to neare} \quad 310 \oplus 1 = 310.$$

The standard for FP numbers is [IEEE-754](#). It defines: FP types, FP functions, rounding modes, exceptions.

Termination ***IS*** Decidable

And *many* non-trivial properties too!

Existence of RTE, most precise error bounds (even with transcendentals), synthesis of smallest possible programs...

This does **not** contradict Turing nor Rice. Contrarily to them, we talk here about machines with **finite amount of memory**.

But **too costly**:

Incompleteness is introduced for practical reasons.








Fig 3: Non-termination if there is a state visited twice

Termination by Abstracting to the Rationals



From left to right, top to bottom:

- **Computing \sqrt{t}**
Newton's method considers $f(x) = \frac{1}{2}(x + \frac{t}{x})$ and stops when $|u_{n+1} - u_n| < \delta$
- **When using FP**
Computations become tricky
Eg from open-vm-tools
- **Risk of non-termination?**
- **No**
Under the condition that $\delta > \Delta$ where $\Delta = w_{\infty}^+ - w_{\infty}^-$ is found after linearly abstracting the rounding function.
If $\delta \leq \Delta$, we don't know.

The PAATERM Algorithm

[Mythesis, Algo 3.2] computes **rational** abstractions $\mathcal{R}^\#$ of the transition relation \mathcal{R} by approximating the rounding function ζ_n with k -Piecewise Affine functions: $\forall x : \mu(x) \leq \zeta_n(x) \leq \nu(x)$.




Fig 4: $k = 1$ and $k = 3$

Refine $\mathcal{R}^\#$ until we can prove termination or until we **give up**.

The refinement increases k .
It is such that the bigger k ,
the bigger the set of Trace-Level Invariants,
the bigger the set of state-level variants,
the higher the “chance”
of proving termination.

The PAATERM Algorithm: Illustration

```

myfloat x = randFP ();
        y = randFP ();
while (x*y > 0 &
      y > -60) {
  x = x - 5.5;
  y = y + 3.5 + 3.5;
  if (x < -100) x = 100;
}

```

Fig 5: The program \mathcal{P}_{my}

\mathcal{P}_{my} always terminates because:

(TLINV1) x is alternately $+$ and $-$

(TLINV2) y eventually becomes $+$

	(TLINV1)	(TLINV2)
$k = 1$	✓	✗
$k = 2$	✗	✓
$k = 3$	✗	✓
$k = 4$	✓	✓

Fig 6: TLINV captured when running PAATERM

The fact that (TLINV1) is captured for $k = 1$ but lost for $k = 2$ is an exceptional case [Thm 3.1-3.2]

Practical Evaluation: the FTerm Analyzer

From	Verdict	Juggernaut	2LS	AProVe Julia	FTerm	Manual
bcache-tools	X	X				
gauche	X	X				
ghostscript	?					
glibc	?					
libpano	✓	✓			✓	
maxlib	?					
maxlib	✓				✓	
postfix	✓				✓	
postfix	✓				✓	
bash	✓	✓	✓		✓	
git	✓					✓
open-vm-tools	✓					✓

Fig 7: ✓: always terminates. X: there is at least one possible value of the variables that leads to non-termination.

?: can decide neither ✓ nor X. CO: time out after 300s on my i5-5200U 2.20GHz. Initial ranges of the variables are given.

Termination by (the Dumbest Possible) Brute-Force

Juggernaut's brute-force¹:

Uses advanced techniques: 2nd order SAT solver, genetic algorithms...
 ... Yet fails to decide termination of Newton's algorithm

What about the following brute-force?

```
float t, tnext = MIN-VALUE;
do {
  t = tnext;

  float v, v1 = t, d;
  do {
    v = v1;
    v1 = (v + t / v) / 2;
  } while (abs(v1 - v) > 0.5f);

  tnext = nextAfter(t);
} while (t < MAX-VALUE);
```

It terminates within minutes!

How to combine cleverness with dumbness?

1. Reduce the search space through clever abstractions
2. Dumbly brute-force when the remaining space is small enough

¹C. David et al. "Unrestricted Termin. and Non-term. Arguments for Bit-Vector Programs". ESOP 2015.

Exhibiting Non-Termination Bugs

Because people are more impressed by bug finding than by correctness proof

A simple **template-based approach** could be very effective.

For example $\text{while}(\underbrace{f + c_1 < c_2}_{\text{in } \mathbb{F}})\{\dots\}$ is quite common in the wild.

For a fully automatic non-termination exhibitor:

- 1 Slicing² for keeping the information related to termination
- 2 Template-based non-termination checking (intra-procedural)
for exhibiting non-termination condition
- 3 Backward value analysis (inter-procedural) for confirming that the bug is
reproducible from the main program

²D. Monniaux. "The Pitfalls of Verifying FP Computations". TOPLAS 2008.

Concluding Words

“Computer scientists worldwide are working hard on schemes to debug and verify software, especially in the context of parallel computation, but not FP. Why not?”

Kahan in “Desperately Needed Remedies for the Undebuggability of Large FP Computations in Science and Engineering” 2014