

From CRLibm to Metalibm : assisting the production of high-performance proven floating-point code


Florent de Dinechin
Arénaire/AriC project



The Arénaire project (soon to be renamed AriC) @ École Normale Supérieure de Lyon :

Computer Arithmetic at large

- Hardware and software
- From addition to linear algebra
- Fixed point, floating-point, multiple-precision, finite fields,
- Pervasive concern of **performance**, **numerical quality** and **validation**



Outline

Introduction : performance versus accuracy

Elementary function evaluation

Formal proof of floating-point code for the masses

Other tools toward MetaLibm

Conclusion

Introduction : performance versus accuracy

Introduction : performance versus accuracy

Elementary function evaluation

Formal proof of floating-point code for the masses

Other tools toward MetaLibm

Conclusion

Bottom line of this talk

Common wisdom

The more accurate you compute, the more expensive it gets

Bottom line of this talk

Common wisdom

The more accurate you compute, the more expensive it gets

In practice

- We (hopefully) remark it when our computation is not accurate enough.
- But do we remark it when it is **too accurate** for our needs?

Bottom line of this talk

Common wisdom

The more accurate you compute, the more expensive it gets

In practice

- We (hopefully) remark it when our computation is not accurate enough.
- But do we remark it when it is **too accurate** for our needs?

Reconciling performance and accuracy ?

Or, regain performance by computing just right ?

Double precision spoils us

The standard binary64 format (formerly known as double-precision) provides roughly **16** decimal digits.

Why should anybody need such accuracy?

Count the digits in the following

- Definition of the second : *the duration of **9,192,631,770** periods of the radiation corresponding to the transition between the two hyperfine levels of the ground state of the cesium 133 atom.*
- Definition of the metre : *the distance travelled by light in vacuum in $1/\mathbf{299,792,458}$ of a second.*
- Most accurate measurement ever (another atomic frequency) to 14 decimal places
- Most accurate measurement of the Planck constant to date : to 7 decimal places
- The gravitation constant G is known to 3 decimal places only

Parenthesis : then why binary64 ?

- This PC computes 10^9 operations per second (1 gigaflops)

Parenthesis : then why binary64 ?

- This PC computes 10^9 operations per second (1 gigaflops)

An allegory due to Kulisch

- print the numbers in 100 lines of 5 columns double-sided :
1000 numbers/sheet
- 1000 sheets \approx a heap of 10 cm
- 10^9 flops \approx heap height speed of 100m/s, or 360km/h
- A teraflops (10^{12} op/s) prints to the moon in one second
- Current top 500 computers reach the petaflop (10^{15} op/s)

Parenthesis : then why binary64 ?

- This PC computes 10^9 operations per second (1 gigaflops)

An allegory due to Kulisch

- print the numbers in 100 lines of 5 columns double-sided :
1000 numbers/sheet
- 1000 sheets \approx a heap of 10 cm
- 10^9 flops \approx heap height speed of 100m/s, or 360km/h
- A teraflops (10^{12} op/s) prints to the moon in one second
- Current top 500 computers reach the petaflop (10^{15} op/s)
- each operation may involve a relative error of 10^{-16} ,
and they accumulate.

Parenthesis : then why binary64 ?

- This PC computes 10^9 operations per second (1 gigaflops)

An allegory due to Kulisch

- print the numbers in 100 lines of 5 columns double-sided :
1000 numbers/sheet
- 1000 sheets \approx a heap of 10 cm
- 10^9 flops \approx heap height speed of 100m/s, or 360km/h
- A teraflops (10^{12} op/s) prints to the moon in one second
- Current top 500 computers reach the petaflop (10^{15} op/s)
- each operation may involve a relative error of 10^{-16} ,
and they accumulate.

Doesn't this sound wrong ?

We would use these 16 digits just to accumulate garbage in them ?

Back to the point

... which was :

Mastering accuracy for performance

When implementing a “computing core”

- A goal : *never compute more accurately than needed*

Back to the point

... which was :

Mastering accuracy for performance

When implementing a “computing core”

- A goal : *never compute more accurately than needed*
- Two sub-goals
 - Know what accuracy you need

... which was :

Mastering accuracy for performance

When implementing a “computing core”

- A goal : *never compute more accurately than needed*
- Two sub-goals
 - Know what accuracy you need
 - Know how accurate you compute

Back to the point

... which was :

Mastering accuracy for performance

When implementing a “computing core”

- A goal : *never compute more accurately than needed*
- Two sub-goals
 - Know what accuracy you need
 - Know how accurate you compute

“Computing cores” considered so far : **elementary functions**, sums of products, linear algebra, Euclidean lattices algorithms.

Elementary function evaluation

Introduction : performance versus accuracy

Elementary function evaluation

Formal proof of floating-point code for the masses

Other tools toward MetaLibm

Conclusion

How does your PC compute elementary functions ?

Rule of the game : use only $+$, $-$, \times

(and maybe $/$ and $\sqrt{\quad}$ but they are expensive).

How does your PC compute elementary functions ?

Rule of the game : use only $+$, $-$, \times

(and maybe $/$ and $\sqrt{\quad}$ but they are expensive).

- **Polynomial approximation** works on a small interval
 - for a fixed approximation error, d° grows with size of the interval
 - typically $x < 2^{-8} \implies d^\circ \approx 3 \dots 10$ ensures $\bar{\epsilon}_{\text{approx}} < 2^{-55}$

How does your PC compute elementary functions ?

Rule of the game : use only $+$, $-$, \times

(and maybe $/$ and $\sqrt{\quad}$ but they are expensive).

- **Polynomial approximation** works on a small interval
 - for a fixed approximation error, d° grows with size of the interval
 - typically $x < 2^{-8} \implies d^\circ \approx 3 \dots 10$ ensures $\bar{\epsilon}_{\text{approx}} < 2^{-55}$
- **Argument reduction** : using mathematical identities, transform large arguments in small ones

How does your PC compute elementary functions ?

Rule of the game : use only $+$, $-$, \times

(and maybe $/$ and $\sqrt{\quad}$ but they are expensive).

- **Polynomial approximation** works on a small interval
 - for a fixed approximation error, d° grows with size of the interval
 - typically $x < 2^{-8} \implies d^\circ \approx 3 \dots 10$ ensures $\bar{\epsilon}_{\text{approx}} < 2^{-55}$
- **Argument reduction** : using mathematical identities, transform large arguments in small ones

Simplistic example : an exponential

- identity : $e^{a+b} = e^a \times e^b$
- split $x = a + b$
 - a : k leading bits of x
 - b : lower bits of x $b \ll 1$
- tabulate all the e^a (2^k entries)
- use a Taylor polynomial for e^b

Know how accurate you compute

- Approximation errors
 - example : approximate a function f with a polynomial p :
 $\|p - f\|_\infty ?$
 - in general : approximate an object by another one

Know how accurate you compute

- Approximation errors
 - example : approximate a function f with a polynomial p :
 $\|p - f\|_\infty ?$
 - in general : approximate an object by another one
- Rounding errors
 - each individual error well specified by IEEE-754
 - but error accumulation difficult to manage

Know how accurate you compute

- Approximation errors
 - example : approximate a function f with a polynomial p :
 $\|p - f\|_\infty ?$
 - in general : approximate an object by another one
- Rounding errors
 - each individual error well specified by IEEE-754
 - but error accumulation difficult to manage
- In physics : time discretization errors, etc

Correctly rounded elementary functions

- IEEE-754 floating-point single or double-precision
- **Elementary functions** : sin, cos, exp, log, implemented in the “standard mathematical library” (`libm`)

Correctly rounded elementary functions

- IEEE-754 floating-point single or double-precision
- **Elementary functions** : sin, cos, exp, log, implemented in the “standard mathematical library” (`libm`)
- **Correctly rounded** : As perfect as can be, considering the finite nature of floating-point arithmetic
 - same standard of quality as $+$, \times , $/$, $\sqrt{\quad}$

Correctly rounded elementary functions

- IEEE-754 floating-point single or double-precision
- **Elementary functions** : sin, cos, exp, log, implemented in the “standard mathematical library” (`libm`)
- **Correctly rounded** : As perfect as can be, considering the finite nature of floating-point arithmetic
 - same standard of quality as $+$, \times , $/$, $\sqrt{\quad}$
- Now recommended by the IEEE754-2008 standard, but long considered **too expensive**
because of the **Table Maker's Dilemma**

The Table Maker's Dilemma


- Finite-precision algorithm for evaluating $f(x)$

The Table Maker's Dilemma

- Finite-precision algorithm for evaluating $f(x)$
- Approximation + rounding errors \rightarrow overall error bound $\bar{\epsilon}$.


The Table Maker's Dilemma

- Finite-precision algorithm for evaluating $f(x)$
- Approximation + rounding errors \rightarrow overall error bound $\bar{\varepsilon}$.
- What we compute : y such that $f(x) \in [y - \bar{\varepsilon}, y + \bar{\varepsilon}]$




The Table Maker's Dilemma

- Finite-precision algorithm for evaluating $f(x)$
- Approximation + rounding errors \rightarrow overall error bound $\bar{\varepsilon}$.
- What we compute : y such that $f(x) \in [y - \bar{\varepsilon}, y + \bar{\varepsilon}]$



The Table Maker's Dilemma


- Finite-precision algorithm for evaluating $f(x)$
- Approximation + rounding errors \rightarrow overall error bound $\bar{\varepsilon}$.
- What we compute : y such that $f(x) \in [y - \bar{\varepsilon}, y + \bar{\varepsilon}]$



Dilemma if this interval contains a midpoint between two FP numbers

The Table Maker's Dilemma

- Finite-precision algorithm for evaluating $f(x)$
- Approximation + rounding errors \rightarrow overall error bound $\bar{\varepsilon}$.
- What we compute : y such that $f(x) \in [y - \bar{\varepsilon}, y + \bar{\varepsilon}]$



Dilemma if this interval contains a midpoint between two FP numbers

The first digital signature algorithm

LOGARITHMICA.

25

Tabula inventimi Logarithmorum inferieur.

1	0,00	100001	0,00000,41419,2
2	0,30102,99915,6	100002	0,00000,86818,0
3	0,47712,12147,2	100003	0,00001,32226,4
4	0,60205,99993,3	100004	0,00001,77714,3
5	0,69897,00002,9	100005	0,00002,17413,8
6	0,77815,12102,8	100006	0,00002,60908,9
7	0,84509,80400,1	100007	0,00002,93995,5
8	0,90308,99869,9	100008	0,00003,47421,7
9	0,95424,35944,4	100009	0,00003,90847,4
10			
11	0,04139,16817,6	1000011	0,00000,2424,9
12	0,07918,12160,5	1000012	0,00000,86818,0
13	0,11918,11121,1	1000013	0,00000,13028,3
14	0,16121,80126,8	1000014	0,00000,17371,7
15	0,17609,12390,6	1000015	0,00000,21714,7
16	0,20411,99826,6	1000016	0,00000,26077,6
17	0,23041,89213,8	1000017	0,00000,30400,5
18	0,25127,23011,0	1000018	0,00000,34743,4
19	0,27877,16609,5	1000019	0,00000,39086,3
20			
10x	0,00431,1727,8	1000001	0,00000,00431,1
102	0,00860,01717,6	1000002	0,00000,00860,6
103	0,01283,72247,1	1000003	0,00000,01283,9
104	0,01701,33393,0	1000004	0,00000,01701,2
105	0,02118,99399,7	1000005	0,00000,02117,5
106	0,02530,60372,5	1000006	0,00000,02530,8
107	0,02938,17770,9	1000007	0,00000,02938,1
108	0,03344,37714,9	1000008	0,00000,03344,4
109	0,03744,64279,6	1000009	0,00000,03744,6
110			
1001	0,00043,40774,8	10000001	0,00000,00043,4
1002	0,00086,77211,1	10000002	0,00000,00086,9
1003	0,00130,09330,2	10000003	0,00000,00130,3
1004	0,00173,37128,1	10000004	0,00000,00173,7
1005	0,00216,60617,6	10000005	0,00000,00217,1
1006	0,00259,79807,2	10000006	0,00000,00256,6
1007	0,00302,94705,5	10000007	0,00000,00304,0
1008	0,00346,05321,1	10000008	0,00000,00347,4
1009	0,00389,11663,4	10000009	0,00000,00390,9
110			
10001	0,00004,14272,8	100000001	0,00000,00004,3
10002	0,00008,68102,1	100000002	0,00000,00008,7
10003	0,00013,02488,1	100000003	0,00000,00013,0
10004	0,00017,16830,6	100000004	0,00000,00017,4
10005	0,00021,70229,7	100000005	0,00000,00021,7
10006	0,00026,04985,5	100000006	0,00000,00026,1
10007	0,00030,15997,8	100000007	0,00000,00030,4
10008	0,00034,73966,9	100000008	0,00000,00034,7
10009	0,00039,08932,8	100000009	0,00000,00039,1

The first digital signature algorithm

LOGARITHMICA.

Tabula inventimi Logarithmorum inferieur.

1	0,00	100001	0,00000,41419,2
2	0,30102,99915,6	100002	0,00000,86818,0
3	0,47712,12147,2	100003	0,00001,32246,4
4	0,60205,99993,3	100004	0,00001,77714,3
5	0,69897,00002,9	100005	0,00002,23143,8
6	0,77815,12103,8	100006	0,00002,68568,9
7	0,84509,80400,1	100007	0,00003,13995,5
8	0,90308,99869,9	100008	0,00003,59421,7
9	0,95424,37944,4	100009	0,00004,04847,4
11	0,04139,16817,6	100011	0,00004,50272,9
12	0,07918,12160,5	100012	0,00004,95698,9
13	0,11918,12121,1	100013	0,00005,41124,3
14	0,16121,80126,8	100014	0,00005,86550,7
15	0,17609,12190,6	100015	0,00006,31976,7
16	0,20411,99826,6	100016	0,00006,77402,6
17	0,23041,89213,8	100017	0,00007,22828,6
18	0,25127,27011,0	100018	0,00007,68254,4
19	0,27877,16099,5	100019	0,00008,13680,3
101	0,00431,17274,8	1000001	0,00000,00431,1
102	0,00860,01717,6	1000002	0,00000,00862,6
103	0,01283,72247,1	1000003	0,00000,01283,9
104	0,01701,33393,0	1000004	0,00000,01727,2
105	0,02118,93999,7	1000005	0,00000,02171,5
106	0,02533,54672,5	1000006	0,00000,02625,8
107	0,02946,15376,9	1000007	0,00000,03080,1
108	0,03344,17714,9	1000008	0,00000,03534,4
109	0,03744,64279,6	1000009	0,00000,03988,6
1001	0,00043,40774,8	10000001	0,00000,00043,4
1002	0,00086,77211,1	10000002	0,00000,00086,9
1003	0,00130,09330,1	10000003	0,00000,00130,3
1004	0,00173,37128,1	10000004	0,00000,00173,7
1005	0,00216,60617,6	10000005	0,00000,00217,1
1006	0,00259,79807,2	10000006	0,00000,00260,6
1007	0,00302,94705,5	10000007	0,00000,00304,0
1008	0,00346,05321,1	10000008	0,00000,00347,4
1009	0,00389,11663,4	10000009	0,00000,00390,9
10001	0,00004,12172,8	100000001	0,00000,00004,3
10002	0,00008,68702,1	100000002	0,00000,00008,7
10003	0,00013,02488,1	100000003	0,00000,00013,0
10004	0,00017,16830,6	100000004	0,00000,00017,4
10005	0,00021,70029,7	100000005	0,00000,00021,7
10006	0,00026,04985,5	100000006	0,00000,00026,1
10007	0,00030,15997,8	100000007	0,00000,00030,4
10008	0,00034,73966,9	100000008	0,00000,00034,7
10009	0,00039,08932,8	100000009	0,00000,00039,1

15

- I want 12 significant digits

The first digital signature algorithm

LOGARITHMICA.

Tabula inventum Logarithmorum inferiorem.

1	0,00	100001	0,0000041429,2
2	0,30102,99915,6	100002	0,00000,86818,0
3	0,47712,12147,2	100003	0,00001,30286,4
4	0,60205,99993,3	100004	0,00001,73714,3
5	0,69897,00002,9	100005	0,00002,17141,8
6	0,77815,12103,8	100006	0,00002,60569,3
7	0,84509,80400,1	100007	0,00003,03997,7
8	0,90308,99869,9	100008	0,00003,47424,7
9	0,95424,35944,4	100009	0,00003,90847,4
10	0,94139,16877,6	100010	0,00004,3424,9
11	0,97128,12160,5	100011	0,00004,7763,9
12	0,98194,12121,2	100012	0,00005,2102,8
13	0,98612,80165,8	100013	0,00005,6441,7
14	0,98769,12190,6	100014	0,00006,0780,7
15	0,98811,99826,6	100015	0,00006,5119,6
16	0,98841,89213,8	100016	0,00007,0458,6
17	0,98857,25011,0	100017	0,00007,4797,5
18	0,98867,36093,7	100018	0,00008,0136,5
19	0,98871,36093,7	100019	0,00008,4475,4
20	0,98871,36093,7	100020	0,00008,8814,3
21	0,98869,01717,6	100021	0,00009,3153,2
22	0,98860,01717,6	100022	0,00009,7492,2
23	0,98851,72247,1	100023	0,00010,1831,1
24	0,98841,33393,0	100024	0,00010,6170,1
25	0,98831,93999,7	100025	0,00011,0509,0
26	0,98821,54606,4	100026	0,00011,4848,0
27	0,98811,15213,1	100027	0,00011,9186,9
28	0,98801,75819,8	100028	0,00012,3525,9
29	0,98791,36426,4	100029	0,00012,7864,8
30	0,98781,42774,8	100030	0,00013,2203,8
31	0,98771,72211,1	100031	0,00013,6542,7
32	0,98761,33357,0	100032	0,00014,0881,7
33	0,98751,93963,7	100033	0,00014,5220,6
34	0,98741,54570,4	100034	0,00014,9559,6
35	0,98731,15177,1	100035	0,00015,3898,5
36	0,98721,75783,8	100036	0,00015,8237,5
37	0,98711,36390,5	100037	0,00016,2576,4
38	0,98701,96997,2	100038	0,00016,6915,4
39	0,98691,57604,0	100039	0,00017,1254,3
40	0,98681,18210,7	100040	0,00017,5593,3
41	0,98671,78817,4	100041	0,00017,9932,2
42	0,98661,39424,1	100042	0,00018,4271,2
43	0,98651,00030,8	100043	0,00018,8610,1
44	0,98641,60637,5	100044	0,00019,2949,1
45	0,98631,21244,2	100045	0,00019,7288,0
46	0,98621,81851,0	100046	0,00020,1627,0
47	0,98611,42457,7	100047	0,00020,5965,9
48	0,98601,03064,4	100048	0,00021,0304,9
49	0,98591,63671,1	100049	0,00021,4643,8
50	0,98581,24277,8	100050	0,00021,8982,8
51	0,98571,84884,5	100051	0,00022,3321,7
52	0,98561,45491,2	100052	0,00022,7660,7
53	0,98551,06097,9	100053	0,00023,2000,6
54	0,98541,66704,6	100054	0,00023,6339,6
55	0,98531,27311,3	100055	0,00024,0678,5
56	0,98521,87918,0	100056	0,00024,5017,5
57	0,98511,48524,7	100057	0,00024,9356,4
58	0,98501,09131,4	100058	0,00025,3695,4
59	0,98491,69738,1	100059	0,00025,8034,3
60	0,98481,30344,8	100060	0,00026,2373,3
61	0,98471,90951,5	100061	0,00026,6712,2
62	0,98461,51558,2	100062	0,00027,1051,2
63	0,98451,12164,9	100063	0,00027,5390,1
64	0,98441,72771,6	100064	0,00027,9729,1
65	0,98431,33378,3	100065	0,00028,4068,0
66	0,98421,93985,0	100066	0,00028,8407,0
67	0,98411,54591,7	100067	0,00029,2746,0
68	0,98401,15198,4	100068	0,00029,7084,9
69	0,98391,75805,1	100069	0,00030,1423,9
70	0,98381,36411,8	100070	0,00030,5762,8
71	0,98371,97018,5	100071	0,00031,0101,8
72	0,98361,57625,2	100072	0,00031,4440,7
73	0,98351,18231,9	100073	0,00031,8779,7
74	0,98341,78838,6	100074	0,00032,3118,6
75	0,98331,39445,3	100075	0,00032,7457,6
76	0,98321,00052,0	100076	0,00033,1796,5
77	0,98311,60658,7	100077	0,00033,6135,5
78	0,98301,21265,4	100078	0,00034,0474,4
79	0,98291,81872,1	100079	0,00034,4813,4
80	0,98281,42478,8	100080	0,00034,9152,3
81	0,98271,03085,5	100081	0,00035,3491,3
82	0,98261,63692,2	100082	0,00035,7830,2
83	0,98251,24298,9	100083	0,00036,2169,2
84	0,98241,84905,6	100084	0,00036,6508,1
85	0,98231,45512,3	100085	0,00037,0847,1
86	0,98221,06119,0	100086	0,00037,5186,0
87	0,98211,66725,7	100087	0,00037,9525,0
88	0,98201,27332,4	100088	0,00038,3864,0
89	0,98191,87939,1	100089	0,00038,8202,9
90	0,98181,48545,8	100090	0,00039,2541,9
91	0,98171,09152,5	100091	0,00039,6880,8
92	0,98161,69759,2	100092	0,00040,1219,8
93	0,98151,30365,9	100093	0,00040,5558,7
94	0,98141,90972,6	100094	0,00040,9897,7
95	0,98131,51579,3	100095	0,00041,4236,6
96	0,98121,12186,0	100096	0,00041,8575,6
97	0,98111,72792,7	100097	0,00042,2914,5
98	0,98101,33399,4	100098	0,00042,7253,5
99	0,98091,94006,1	100099	0,00043,1592,4
100	0,98081,54612,8	100100	0,00043,5931,4

15

- I want 12 significant digits
- I have an approximation scheme that provides 14 digits

The first digital signature algorithm

LOGARITHMICA.

Tabula inventimi Logarithmorum inferiour.

1	0,00	100001	0,00000,41419,2
2	0,30102,99915,6	100002	0,00000,82838,0
3	0,47712,12147,2	100003	0,00001,24256,4
4	0,60205,99903,3	100004	0,00001,65674,8
5	0,69897,00002,9	100005	0,00002,07093,2
6	0,77815,12103,8	100006	0,00002,48511,6
7	0,84509,80400,1	100007	0,00002,89929,9
8	0,90308,99869,9	100008	0,00003,31348,2
9	0,95424,29944,4	100009	0,00003,72766,6
10	1,00000,00000,0	100010	0,00004,14185,0
11	0,04139,26877,6	100011	0,00004,55603,4
12	0,07918,12460,5	100012	0,00004,97021,8
13	0,11918,21212,1	100013	0,00005,38440,2
14	0,16122,30126,8	100014	0,00005,79858,6
15	0,17609,12390,6	100015	0,00006,21277,0
16	0,20411,99826,6	100016	0,00006,62695,4
17	0,23041,89213,8	100017	0,00007,04113,8
18	0,25122,29071,0	100018	0,00007,45532,2
19	0,27877,26093,3	100019	0,00007,86950,6
20	0,30213,17272,8	100020	0,00008,28369,0
21	0,32186,01717,6	100021	0,00008,69787,4
22	0,33812,72247,1	100022	0,00009,11205,8
23	0,35170,23393,0	100023	0,00009,52624,2
24	0,36301,99997,7	100024	0,00009,94042,6
25	0,37239,86872,5	100025	0,00010,35461,0
26	0,38033,27705,9	100026	0,00010,76879,4
27	0,38714,77714,9	100027	0,00011,18297,8
28	0,39314,64279,6	100028	0,00011,59716,2
29	0,39843,40774,8	100029	0,00012,01134,6
30	0,40308,77215,1	100030	0,00012,42553,0
31	0,40713,37128,1	100031	0,00012,83971,4
32	0,41066,60617,6	100032	0,00013,25389,8
33	0,41379,79807,2	100033	0,00013,66808,2
34	0,41653,04707,5	100034	0,00014,08226,6
35	0,41896,07321,1	100035	0,00014,49645,0
36	0,42119,11663,4	100036	0,00014,91063,4
37	0,42324,14272,8	100037	0,00015,32481,8
38	0,42511,68702,1	100038	0,00015,73899,2
39	0,42681,30283,1	100039	0,00016,15317,6
40	0,42835,16830,6	100040	0,00016,56736,0
41	0,42975,70229,7	100041	0,00016,98154,4
42	0,43104,04875,5	100042	0,00017,39572,8
43	0,43223,19977,8	100043	0,00017,80991,2
44	0,43334,79665,9	100044	0,00018,22409,6
45	0,43438,08923,8	100045	0,00018,63828,0
46	0,43534,18751,5	100046	0,00019,05246,4
47	0,43623,19149,7	100047	0,00019,46664,8
48	0,43706,19117,2	100048	0,00019,88083,2
49	0,43783,18754,9	100049	0,00020,29501,6
50	0,43855,18071,6	100050	0,00020,70920,0

15

- I want 12 significant digits
- I have an approximation scheme that provides 14 digits
- or,

$$y = \log(x) \pm 10^{-14}$$

The first digital signature algorithm

LOGARITHMICA.

Tabula inventimi Logarithmorum inferiour.

1	0,00	100001	0,00000,41419,2
2	0,30102,99915,6	100002	0,00000,82838,0
3	0,47712,12147,2	100003	0,00001,24256,4
4	0,60205,99903,3	100004	0,00001,65674,8
5	0,69897,00002,9	100005	0,00002,07093,2
6	0,77815,12102,8	100006	0,00002,48511,6
7	0,84509,80400,1	100007	0,00002,89929,9
8	0,90308,99869,9	100008	0,00003,31348,2
9	0,95424,25044,4	100009	0,00003,72766,6
10			
11	0,04139,16877,6	100010	0,00004,14185,0
12	0,07918,12160,5	100011	0,00004,55603,4
13	0,11918,12121,1	100012	0,00004,97021,8
14	0,16121,30126,8	100013	0,00005,38440,2
15	0,17609,12130,6	100014	0,00005,79858,6
16	0,20411,99826,6	100015	0,00006,21277,0
17	0,23041,89213,8	100016	0,00006,62695,4
18	0,25127,25071,0	100017	0,00007,04113,8
19	0,27877,16099,5	100018	0,00007,45532,2
20			
10x	0,00432,12727,8	100019	0,00007,86950,6
102	0,00860,01717,6	100020	0,00008,28369,0
103	0,01283,72247,1	100021	0,00008,69787,4
104	0,01703,33393,0	100022	0,00009,11205,8
105	0,02118,93999,7	100023	0,00009,52624,2
106	0,02533,54606,4	100024	0,00009,94042,6
107	0,02948,15213,1	100025	0,00010,35461,0
108	0,03363,75819,9	100026	0,00010,76879,4
109	0,03778,36426,6	100027	0,00011,18297,8
110			
1001	0,00043,40774,8	100028	0,00011,59716,2
1002	0,00086,72151,1	100029	0,00012,01134,6
1003	0,00130,03527,4		
1004	0,00173,34903,7	100030	0,00012,42553,0
1005	0,00216,66279,9	100031	0,00012,83971,4
1006	0,00259,97656,2	100032	0,00013,25389,8
1007	0,00302,29032,5	100033	0,00013,66808,2
1008	0,00345,60408,8	100034	0,00014,08226,6
1009	0,00389,91785,1	100035	0,00014,49645,0
1100			
10001	0,00004,14272,8	100036	0,00014,91063,4
10002	0,00008,28545,1	100037	0,00015,32481,8
10003	0,00013,42818,1	100038	0,00015,73899,2
10004	0,00017,57091,6	100039	0,00016,15317,6
10005	0,00021,71364,2		
10006	0,00025,85637,7	100040	0,00016,56736,0
10007	0,00030,00000,0	100041	0,00016,98154,4
10008	0,00034,14272,8	100042	0,00017,39572,8
10009	0,00038,28545,1	100043	0,00017,80991,2
10010	0,00042,42818,1	100044	0,00018,22409,6
10011	0,00046,57091,6	100045	0,00018,63828,0
10012	0,00050,71364,2	100046	0,00019,05246,4
10013	0,00054,85637,7	100047	0,00019,46664,8
10014	0,00058,00000,0	100048	0,00019,88083,2
10015	0,00062,14272,8	100049	0,00020,29501,6
10016	0,00066,28545,1		
10017	0,00070,42818,1	100050	0,00020,70920,0
10018	0,00074,57091,6	100051	0,00021,12338,4
10019	0,00078,71364,2	100052	0,00021,53756,8
10020	0,00082,85637,7	100053	0,00021,95175,2
10021	0,00086,00000,0	100054	0,00022,36593,6
10022	0,00090,14272,8	100055	0,00022,78012,0
10023	0,00094,28545,1	100056	0,00023,19430,4
10024	0,00098,42818,1	100057	0,00023,60848,8
10025	0,00102,57091,6	100058	0,00024,02267,2
10026	0,00106,71364,2	100059	0,00024,43685,6
10027	0,00110,85637,7		
10028	0,00114,00000,0	100060	0,00024,85104,0
10029	0,00118,14272,8	100061	0,00025,26522,4
10030	0,00122,28545,1	100062	0,00025,67940,8
10031	0,00126,42818,1	100063	0,00026,09359,2
10032	0,00130,57091,6	100064	0,00026,50777,6
10033	0,00134,71364,2	100065	0,00026,92196,0
10034	0,00138,85637,7	100066	0,00027,33614,4
10035	0,00142,00000,0	100067	0,00027,75032,8
10036	0,00146,14272,8	100068	0,00028,16451,2
10037	0,00150,28545,1	100069	0,00028,57869,6
10038	0,00154,42818,1		
10039	0,00158,57091,6	100070	0,00028,99288,0
10040	0,00162,71364,2	100071	0,00029,40706,4
10041	0,00166,85637,7	100072	0,00029,82124,8
10042	0,00170,00000,0	100073	0,00030,23543,2
10043	0,00174,14272,8	100074	0,00030,64961,6
10044	0,00178,28545,1	100075	0,00031,06380,0
10045	0,00182,42818,1	100076	0,00031,47798,4
10046	0,00186,57091,6	100077	0,00031,89216,8
10047	0,00190,71364,2	100078	0,00032,30635,2
10048	0,00194,85637,7	100079	0,00032,72053,6
10049	0,00198,00000,0	100080	0,00033,13472,0
10050	0,00202,14272,8	100081	0,00033,54890,4
10051	0,00206,28545,1	100082	0,00033,96308,8
10052	0,00210,42818,1	100083	0,00034,37727,2
10053	0,00214,57091,6	100084	0,00034,79145,6
10054	0,00218,71364,2	100085	0,00035,20564,0
10055	0,00222,85637,7	100086	0,00035,61982,4
10056	0,00226,00000,0	100087	0,00036,03400,8
10057	0,00230,14272,8	100088	0,00036,44819,2
10058	0,00234,28545,1	100089	0,00036,86237,6
10059	0,00238,42818,1	100090	0,00037,27656,0
10060	0,00242,57091,6		

15

- I want 12 significant digits
- I have an approximation scheme that provides 14 digits
- or,

$$y = \log(x) \pm 10^{-14}$$

- "Usually" that's enough to round

$$y = x, \text{xxxxxxxxxxxx}17 \pm 10^{-14}$$

$$y = x, \text{xxxxxxxxxxxx}83 \pm 10^{-14}$$

The first digital signature algorithm

LOGARITHMICA.

Tabula inventimi Logarithmorum inferiour.

1	0,00	100001	0,0000041419,1
2	0,30102,99915,6	100002	0,00000,86818,0
3	0,47712,12147,1	100003	0,00001,13236,4
4	0,60205,99903,3	100004	0,00001,17314,3
5	0,69897,00149,4	100005	0,00002,17641,8
6	0,77815,12101,8	100006	0,00002,60968,9
7	0,84509,86400,1	100007	0,00001,35991,5
8	0,90308,99869,9	100008	0,00002,47421,7
9	0,95424,15094,4	100009	0,00002,90847,4
11	0,04139,16817,6	1000011	0,00000,24542,9
12	0,07918,12160,5	1000012	0,00000,38639,9
13	0,11918,11311,1	1000013	0,00000,13028,3
14	0,14612,30165,8	1000014	0,00001,17127,7
15	0,17609,12390,6	1000015	0,00000,21714,7
16	0,20411,99826,6	1000016	0,00000,26077,6
17	0,23041,89213,8	1000017	0,00000,30400,5
18	0,25527,23011,0	1000018	0,00000,34741,4
19	0,27877,16609,5	1000019	0,00000,39083,3
10x	0,00431,13727,8	1000001	0,00000,00431,1
102	0,00860,01717,6	1000002	0,00000,00860,6
103	0,01283,72247,1	1000003	0,00000,01283,9
104	0,01701,33393,0	1000004	0,00000,01701,3
105	0,02118,99999,7	1000005	0,00000,02118,1
106	0,02533,68873,5	1000006	0,00000,02533,8
107	0,02948,37776,9	1000007	0,00000,02948,1
108	0,03364,17714,9	1000008	0,00000,03364,4
109	0,03781,64279,6	1000009	0,00000,03781,6
1001	0,00043,40774,8	10000011	0,00000,00043,4
1002	0,00086,77211,1	10000012	0,00000,00086,9
1003	0,00130,93302,1	10000013	0,00000,00130,3
1004	0,00173,17128,1	10000014	0,00000,00173,7
1005	0,00216,40617,6	10000015	0,00000,00216,1
1006	0,00259,79807,2	10000016	0,00000,00259,6
1007	0,00302,94701,5	10000017	0,00000,00302,0
1008	0,00346,01311,1	10000018	0,00000,00346,4
1009	0,00389,11663,4	10000019	0,00000,00389,9
10001	0,00004,12172,8	100000001	0,00000,00004,3
10002	0,00008,68702,1	100000002	0,00000,00008,7
10003	0,00013,02328,1	100000003	0,00000,00013,0
10004	0,00017,16830,6	100000004	0,00000,00017,4
10005	0,00021,70029,7	100000005	0,00000,00021,7
10006	0,00026,04981,5	100000006	0,00000,00026,1
10007	0,00030,15997,8	100000007	0,00000,00030,4
10008	0,00034,73966,9	100000008	0,00000,00034,7
10009	0,00039,08913,8	100000009	0,00000,00039,1

15

- I want 12 significant digits
- I have an approximation scheme that provides 14 digits

- or,

$$y = \log(x) \pm 10^{-14}$$

- “Usually” that’s enough to round

$$y = x, \text{xxxxxxxxxxxx}17 \pm 10^{-14}$$

$$y = x, \text{xxxxxxxxxxxx}83 \pm 10^{-14}$$

- **Dilemma** when

$$y = x, \text{xxxxxxxxxxxx}50 \pm 10^{-14}$$

The first digital signature algorithm

LOGARITHMICA.

Tabula inventimi Logarithmorum inferiorem.

1	0,00	100001	0,00000,014129,2
2	0,30102,99915,6	100002	0,00000,00818,0
3	0,47712,12147,2	100003	0,00001,00226,4
4	0,60205,99995,3	100004	0,00001,00714,3
5	0,69897,00002,9	100005	0,00002,00214,8
6	0,77815,12103,8	100006	0,00002,00698,9
7	0,84509,80400,1	100007	0,00002,00995,5
8	0,90308,99869,9	100008	0,00002,00821,7
9	0,95424,35944,4	100009	0,00002,00847,4
10			
11	0,04139,16877,6	1000011	0,00000,02542,9
12	0,07918,12160,5	1000012	0,00000,00833,9
13	0,11918,12121,1	1000013	0,00000,01928,3
14	0,16112,80165,8	1000014	0,00001,00717,7
15	0,17609,12190,6	1000015	0,00000,01714,7
16	0,20411,99826,6	1000016	0,00000,02607,6
17	0,23041,89213,8	1000017	0,00000,03040,5
18	0,25129,25071,0	1000018	0,00000,04741,4
19	0,27877,16609,5	1000019	0,00000,09085,3
20			
101	0,00432,13727,8	1000001	0,00000,00434,1
102	0,00860,01717,6	1000002	0,00000,00868,6
103	0,01287,02247,4	1000003	0,00000,01302,9
104	0,01703,03393,0	1000004	0,00000,01737,2
105	0,02118,09999,7	1000005	0,00000,02171,5
106	0,02533,08177,5	1000006	0,00000,02605,8
107	0,02948,03776,9	1000007	0,00000,03040,1
108	0,03364,07754,9	1000008	0,00000,03474,4
109	0,03781,04279,6	1000009	0,00000,03908,6
110			
1001	0,00043,40774,8	10000001	0,00000,00043,4
1002	0,00086,77215,1	10000002	0,00000,00086,9
1003	0,00130,09330,2	10000003	0,00000,00130,3
1004	0,00173,17128,1	10000004	0,00000,00173,7
1005	0,00216,46617,6	10000005	0,00000,00217,1
1006	0,00259,79807,2	10000006	0,00000,00260,6
1007	0,00302,94705,5	10000007	0,00000,00304,0
1008	0,00346,07321,1	10000008	0,00000,00347,4
1009	0,00389,11663,4	10000009	0,00000,00390,9
1100			
10001	0,00004,14272,8	100000001	0,00000,00004,3
10002	0,00008,08170,1	100000002	0,00000,00008,7
10003	0,00013,02328,1	100000003	0,00000,00013,0
10004	0,00017,16830,6	100000004	0,00000,00017,4
10005	0,00021,70029,2	100000005	0,00000,00021,7
10006	0,00026,04985,5	100000006	0,00000,00026,1
10007	0,00030,15997,8	100000007	0,00000,00030,4
10008	0,00034,73966,9	100000008	0,00000,00034,7
10009	0,00039,08935,8	100000009	0,00000,00039,1

15

- I want 12 significant digits
- I have an approximation scheme that provides 14 digits
- or,

$$y = \log(x) \pm 10^{-14}$$

- “Usually” that’s enough to round

$$y = x, \text{xxxxxxxxxxxx}17 \pm 10^{-14}$$

$$y = x, \text{xxxxxxxxxxxx}83 \pm 10^{-14}$$

- **Dilemma** when

$$y = x, \text{xxxxxxxxxxxx}50 \pm 10^{-14}$$


The first table-makers rounded these cases randomly, and recorded them to confound copiers.

Solving the table maker's dilemma

Ziv's onion peeling algorithm

1. Initialisation : $\varepsilon = \varepsilon_1$


Solving the table maker's dilemma



Ziv's onion peeling algorithm

1. Initialisation : $\varepsilon = \varepsilon_1$
2. Compute y such that $f(x) = y \pm \varepsilon$


Solving the table maker's dilemma



Ziv's onion peeling algorithm

1. Initialisation : $\varepsilon = \varepsilon_1$
2. Compute y such that $f(x) = y \pm \varepsilon$
3. Does $y \pm \varepsilon$ contain the middle point between two FP numbers?


Solving the table maker's dilemma



Ziv's onion peeling algorithm

1. Initialisation : $\varepsilon = \varepsilon_1$
2. Compute y such that $f(x) = y \pm \varepsilon$
3. Does $y \pm \varepsilon$ contain the middle point between two FP numbers?
 - If no, return $\text{RN}(y)$


Solving the table maker's dilemma



Ziv's onion peeling algorithm

1. Initialisation : $\varepsilon = \varepsilon_1$
2. Compute y such that $f(x) = y \pm \varepsilon$
3. Does $y \pm \varepsilon$ contain the middle point between two FP numbers?
 - If no, return $\text{RN}(y)$
 - If yes,


Solving the table maker's dilemma



Ziv's onion peeling algorithm

1. Initialisation : $\varepsilon = \varepsilon_1$
2. Compute y such that $f(x) = y \pm \varepsilon$
3. Does $y \pm \varepsilon$ contain the middle point between two FP numbers?
 - If no, return $\text{RN}(y)$
 - If yes, dilemma !


Solving the table maker's dilemma



Ziv's onion peeling algorithm

1. Initialisation : $\varepsilon = \varepsilon_1$
2. Compute y such that $f(x) = y \pm \varepsilon$
3. Does $y \pm \varepsilon$ contain the middle point between two FP numbers?
 - If no, return $\text{RN}(y)$
 - If yes, dilemma ! Reduce ε , and go back to 2


Solving the table maker's dilemma



Ziv's onion peeling algorithm

1. Initialisation : $\varepsilon = \varepsilon_1$
2. Compute y such that $f(x) = y \pm \varepsilon$
3. Does $y \pm \varepsilon$ contain the middle point between two FP numbers?
 - If no, return $\text{RN}(y)$
 - If yes, dilemma ! Reduce ε , and go back to 2

Solving the table maker's dilemma



Ziv's onion peeling algorithm

1. Initialisation : $\varepsilon = \varepsilon_1$
2. Compute y such that $f(x) = y \pm \varepsilon$
3. Does $y \pm \varepsilon$ contain the middle point between two FP numbers?
 - If no, return $\text{RN}(y)$
 - If yes, dilemma ! Reduce ε , and go back to 2

It is a *while* loop... we have to show it terminates, a topic in itself.

Accuracy versus performance

When we know that the loop terminates...

CRLibm : 2-step approximation process

- first step **fast** but accurate to $\bar{\epsilon}_1$
sometimes not accurate enough
- (rarely) second step slower but **always accurate enough**

Accuracy versus performance

When we know that the loop terminates...

CRLibm : 2-step approximation process

- first step **fast** but accurate to $\bar{\epsilon}_1$
sometimes not accurate enough
- (rarely) second step slower but **always accurate enough**

$$T_{\text{avg}} = T_1 + p_2 T_2$$

Accuracy versus performance

When we know that the loop terminates...

CRLibm : 2-step approximation process

- first step **fast** but accurate to $\bar{\epsilon}_1$
sometimes not accurate enough
- (rarely) second step slower but **always accurate enough**

$$T_{\text{avg}} = T_1 + p_2 T_2$$

For each step, we want to prove a **tight** bound $\bar{\epsilon}$ such that

$$\left| \frac{F(x) - f(x)}{f(x)} \right| \leq \bar{\epsilon}$$

Accuracy versus performance

When we know that the loop terminates...

CRLibm : 2-step approximation process

- first step **fast** but accurate to $\bar{\epsilon}_1$
sometimes not accurate enough
- (rarely) second step slower but **always accurate enough**

$$T_{\text{avg}} = T_1 + p_2 T_2$$

For each step, we want to prove a **tight** bound $\bar{\epsilon}$ such that

$$\left| \frac{F(x) - f(x)}{f(x)} \right| \leq \bar{\epsilon}$$

- Overestimating $\bar{\epsilon}_2$ degrades T_2 ! (common wisdom)

Accuracy versus performance

When we know that the loop terminates...

CRLibm : 2-step approximation process


- first step **fast** but accurate to $\bar{\varepsilon}_1$ sometimes not accurate enough
- (rarely) second step slower but **always accurate enough**

$$T_{\text{avg}} = T_1 + p_2 T_2$$

For each step, we want to prove a **tight** bound $\bar{\varepsilon}$ such that

$$\left| \frac{F(x) - f(x)}{f(x)} \right| \leq \bar{\varepsilon}$$

- Overestimating $\bar{\varepsilon}_2$ degrades T_2 ! (common wisdom)
- Overestimating $\bar{\varepsilon}_1$ degrades p_2 !



First function development in Arénaire

First correctly rounded elementary function in CRLibm

- exp by David Defour
- worst-case time $T_2 \approx 10,000$ cycles
- complex, hand-written proof

First function development in Arénaire

First correctly rounded elementary function in CRLibm

- exp by David Defour
- worst-case time $T_2 \approx 10,000$ cycles
- complex, hand-written proof
- duration : a Ph.D. thesis (2002)

First function development in Arénaire

First correctly rounded elementary function in CRLibm

- exp by David Defour
- worst-case time $T_2 \approx 10,000$ cycles
- complex, hand-written proof
- duration : a Ph.D. thesis (2002)

Conclusion was :

- performance and memory consumption of CR elem function is OK

First function development in Arénaire

First correctly rounded elementary function in CRLibm

- exp by David Defour
- worst-case time $T_2 \approx 10,000$ cycles
- complex, hand-written proof
- duration : a Ph.D. thesis (2002)

Conclusion was :

- performance and memory consumption of CR elem function is OK
- problem now is : performance and coffee consumption of the programmer

Latest function developments in Arénaire

C. Lauter at the end of his PhD,

- development time for `sinpi`, `cospi`, `tanpi` :

C. Lauter at the end of his PhD,

- development time for sinpi, cospi, tanpi : 2 days
- worst-case time $T_2 \approx 1,000$ cycles

Latest function developments in Arénaire

C. Lauter at the end of his PhD,

- development time for sinpi, cospi, tanpi : 2 days
- worst-case time $T_2 \approx 1,000$ cycles

(but as a result of three more PhDs)

Summary of the progress made

$$T_{\text{avg}} = T_1 + p_2 T_2$$

- Reduction of T_1 by learning from Intel
- Reduction of p_2 by automating the computation of tight $\bar{\epsilon}_1$
(p_2 is proportional to $\bar{\epsilon}_1$)
- Reduction of T_2 by computing just right
- Reduction of coffee consumption by automating the whole thing

Summary of the progress made

$$T_{\text{avg}} = T_1 + p_2 T_2$$

- Reduction of T_1 by learning from Intel
- Reduction of p_2 by automating the computation of tight $\bar{\epsilon}_1$
(p_2 is proportional to $\bar{\epsilon}_1$)
- Reduction of T_2 by computing just right
- Reduction of coffee consumption by automating the whole thing

The MetaLibm vision

Automate libm expertise so that a new, correct libm can be written for a new processor/context in minutes instead of months.

Formal proof of floating-point code for the masses

Introduction : performance versus accuracy

Elementary function evaluation

Formal proof of floating-point code for the masses

Other tools toward MetaLibm

Conclusion


```

1  yh2 = yh*yh;
2  ts = yh2 * (s3.d + yh2*(s5.d + yh2*s7.d));
3  Add12(*psh,*psl, yh, y1+ts*yh);

```

Upon entering DoSinZero, we have in $y_h + y_l$ an approximation to the ideal reduced value $\hat{y} = x - k\frac{\pi}{256}$ with a relative accuracy $\varepsilon_{\text{argred}}$:

$$y_h + y_l = \left(x - k\frac{\pi}{256}\right)(1 + \varepsilon_{\text{argred}}) = \hat{y}(1 + \varepsilon_{\text{argred}}) \quad (1)$$

with, depending on the quadrant, $\sin(\hat{y}) = \pm \sin(x)$ or $\sin(\hat{y}) = \pm \cos(x)$ and similarly for $\cos(\hat{y})$. This just means that \hat{y} is the ideal, errorless reduced value.

In the following we will assume we are in the case $\sin(\hat{y}) = \sin(x)$, (the proof is identical in the other cases), therefore the relative error that we need to compute is

$$\varepsilon_{\text{sinkzero}} = \frac{(*\text{psh} + *\text{psl})}{\sin(x)} - 1 = \frac{(*\text{psh} + *\text{psl})}{\sin(\hat{y})} - 1 \quad (2)$$

One may remark that we almost have the same code as we have for computing the sine of a small argument (without range reduction). The difference is that we have as input a double-double $y_h + y_l$, which is itself an inexact term.

At Line 4, the error of neglecting y_l and the rounding error in the multiplication each amount to half an ulp :

$$y_h2 = y_h^2(1 + \varepsilon_{-53}), \text{ with } y_h = (y_h + y_l)(1 + \varepsilon_{-53}) = \hat{y}(1 + \varepsilon_{\text{argred}})(1 + \varepsilon_{-53})$$

Therefore

$$yh2 = \hat{y}^2(1 + \varepsilon_{yh2}) \quad (3)$$

with

$$\bar{\varepsilon}_{yh2} = (1 + \bar{\varepsilon}_{argred})^2(1 + \bar{\varepsilon}_{-53})^3 - 1 \quad (4)$$

Line 5 is a standard Horner evaluation. Its approximation error is defined by :

$$P_{ts}(\hat{y}) = \frac{\sin(\hat{y}) - \hat{y}}{\hat{y}}(1 + \varepsilon_{\text{approx}ts})$$

This error is computed in **Maple** as previously, only the interval changes :

$$\bar{\varepsilon}_{\text{approx}ts} = \left\| \frac{xP_{ts}(x)}{\sin(x) - x} - 1 \right\|_{\infty}$$

We also compute $\bar{\varepsilon}_{\text{hornert}ts}$, the bound on the relative error due to rounding in the Horner evaluation thanks to the `compute_horner_rounding_error` procedure. This time, this procedure takes into account the relative error carried by `yh2`, which is $\bar{\varepsilon}_{yh2}$ computed above. We thus get the total relative error on `ts` :

$$ts = P_{ts}(\hat{y})(1 + \varepsilon_{\text{hornert}ts}) = \frac{\sin(\hat{y}) - \hat{y}}{\hat{y}}(1 + \varepsilon_{\text{approx}ts})(1 + \varepsilon_{\text{hornert}ts}) \quad (5)$$

The final Add12 is exact. Therefore the overall relative error is :

$$\begin{aligned}
 \varepsilon_{\text{sinkzero}} &= \frac{((y_h \otimes t_s) \oplus y_l) + y_h}{\sin(\hat{y})} - 1 \\
 &= \frac{(y_h \otimes t_s + y_l)(1 + \varepsilon_{-53}) + y_h}{\sin(\hat{y})} - 1 \\
 &= \frac{y_h \otimes t_s + y_l + y_h + (y_h \otimes t_s + y_l) \cdot \varepsilon_{-53}}{\sin(\hat{y})} - 1
 \end{aligned}$$

Let us define for now

$$\delta_{\text{addsin}} = (y_h \otimes t_s + y_l) \cdot \varepsilon_{-53} \quad (6)$$

Then we have

$$\varepsilon_{\text{sinkzero}} = \frac{(y_h + y_l)t_s(1 + \varepsilon_{-53})^2 + y_l + y_h + \delta_{\text{addsin}}}{\sin(\hat{y})} - 1$$

Using (1) and (5) we get :

$$\varepsilon_{\text{sinkzero}} = \frac{\hat{y}(1 + \varepsilon_{\text{argred}}) \times \frac{\sin(\hat{y}) - \hat{y}}{\hat{y}} (1 + \varepsilon_{\text{approx}t_s})(1 + \varepsilon_{\text{hornert}t_s})(1 + \varepsilon_{-53})^2 + y_l + y_h + \delta_{\text{addsin}}}{\sin(\hat{y})} - 1$$

To lighten notations, let us define

$$\varepsilon_{\text{sin1}} = (1 + \varepsilon_{\text{approx}t_s})(1 + \varepsilon_{\text{hornert}t_s})(1 + \varepsilon_{-53})^2 - 1 \quad (7)$$

We get

$$\begin{aligned}\varepsilon_{\text{sinkzero}} &= \frac{(\sin(\hat{y}) - \hat{y})(1 + \varepsilon_{\text{sin1}}) + \hat{y}(1 + \varepsilon_{\text{argred}}) + \delta_{\text{addsin}} - \sin(\hat{y})}{\sin(\hat{y})} \\ &= \frac{(\sin(\hat{y}) - \hat{y}) \cdot \varepsilon_{\text{sin1}} + \hat{y} \cdot \varepsilon_{\text{argred}} + \delta_{\text{addsin}}}{\sin(\hat{y})}\end{aligned}$$

Using the following bound :

$$|\delta_{\text{addsin}}| = |(\mathbf{yh} \otimes \mathbf{ts} + \mathbf{y1}) \cdot \varepsilon_{-53}| < 2^{-53} \times |y|^3/3 \quad (8)$$

we may compute the value of $\bar{\varepsilon}_{\text{sinkzero}}$ as an infinite norm under **Maple**. We get an error smaller than 2^{-67} .

4 pages for 3 lines of code...

Two years of experience showed that nobody (including myself) should trust such a proof

4 pages for 3 lines of code...

Two years of experience showed that nobody (including myself) should trust such a proof (and that nobody reads it anyway).

4 pages for 3 lines of code...

Two years of experience showed that nobody (including myself) should trust such a proof (and that nobody reads it anyway).

We wish we had an **automatic** tool that

- takes a set of C files,
- parses them,
- and outputs “The overall error of the computation is ...”.

4 pages for 3 lines of code...

Two years of experience showed that nobody (including myself) should trust such a proof (and that nobody reads it anyway).

We wish we had an **automatic** tool that

- takes a set of C files,
- parses them,
- and outputs “The overall error of the computation is ...”.

It's hopeless, of course :

4 pages for 3 lines of code...

Two years of experience showed that nobody (including myself) should trust such a proof (and that nobody reads it anyway).

We wish we had an **automatic** tool that

- takes a set of C files,
- parses them,
- and outputs “The overall error of the computation is ...”.

It's hopeless, of course :

- Where, in your code, can you read **what it is supposed to compute** ?

4 pages for 3 lines of code...

Two years of experience showed that nobody (including myself) should trust such a proof (and that nobody reads it anyway).

We wish we had an **automatic** tool that

- takes a set of C files,
- parses them,
- and outputs “The overall error of the computation is ...”.

It's hopeless, of course :

- Where, in your code, can you read **what it is supposed to compute** ?
- Most of the knowledge used to build the code is **not** in the code

Trusted error computation means : formal proof

but... automatic proof assistants are not there yet

- Research on formal proofs for arithmetic
 - John Harrison at Intel (HOL light)
 - Marc Daumas and Sylvie Boldo in the Arénaire project (Coq, PVS)
 - And many others...

Trusted error computation means : formal proof

but... automatic proof assistants are not there yet

- Research on formal proofs for arithmetic
 - John Harrison at Intel (HOL light)
 - Marc Daumas and Sylvie Boldo in the Arénaire project (Coq, PVS)
 - And many others...
- Proving Sterbenz Lemma (one operation) is worth a full paper.

Trusted error computation means : formal proof

but... automatic proof assistants are not there yet

- Research on formal proofs for arithmetic
 - John Harrison at Intel (HOL light)
 - Marc Daumas and Sylvie Boldo in the Arénaire project (Coq, PVS)
 - And many others...
- Proving Sterbenz Lemma (one operation) is worth a full paper.
- Here is the typical `crlibm` code for which I want the relative error :

```
1  yh2 = yh*yh ;
2  ts = yh2 * (s3 + yh2*(s5 + yh2*s7));
3  tc = yh2 * (c2 + yh2*(c4 + yh2*c6 ));
4  Mul12(&cahyh_h,&cahyh_l, cah, yh);
5  Add12(thi, tlo, sah,cahyh_h);
6  tlo = tc*sah+(ts*cahyh_h+(sal+(tlo+(cahyh_l+(cal*yh +
   cah*y1)))))) ;
7  Add12(*psh,*psl, thi, tlo);
```

Trusted error computation means : formal proof

but... automatic proof assistants are not there yet

- Research on formal proofs for arithmetic
 - John Harrison at Intel (HOL light)
 - Marc Daumas and Sylvie Boldo in the Arénaire project (Coq, PVS)
 - And many others...
- Proving Sterbenz Lemma (one operation) is worth a full paper.
- Here is the typical `crlibm` code for which I want the relative error :

```
1  yh2 = yh*yh ;
2  ts = yh2 * (s3 + yh2*(s5 + yh2*s7));
3  tc = yh2 * (c2 + yh2*(c4 + yh2*c6 ));
4  Mul12(&cahyh_h,&cahyh_l, cah, yh);
5  Add12(thi, tlo, sah, cahyh_h);
6  tlo = tc*sah+(ts*cahyh_h+(sal+(tlo+(cahyh_l+(cal*yh +
   cah*y1)))))) ;
7  Add12(*psh,*psl, thi, tlo);
```

... and it changes all the time as we optimize it.

Rounding errors piled over approximations

```
y2 = y * y;  
ts = y2 * (s3 + y2*(s5 + y2*s7));  
r = y + y*ts
```

- This polynomial is an approximation to $\sin(y)$

Rounding errors piled over approximations

```
y2 = y * y;  
ts = y2 * (s3 + y2*(s5 + y2*s7));  
r = y + y*ts
```

- This polynomial is an approximation to $\sin(y)$
- Oops, I wrote its coefficients in decimal !

Rounding errors piled over approximations

```
y2 = y * y;  
ts = y2 * (s3 + y2*(s5 + y2*s7));  
r = y + y*ts
```

- This polynomial is an approximation to $\sin(y)$
- Oops, I wrote its coefficients in decimal!
- y is not the ideal reduced argument Y (such that $x = Y + k\frac{\pi}{256}$)

Rounding errors piled over approximations

```
y2 = y * y;  
ts = y2 * (s3 + y2*(s5 + y2*s7));  
r = y + y*ts
```

- This polynomial is an approximation to $\sin(y)$
- Oops, I wrote its coefficients in decimal!
- y is not the ideal reduced argument Y (such that $x = Y + k\frac{\pi}{256}$)
- We have a rounding error in computing y^2

Rounding errors piled over approximations

```
y2 = y * y;  
ts = y2 * (s3 + y2*(s5 + y2*s7));  
r = y + y*ts
```

- This polynomial is an approximation to $\sin(y)$
- Oops, I wrote its coefficients in decimal!
- y is not the ideal reduced argument Y (such that $x = Y + k\frac{\pi}{256}$)
- We have a rounding error in computing y^2
- $y2$ already stacks two errors. We evaluate ts out of it

Rounding errors piled over approximations

```
y2 = y * y;  
ts = y2 * (s3 + y2*(s5 + y2*s7));  
r = y + y*ts
```

- This polynomial is an approximation to $\sin(y)$
- Oops, I wrote its coefficients in decimal!
- y is not the ideal reduced argument Y (such that $x = Y + k\frac{\pi}{256}$)
- We have a rounding error in computing y^2
- $y2$ already stacks two errors. We evaluate ts out of it
- There is a rounding error hidden in each operation.

Rounding errors piled over approximations

```
y2 = y * y;  
ts = y2 * (s3 + y2*(s5 + y2*s7));  
r = y + y*ts
```

- This polynomial is an approximation to $\sin(y)$
- Oops, I wrote its coefficients in decimal!
- y is not the ideal reduced argument Y (such that $x = Y + k\frac{\pi}{256}$)
- We have a rounding error in computing y^2
- $y2$ already stacks two errors. We evaluate ts out of it
- There is a rounding error hidden in each operation.

How many correct bits at the end?

My programmer's genius is hidden in this code

$y*(1+ts)$ is a bit less accurate than $y + y*ts$ in floating-point

That's because $|t| < 2^{-14}$ because $|y| < 2^{-7}$ (not in the code)

$$\begin{array}{l} \boxed{1} \\ + \quad \boxed{t} \\ = \quad \boxed{1+t} \end{array}$$

$$\begin{array}{l} \boxed{y} \\ + \quad \boxed{y*t} \\ = \quad \boxed{y+y*t} \end{array}$$

Gappa

Written by Guillaume Melquiond, Gappa is a tool that

- takes an input that closely matches your C file,

Gappa

Written by Guillaume Melquiond, Gappa is a tool that

- takes an input that closely matches your C file,
- forces you to express **what this code is supposed to compute**

Gappa

Written by Guillaume Melquiond, Gappa is a tool that

- takes an input that closely matches your C file,
- forces you to express **what this code is supposed to compute**
- ... and some numerical property to prove (expressed in terms of intervals)

Gappa

Written by Guillaume Melquiond, Gappa is a tool that

- takes an input that closely matches your C file,
- forces you to express **what this code is supposed to compute**
- ... and some numerical property to prove (expressed in terms of intervals)
- and eventually outputs a proof of this property suitable for checking by Coq or HOL Light

Try it, it's free software

Should I present interval arithmetic?

Using a machine's finite precision, manipulate reals safely

Should I present interval arithmetic?

Using a machine's finite precision, manipulate reals safely

- represent a **real** x in a machine as an interval $[x_l, x_r]$
guaranteed to enclose it
 - x_l and x_r are finitely representable numbers (e.g. floating-point)
 - Example : π represented by $[3.14, 3.15]$

Should I present interval arithmetic?

Using a machine's finite precision, manipulate reals safely

- represent a **real** x in a machine as an interval $[x_l, x_r]$
guaranteed to enclose it
 - x_l and x_r are finitely representable numbers (e.g. floating-point)
 - Example : π represented by $[3.14, 3.15]$
- Operation \oplus on the reals \rightarrow its interval counterpart

Guarantees based on the **inclusion property**

$I_x \oplus I_y$ must be an interval I_z such that

$$\forall x \in I_x, \forall y \in I_y, \quad x \oplus y \in I_z$$

Should I present interval arithmetic?

Using a machine's finite precision, manipulate reals safely

- represent a **real** x in a machine as an interval $[x_l, x_r]$
guaranteed to enclose it
 - x_l and x_r are finitely representable numbers (e.g. floating-point)
 - Example : π represented by $[3.14, 3.15]$
- Operation \oplus on the reals \rightarrow its interval counterpart

Guarantees based on the **inclusion property**

$I_x \oplus I_y$ must be an interval I_z such that

$$\forall x \in I_x, \forall y \in I_y, \quad x \oplus y \in I_z$$

- Example : interval addition using floating-point arithmetic

$$[a, b] + [c, d] \quad \text{is} \quad [\text{RoundDown}(a + c), \text{RoundUp}(b + d)]$$

- (multiplication, division similar but more complex)


```
$ gappa < tutorial1.gappa
Results for Y in [-0.00615, 0.00615] and y - Y in [-2.53e-23, 2.53e-23]
r - SinY in [-2-60.9998, 2-60.9998]
Warning: some enclosures were not satisfied.
Missing (r - SinY) / SinY
$
```

- A tight bound on the absolute error
- No bound for the relative error
 - of course, I have to prove that SinY cannot come close to zero
 - that's formal proof for you

We should now try gappa -Bcoq

How does Gappa work ?

- Gappa tries to associate an interval with each expression.
- Interval arithmetic is used to combine these intervals, until the goal is reached.

How does Gappa work ?

- Gappa tries to associate an interval with each expression.
- Interval arithmetic is used to combine these intervals, until the goal is reached.
- Naively, it would lead to interval bloat. Here for instance
 - $r \approx \text{Sin}Y \in [-2^{-7}, 2^{-7}]$
 - so $r - \text{Sin}Y \in [-2^{-6}, 2^{-6}]$ using naive IA.

How does Gappa work ?

- Gappa tries to associate an interval with each expression.
- Interval arithmetic is used to combine these intervals, until the goal is reached.
- Naively, it would lead to interval bloat. Here for instance
 - $r \approx \text{SinY} \in [-2^{-7}, 2^{-7}]$
 - so $r - \text{SinY} \in [-2^{-6}, 2^{-6}]$ using naive IA.
- Gappa uses **rewriting** of expressions

As `r = float64ne(E);`

try and use the rule

`float64ne(E) - SinY -> (float64ne(E) - E) + (E - SinY);`

(hopefully now the sum of two smaller intervals)

How does Gappa work ?

- Gappa tries to associate an interval with each expression.
- Interval arithmetic is used to combine these intervals, until the goal is reached.
- Naively, it would lead to interval bloat. Here for instance
 - $r \approx \text{SinY} \in [-2^{-7}, 2^{-7}]$
 - so $r - \text{SinY} \in [-2^{-6}, 2^{-6}]$ using naive IA.
- Gappa uses **rewriting** of expressions
As `r = float64ne(E);`
try and use the rule
`float64ne(E) - SinY -> (float64ne(E) - E) + (E - SinY);`
(hopefully now the sum of two smaller intervals)
- Add **user-defined** rewriting rules when Gappa is stuck
 - That's how you explain your floating-point tricks to the tool

How does Gappa work ?

- Gappa tries to associate an interval with each expression.
- Interval arithmetic is used to combine these intervals, until the goal is reached.
- Naively, it would lead to interval bloat. Here for instance
 - $r \approx \text{Sin}Y \in [-2^{-7}, 2^{-7}]$
 - so $r - \text{Sin}Y \in [-2^{-6}, 2^{-6}]$ using naive IA.
- Gappa uses **rewriting** of expressions
As `r = float64ne(E);`
try and use the rule
`float64ne(E) - SinY -> (float64ne(E) - E) + (E - SinY);`
(hopefully now the sum of two smaller intervals)
- Add **user-defined** rewriting rules when Gappa is stuck
 - That's how you explain your floating-point tricks to the tool
- Internally, construction of a proof graph
 - Branches are cut when a shorter path or a better bound are found.

How does Gappa work ?

- Gappa tries to associate an interval with each expression.
- Interval arithmetic is used to combine these intervals, until the goal is reached.
- Naively, it would lead to interval bloat. Here for instance
 - $r \approx \text{Sin}Y \in [-2^{-7}, 2^{-7}]$
 - so $r - \text{Sin}Y \in [-2^{-6}, 2^{-6}]$ using naive IA.
- Gappa uses **rewriting** of expressions
As `r = float64ne(E);`
try and use the rule
 $\text{float64ne}(E) - \text{Sin}Y \rightarrow (\text{float64ne}(E) - E) + (E - \text{Sin}Y);$
(hopefully now the sum of two smaller intervals)
- Add **user-defined** rewriting rules when Gappa is stuck
 - That's how you explain your floating-point tricks to the tool
- Internally, construction of a proof graph
 - Branches are cut when a shorter path or a better bound are found.
 - The final graph will be used to generate the formal proof.

Gappa's theorem library

- Predefined set of rewriting rules :
 - `float64ne(a)- b ->(float64ne(a)- a)+ (a - b);`
 - ...
- Support library of theorems (**with their Coq proofs**) :
 - Theorems giving the errors when rounding
 - ▶ `a in [...] ->(float64ne(a)-a)/a in [...]`
Note how this takes care of dangerous cases (subnormal numbers, over/underflows...)

Gappa's theorem library

- Predefined set of rewriting rules :
 - `float64ne(a)- b ->(float64ne(a)- a)+ (a - b);`
 - ...
- Support library of theorems (**with their Coq proofs**) :
 - Theorems giving the errors when rounding
 - ▶ `a in [...] ->(float64ne(a)-a)/a in [...]`
Note how this takes care of dangerous cases (subnormal numbers, over/underflows...)
 - Classical theorems like Sterbenz Lemma
 - ...

Gappa's theorem library

- Predefined set of rewriting rules :
 - `float64ne(a)- b ->(float64ne(a)- a)+ (a - b);`
 - ...
- Support library of theorems (**with their Coq proofs**) :
 - Theorems giving the errors when rounding
 - ▶ `a in [...] ->(float64ne(a)-a)/a in [...]`
Note how this takes care of dangerous cases (subnormal numbers, over/underflows...)
 - Classical theorems like Sterbenz Lemma
 - ...

To obtain a good relative error, Gappa will demand to prove that y may not be subnormal...

$y + y*ts$ is a bit more accurate than $y*(1+ts)$

```
14 r1 float<ieee_64,ne>= y*(1+ts);
15 r2 float<ieee_64,ne>= y+y*ts;
16
17 yts float<ieee_64,ne>= y*ts; # for lighter hints
18
19 #----- Mathematical definition of what we are approximating -----
20 # (The same expression as in the code, but without rounding errors)
21 Y2 = y*y;
22 Ts = Y2 * (s3 + Y2*(s5 + Y2*s7));
23 Poly = y*(1+Ts);
24 #----- The theorem to prove -----
25 {
26 # Hypotheses (numerical values computed by Sollya)
27 y in [1b-200, 6.15e-3] # left: Kahan/Douglas algorithm. Right: Pi/512, rounded up
28 ->
29 r1-/Poly in ? # relative error
30 /\
31 r2-/Poly in ? # relative error
32 }
33
34 #----- Loads of rewriting hints needed for r2 -----
35 y+yts -> y* ( (1+ts) + ts*((yts-y*ts) / (y*ts)) {y*ts <> 0};
36
37 (r2-Poly)/Poly -> ((r2 - (y+yts))/(y+yts) + 1) * ( ((y+yts)/y) / (1+Ts)) -1 {1+Ts
  <>0};
38
39 (y+yts)/y ->
40 # (y+y*ts-y*ts+yts) /y;
41 # 1+ts + (yts-y*ts)/y;
42 1+ts + ts*( (yts-y*ts)/(y*ts) ) {y*ts <> 0};
43
44 ((y+yts)/y) / (1+Ts) -> (1+ts)/(1+Ts) + ts*( (yts-y*ts)/(y*ts) )/(1+Ts) {1+Ts<>0};
45
46 (1+ts)/(1+Ts) -> 1 + (Ts*((ts-Ts)/Ts))/(1+Ts) {1+Ts<>0};
```

```
$ gappa < tutorial2.gappa
```

```
Results for y in [7.88861e-31, 0.00615]:
```

```
(r1 - Poly) / Poly in [-2(-52.415), 2(-52.415)]
```

```
(r2 - Poly) / Poly in [-2(-52.9777), 2(-52.9339)]
```

```
$
```

Conclusion on Gappa

- I probably failed to convey this, but...

Gappa is surprisingly easy to use.

(if you didn't understand my Gappa proof, you just don't understand my C code)

- if you don't know where it is stuck, ask it (by adding goals)
- then add rewriting rules to help it

Conclusion on Gappa

- I probably failed to convey this, but...
Gappa is surprisingly easy to use.
(if you didn't understand my Gappa proof, you just don't understand my C code)
 - if you don't know where it is stuck, ask it (by adding goals)
 - then add rewriting rules to help it
- It is built upon very solid theoretical foundations

Conclusion on Gappa

- I probably failed to convey this, but...
Gappa is surprisingly easy to use.
(if you didn't understand my Gappa proof, you just don't understand my C code)
 - if you don't know where it is stuck, ask it (by adding goals)
 - then add rewriting rules to help it
- It is built upon very solid theoretical foundations
- What we have now is generators of code + Gappa proof
 - The same RR work for large classes of generated codes.

Conclusion on Gappa

- I probably failed to convey this, but...
Gappa is surprisingly easy to use.
(if you didn't understand my Gappa proof, you just don't understand my C code)
 - if you don't know where it is stuck, ask it (by adding goals)
 - then add rewriting rules to help it
- It is built upon very solid theoretical foundations
- What we have now is generators of code + Gappa proof
 - The same RR work for large classes of generated codes.
- Also support for arbitrary-precision fixed-point.

Other tools toward MetaLibm

Introduction : performance versus accuracy

Elementary function evaluation

Formal proof of floating-point code for the masses

Other tools toward MetaLibm

Conclusion

Multiple Precision Floating-point correctly Rounded

MPFI : interval arithmetic on top of MPFR

The Swiss Army Knife of the libm developer (Lauter, Chevillard, Joldes)

- multiple-precision, last-bit accurate evaluation of arbitrary expressions
 - apologizes each time it rounds something
 - a demo ?

The Swiss Army Knife of the libm developer (Lauter, Chevillard, Joldes)

- multiple-precision, last-bit accurate evaluation of arbitrary expressions
 - apologizes each time it rounds something
 - a demo?
- guaranteed infinite norm $\|f(x)\|_\infty$ even in degenerate cases
 - $\|f(x) - P(x)\|_\infty$ is a degenerate case...
 - Gappa bounds the **rounding** errors, this bounds the **approximation** error

The Swiss Army Knife of the libm developer (Lauter, Chevillard, Joldes)

- multiple-precision, last-bit accurate evaluation of arbitrary expressions
 - apologizes each time it rounds something
 - a demo?
- guaranteed infinite norm $\|f(x)\|_\infty$ even in degenerate cases
 - $\|f(x) - P(x)\|_\infty$ is a degenerate case...
 - Gappa bounds the **rounding** errors, this bounds the **approximation** error
- Machine-efficient polynomial approximation

The Patriot bug

In 1991, a Patriot missile failed to intercept a Scud, and 28 people were killed.

- The code worked with time increments of 0.1 s.
- But 0.1 is not representable in binary.
- In the 24-bit format used, the number stored was 0.099999904632568359375
- The error was 0.0000000953.
- After 100 hours = 360,000 seconds, time is wrong by 0.34s.
- In 0.34s, a Scud moves 500m

(similar problems have been discovered in civilian air traffic control systems, after near-miss incidents)

Test : which of the following increments should you use ?

10 5 3 1 0.5 0.25 0.2 0.125 0.1

Machine-efficient polynomial approximation

- Remez' minimax algorithm finds the best polynomial approximation **over the reals**
- But we need polynomials with **machine** coefficients
 - float, double, fixed-point, ...
- Rounding Remez coefficients does **not** provide the best polynomial among polynomial with machine coefficients.

Machine-efficient polynomial approximation

- Remez' minimax algorithm finds the best polynomial approximation **over the reals**
- But we need polynomials with **machine** coefficients
 - float, double, fixed-point, ...
- Rounding Remez coefficients does **not** provide the best polynomial among polynomial with machine coefficients.
- Sollya does (almost).
 - this saves a few bits of accuracy
 - especially relevant for small precisions (FPGAs)
 - that's how we get our polynomials

Machine-efficient polynomial approximation

- Remez' minimax algorithm finds the best polynomial approximation **over the reals**
- But we need polynomials with **machine** coefficients
 - float, double, fixed-point, ...
- Rounding Remez coefficients does **not** provide the best polynomial among polynomial with machine coefficients.
- Sollya does (almost).
 - this saves a few bits of accuracy
 - especially relevant for small precisions (FPGAs)
 - that's how we get our polynomials

Nice number theory behind.


Classical doubled FP

- Store a $2p$ -digit number y as two p -digit numbers y_h and y_l
- $y = y_h + y_l$
- $\text{exponent}(y_l) \leq \text{exponent}(y_h) - p$



Example


Decimal format, $p = 3$ digits,
3.14159 stored as $y_h = 3.14$, $y_l = 1.59e - 3$



A lot of litterature to compute efficiently on doubled-FP.

Never compute more accurately than you need

Polynomial evaluation $P(y)$ when $y < 2^{-k}$



For CRLibm

- doubled-binary64 (106 bits) is not enough,
- but triple-binary64 (159 bits) is overkill

An example of overlapping triple-double arithmetic

Add233 : add a double-FP to a triple-FP

Require: $a_h + a_\ell$ is a double-double number and $b_h + b_m + b_\ell$ is a triple-double number such that $|b_h| \leq 2^{-2} \cdot |a_h|$, $|a_\ell| \leq 2^{-53} \cdot |a_h|$, $|b_m| \leq 2^{-\beta_o} \cdot |b_h|$, $|b_\ell| \leq 2^{-\beta_u} \cdot |b_m|$.

Ensure: $r_h + r_m + r_\ell$ is a triple-double number approximating $a_h + a_\ell + b_h + b_m + b_\ell$ with a relative error given by the Theorem on next slide.

$$(r_h, t_1) \leftarrow \text{Fast2Sum}(a_h, b_h)$$

$$(t_2, t_3) \leftarrow \text{Fast2Sum}(a_\ell, b_m)$$

$$(t_4, t_5) \leftarrow \text{Fast2Sum}(t_1, t_2)$$

$$t_6 \leftarrow \text{RN}(t_3 + b_\ell)$$

$$t_7 \leftarrow \text{RN}(t_6 + t_5)$$

$$(r_m, r_\ell) \leftarrow \text{Fast2Sum}(t_4, t_7)$$

β_o and β_u measure the possible overlap of the significands of the inputs.

The associated theorem

Theorem (Result overlap and relative error of Add233)

Under the conditions on previous slide, the values r_h , r_m , and r_ℓ returned by the algorithm satisfy

$$r_h + r_m + r_\ell = ((a_h + a_\ell) + (b_h + b_m + b_\ell)) \cdot (1 + \varepsilon),$$

where ε is bounded by

$$|\varepsilon| \leq 2^{-\beta_o - \beta_u - 52} + 2^{-\beta_o - 104} + 2^{-153}.$$

The values r_m and r_ℓ will not overlap at all, and the overlap of r_h and r_m will be bounded by

$$|r_m| \leq 2^{-\gamma} \cdot |r_h|$$

with

$$\gamma \geq \min(45, \beta_o - 4, \beta_o + \beta_u - 2).$$

30 more, but who will read the proofs?

- See `crlibm` source and documentation for the operators themselves.
- Manipulating these theorems by hand is painful : Lauter's `metalibm` assembles such operators automatically for polynomial evaluation.

Code generation for polynomial evaluation

- explores different parallelizations of a polynomial on a VLIW processor
- generates **code** and **Gappa proof of the evaluation error**

Code generation for polynomial evaluation

- explores different parallelizations of a polynomial on a VLIW processor
- generates **code** and **Gappa proof of the evaluation error**

Used to generate the code for the division and square root of FLIP, a Floating-Point Library for Integer Processors (collaboration with ST Microelectronics)

Conclusion

Introduction : performance versus accuracy

Elementary function evaluation

Formal proof of floating-point code for the masses

Other tools toward MetaLibm

Conclusion

Main messages

- Are you able to express what your code is supposed to compute?

Main messages

- Are you able to express what your code is supposed to compute?
If yes, we can help you sort out the gory floating-point issues.

Main messages

- Are you able to express what your code is supposed to compute?
If yes, we can help you sort out the gory floating-point issues.

- If you're computing accurately enough, you're probably computing too accurately.

The Arénaire Touch




Library

Program

---> developed using

—> links against




All these developments are free software.

More automation means more optimization

- $\log(1 + x)$
- Two parameters
 - k from 1 to 13, defines table size
 - target accuracy, between 20 and 120 bits
- 1203 implementations, **all formally checked**


z axis : timings in arbitrary units



My other research project

Computing just right for FPGAs

- Finer granularity : never compute 1 bit that you don't need
- More qualitative freedom : build the operators you need
 - A squarer, a multiplier by $\ln(2)$, a divider by 3...
- Compute more efficiently ?



<http://flopoco.gforge.inria.fr/>

Thank you for your attention