

Worst Cases for the Exponential Function in the IEEE 754r decimal64 Format

Vincent Lefèvre and Damien Stehlé and Paul Zimmermann

LORIA/INRIA Lorraine, Technopôle de Nancy-Brabois,
615 rue du jardin botanique, F-54602 Villers-lès-Nancy Cedex, France

Vincent.Lefevre@inria.fr

stehle@maths.usyd.edu.au

Paul.Zimmermann@loria.fr

1 Introduction

Most computers nowadays support the IEEE 754-1985 standard for binary floating-point arithmetic [1], which requires that all four arithmetic operations (+, −, ×, ÷) and the square root are *correctly rounded*. However radix 10 is more suited to some applications, such as financial and commercial ones, and there have been propositions to normalize it as well and also design hardware implementations. The IEEE 854-1987 standard for radix-independent floating-point arithmetic [2] has been a first step in this direction, but this standard just gives some constraints on the value sets and is not even specific to radix 10. The article [3] describes a first specification of a decimal floating-point arithmetic; it has been improved and the specification included in the current working draft of the revision of the IEEE 754 standard (754r) is described in [4].

One also seeks to extend the IEEE 754 standard to elementary functions, such as the exponential, logarithm and trigonometric functions, by requiring correct rounding on these functions too. Unfortunately fulfilling this requirement is much more complicated than with the basic operations. Indeed, while efficient algorithms to guarantee the correct rounding are known for these basic operations, the only known way to evaluate $f(x)$, where f is an elementary function and x is a machine number¹, is to compute an approximation to $f(x)$ without any useful knowledge except an error bound; and the exact result $f(x)$ may be very close to a machine number or to the middle of two consecutive machine numbers (which are the discontinuity points of the rounding functions), in which case correct rounding can be guaranteed only if the error on the approximation is small enough. This problem is known as the *Table Maker's Dilemma* (TMD). Some cases can be decided easily, but the only known way to obtain a bound on the acceptable error for any input value is to perform an exhaustive search (with a 64-bit format, as considered below, there are at most 2^{64} possible input values). The arguments x for which the values $f(x)$ are the hardest to round are called *worst cases*.

¹ A number that is exactly representable in the floating-point system.

Systematic work on the TMD in radix 2 was first done by Lefèvre and Muller [5], who published worst cases for many elementary functions in double precision, over the full IEEE 754 range for some functions. And correct rounding requirements for some functions in some domains have been added to the 754r working draft. The present authors improved algorithms to deal with higher precisions [6], and in the present paper, the practical feasibility of the method for decimal formats is demonstrated. Indeed the worst cases depend on the representation (radix and precision) and the mathematical function.

Section 2 describes the decimal formats, how worst cases are expressed and briefly recalls the algorithms (in the decimal context) to search for these worst cases. Section 3 gives the example of the exponential function in the 64-bit decimal format.

2 The Table Maker’s Dilemma in Decimal

In this section, the decimal formats are described in Section 2.1, then the general form of worst cases is given, along with a few illustrating examples (Section 2.2). Finally, the algorithms to search for these worst cases are briefly recalled and applied to radix 10 (Section 2.3).

2.1 The Decimal Formats

As specified by the IEEE 854 standard [2], a non-special decimal floating-point number x in precision n has the form:

$$x = (-1)^s 10^E d_0.d_1d_2 \dots d_{n-1}$$

where $s \in \{0, 1\}$, the exponent E is an integer between two given integers E_{\min} and E_{\max} , and the mantissa $d_0.d_1d_2 \dots d_{n-1}$ is a fixed-point number written in radix 10; i.e., for i between 0 and $n - 1$, one has: $0 \leq d_i \leq 9$.

As d_0 may be equal to 0, some numbers have several representations and the standard does not distinguish them. Without changing the value set, one can require that if $E \neq E_{\min}$, then $d_0 \neq 0$, and this will be done in the following for the sake of simplicity. A number such that $d_0 \neq 0$ is called a *normal number*, and a number such that $d_0 = 0$ (in which case $E = E_{\min}$) is called a *subnormal number*. In this way, the representation of a floating-point number is uniquely defined.

Below $\text{ulp}(x)$ denotes the weight of the digit d_{n-1} in this unique representation; i.e., $\text{ulp}(x) = 10^{E-n+1}$.

The document [4], based on the IEEE 854 standard, defines three decimal formats, whose parameters are given in Table 1: decimal32, decimal64 and decimal128, with an encoding on 32, 64 and 128 bits respectively. This specification has been included in the 754r working draft.

Table 1. The parameters of the three 754r decimal formats.

Format	decimal32	decimal64	decimal128
Precision n (digits)	7	16	34
E_{\min}	-95	-383	-6143
E_{\max}	96	384	6144

2.2 The Bad and Worst Cases

Given a floating-point format, let us call a *breakpoint* a value where the rounding changes in one of the rounding modes, i.e., the discontinuity points of the rounding functions. A breakpoint is either a machine number (for the directed rounding modes) or the middle of two consecutive machine numbers (for the rounding-to-nearest mode).

For a given function f and a “small” positive number ε , a machine number x is a *bad case* when the distance between the exact value of $f(x)$ and the nearest breakpoint(s) is less than $\varepsilon \cdot \text{ulp}(f(x))$. For instance, if f is the exponential function in the decimal64 format ($n = 16$ digits), then the machine numbers 0.5091077534282133 and 0.7906867968553504 are bad cases for $\varepsilon \geq 10^{-16}$, since for these values, $\exp(x)$ is close enough to the middle of two consecutive machine numbers:

$$\exp(0.5091077534282133) = \underbrace{1.663806007261509}_{16 \text{ digits}} \underbrace{5000000000000000}_{16 \text{ digits}} 49 \dots$$

and

$$\exp(0.7906867968553504) = \underbrace{2.204910231771509}_{16 \text{ digits}} \underbrace{4999999999999999}_{16 \text{ digits}} 16 \dots$$

i.e., rounding $\exp(x)$ in the rounding-to-nearest mode requires to evaluate $\exp(x)$ in a precision significantly higher than the target precision. Similarly, with the following bad cases, $\exp(x)$ is very close to a machine number, so that rounding it in directed rounding modes also requires to evaluate it in a precision significantly higher than the target precision:

$$\exp(0.001548443067391468) = \underbrace{1.001549642524374}_{16 \text{ digits}} \underbrace{9999999999999999}_{16 \text{ digits}} 26 \dots$$

and

$$\exp(0.2953379504777270) = \underbrace{1.343580345589067}_{16 \text{ digits}} \underbrace{0000000000000000}_{16 \text{ digits}} 86 \dots$$

2.3 Searching for Bad and Worst Cases

Searching for bad cases in decimal is very similar to the search in binary. First the domain of the tested function is selected: arguments that give an underflow or an overflow are not tested, and some other arguments do not need to be tested either when a simple reasoning can be carried out (see Section 3.1 as an example). And like in binary [7–9], probabilistic hypotheses allow us to guess that the smallest distance amongst all the arguments to be tested is of the order of 10^{-n} ulp (divided by the number of exponents E), so that we are interested in $\varepsilon \sim 10^{-n}$ to get only a few bad cases²; i.e., we are interested in bad cases with about at least n identical digits 0 or 9 (possibly except the first one, which may be respectively 5 or 4) after the n -digit mantissa.

In the decimal32 format, the number of arguments to be tested is small enough for a naive algorithm to be sufficient: for each argument x , one computes $f(x)$ in a higher precision to eliminate the values x for which the distance between $f(x)$ and the nearest breakpoint(s) is larger than $\varepsilon \cdot \text{ulp}(f(x))$. Of course, a faster algorithm, as those needed for higher precisions (see below), can be used. Anyway, since finding bad cases is rather easy for the decimal32 format, this paper will not focus on this format; the reader may find some results for the exponential function at <http://www.loria.fr/~zimmerma/wc/decimal32.html>.

In the decimal64 format, the number of the remaining arguments after reducing the domain is still very large (say, between 10^{17} and 7×10^{18} , depending on the function), and a naive algorithm would require several centuries of computations. Like in the binary double precision, one needs specific algorithms, and since the decimal arithmetic has the same important properties as the binary one (the machine numbers are in arithmetic progression except at exponent changes, the breakpoints have a similar form...), the same methods can be applied.

In radix 2, bad cases for precision n and any rounding mode are the same as bad cases for precision $n + 1$ and directed rounding modes³, so that the problem was restricted to directed rounding modes in [6]. This property is no longer true in radix 10, but the breakpoints are still in an arithmetic progression (except when the exponent changes, just like in radix 2), which is the only important property used by our algorithms. Indeed in each domain where the exponent of $f(x)$ does not change, one needs to search for the solutions of:

$$|f(x) \bmod u/2| < \varepsilon u,$$

where $u = \text{ulp}(f(x))$, which is a constant in the considered domain.

² This may not be true in some domains, for instance when the function can be approximated accurately by a simple degree-2 polynomial, such as $\exp(x) \simeq 1 + x + x^2/2$ for x sufficiently close to 0; in this case, one can get bad cases which are much closer to breakpoints than what can be estimated with the probabilistic hypotheses. This is not a problem in practice: A simple reasoning may be sufficient instead of an exhaustive search in this domain, or the search can be done anyway, possibly with different parameters.

³ Said otherwise, in radix 2, the breakpoints for precision n and all rounding modes are the machine numbers in precision $n + 1$.

To solve this problem, one splits the domain into subintervals, and in each subinterval, one approximates the function f by a polynomial P of small degree and scales/translates the input and output values to reduce the problem to the following (as in the binary case [6]):

Real Small Value Problem (Real SValP). Given positive integers M and T , and a polynomial P with real coefficients, find all integers $|t| < T$ such that:

$$|P(t) \bmod 1| < \frac{1}{M}. \quad (1)$$

The coefficients of the polynomial are computed using the MPFR library [10] in order to obtain guaranteed error bounds.

Then several fast algorithms can be used to solve the Real SValP. Lefèvre’s algorithm needs degree-1 polynomial approximations; as these approximations are valid on very small intervals, one also needs a way to determine these approximations very quickly [11]. The Stehlé-Lefèvre-Zimmermann (SLZ) algorithm allows to have polynomials of higher degrees and has a smaller asymptotic complexity [6], but with a high constant factor. It is based on Coppersmith’s technique, which uses the LLL algorithm for lattice reduction.

In order to make the implementation of the SLZ algorithm as efficient as possible, it is crucial to use an efficient LLL code. For instance, one should avoid using the text-book LLL algorithm making use of a rational arithmetic. In the implementation of the SLZ algorithm, it is better to use variants of the LLL algorithm relying on floating-point arithmetic rather than rational arithmetic.

In his PhD thesis [12], Stehlé describes three floating-point variants of LLL, respectively called “fast”, “heuristic” and “proved”. The corresponding codes are available at <http://www.loria.fr/~stehle/>. The proved variant implements the algorithm described in [13], whereas the other two can fail⁴ but are usually more efficient.

Remark 1. The above methods may no longer work well for the lowest subnormals, due to the loss of precision for these numbers. For instance, a low-degree polynomial approximation may be valid on an interval that contains only very few machine numbers. Nevertheless these few values may be tested separately with a naive algorithm, if need be.

3 The Application to the Exponential Function

An application is now presented: the correct rounding of the exponential function, denoted `exp`, in the decimal64 format. This is just an example: a similar work can be carried out for other functions. After a simple analysis of the function (Section 3.1), we consider the search for bad cases (Section 3.2).

⁴ In practice, when they fail, they loop forever; they may also return a badly-reduced basis. But in both cases, no bad cases will be missed.

3.1 Correctly Rounding the Exponential Function

Let us first recall the parameters of the decimal64 format, with a few more details. A non-special floating-point number x has the form:

$$x = (-1)^s 10^E d_0.d_1d_2 \dots d_{15}$$

where $s \in \{0, 1\}$ and $-383 \leq E \leq 384$. So, the largest finite machine number is $10^{385} - 10^{369}$, the smallest positive normal machine number is 10^{-383} and the smallest positive machine number is 10^{-398} .

Now let us briefly analyze the exponential function, assuming that the argument is a finite number, to eliminate the special cases. The exponential function is mathematically defined on the whole domain of real numbers, so that the value will never be a NaN. It is increasing, with $e^x \rightarrow +\infty$ when $x \rightarrow +\infty$, and $e^x \rightarrow 0$ when $x \rightarrow -\infty$. And the mathematical properties of the exponential function are such that there will be an overflow when x is larger than some value and an underflow when x is smaller than some value. Moreover, $e^0 = 1$, meaning that for values of x close to 0, the rounding of $\exp(x)$ is determined only by the rounding mode and the sign of x .

So, there are four couples of consecutive machine numbers (a^-, a^+) , (b^-, b^+) , (c^-, c^+) and (d^-, d^+) that determine the following five intervals:

$$\underbrace{-\infty \dots a^-}_{+0} \quad \underbrace{a^+ \dots b^-}_{\text{search}} \quad \underbrace{b^+ \dots c^-}_1 \quad \underbrace{c^+ \dots d^-}_{\text{search}} \quad \underbrace{d^+ \dots +\infty}_{+\infty}$$

where in intervals 1, 3 and 5, the rounded values in the rounding-to-nearest mode are respectively +0, 1 and + ∞ (the rounded values in the directed rounding modes can also be determined, keeping the same interval bounds for the sake of simplicity), and in intervals 2 and 4, a search for bad cases is needed. These interval bounds are determined below.

An argument x generates an overflow when the *rounded* result obtained assuming an unbounded exponent range exceeds the largest finite machine number $10^{385} - 10^{369}$. One has:

$$\log(10^{385} - 10^{369}) = \underbrace{886.4952608027075}_{16 \text{ digits}} 882469 \dots$$

and

$$\log(10^{385}) = \underbrace{886.4952608027075}_{16 \text{ digits}} 883469 \dots,$$

so that one gets an overflow if and only if $x \geq d^+$, with $d^+ = 886.4952608027076$ (x being a machine number).

Concerning a^- , one has:

$$\log(10^{-398}/2) = -\underbrace{917.1220141921901}_{16 \text{ digits}} 2 \dots,$$

so that in any rounding mode, $\exp x$ is rounded to the same value for any $x \leq a^-$, with $a^- = -917.1220141921902$: It is rounded to 10^{-398} in the rounding to $+\infty$ mode, and $+0$ in the other rounding modes.

Concerning b^+ and c^- , one has:

$$\log(1 - 10^{-16}/2) = -\underbrace{5.000000000000000}_{16 \text{ digits}} 125 \dots \times 10^{-17}$$

and

$$\log(1 + 10^{-15}/2) = \underbrace{4.999999999999998}_{16 \text{ digits}} 750 \dots \times 10^{-16},$$

so that one chooses $b^+ = -5 \times 10^{-17}$ and $c^- = 4.999999999999998 \times 10^{-16}$.

Finally, in the other domains, that is for x in

$$[a^+, b^-] = [-917.1220141921901, -5.000000000000001 \times 10^{-17}]$$

and in

$$[c^+, d^-] = [4.999999999999999 \times 10^{-16}, 886.4952608027075],$$

a search for bad cases needs to be done to be able to round $\exp x$ correctly in any rounding mode.

Remark 2. When x is close enough to 0, one could use the approximation $e^x \simeq 1 + x + x^2/2$ to find bad cases with much less computing time in this domain. But globally, one would gain very little since this is an easy domain (as the error on a polynomial approximation is very small compared to higher values of x , and the algorithms work much better).

3.2 Searching for Bad and Worst Cases of the Exponential Function

To search for bad cases, one first splits the tested domain into intervals in which both the argument x and the result $\exp(x)$ have a constant (possibly different) exponent. This has been done with a small Maple program.

As said in [11] and [5], one could test the inverse function, i.e., the logarithm, instead of the exponential when x is small enough (say, $|x| < 1$). The reason is that there are fewer machine numbers to test in this domain for the inverse function. However this domain requires very little computation time compared to those with high values of x .

The search for bad cases is performed with BaCSeL⁵, which is running on a few machines. The chosen parameters are: a working precision of 200 bits, $m = 14.6$ (the quality of the bad cases, i.e., $-\log_{10}(2\varepsilon)$), $t = 5.5$ (a parameter that fixes the size of the sub-intervals), $d = 3$ (the degree of the polynomials) and $\alpha = 2$ (a parameter for Coppersmith's technique). For values of x close enough to 0, the fast LLL variant fails, so that the proved variant is used in this domain.

Table 2. Some bad cases of the exponential function in the decimal64 format. The notation d^k means that the digit d is repeated k times.

x	$\exp x$
$6.581539478341669 \times 10^{-9}$	1.000000006581539 5 0 ¹⁵ 177...
$2.662858264545929 \times 10^{-8}$	1.000000026628583 0 0 ¹⁵ 318...
$3.639588333766983 \times 10^{-8}$	1.000000036395884 0 0 ¹⁵ 240...
$6.036998017773271 \times 10^{-8}$	1.000000060369982 0 0 ¹⁵ 379...
$6.638670361402304 \times 10^{-7}$	1.000000663867256 4 9 ¹⁵ 569...
$9.366572213364879 \times 10^{-7}$	1.000000936657659 9 9 ¹⁵ 883...
$7.970613003079781 \times 10^{-6}$	1.000007970644768 5 0 ¹⁵ 362...
$3.089765552852523 \times 10^{-5}$	1.000030898132866 0 0 ¹⁵ 241...
$1.302531956641873 \times 10^{-4}$	1.000130261678980 0 0 ¹⁶ 798...
$2.241856702421245 \times 10^{-4}$	1.000224210801727 5 0 ¹⁵ 118...
$7.230293679121590 \times 10^{-4}$	1.000723290816653 4 9 ¹⁶ 127...
$5.259640428979129 \times 10^{-3}$	1.005273496619909 4 9 ¹⁵ 739...
$9.407822313572878 \times 10^{-2}$	1.098645682066338 5 0 ¹⁶ 278...
$1.267914924960933 \times 10^{-1}$	1.135180299492843 0 0 ¹⁶ 706...
$5.091077534282133 \times 10^{-1}$	1.663806007261509 5 0 ¹⁵ 492...
3.359104074009002	28.76340944572687 5 0 ¹⁶ 904...
294.9551257293143	1.251363586659789 5 0 ¹⁵ 108... $\times 10^{128}$

Table 2 presents some bad cases that have been found so far in the domain $[10^{-9}, 6.907755278982137]$ (completely checked) and some subintervals of $[6.907755278982138, 421.1499284665963]$.

For $c^+ \leq x < 10^{-8}$, many bad cases have some patterns in their mantissa. For instance, one has the following bad cases with $\varepsilon = 3 \times 10^{-15}$ (look at the 8th, 9th and 10th digits):

$$\begin{aligned} &3.897940992403028 \times 10^{-9}, \\ &4.230932991049603 \times 10^{-9}, \\ &4.291382990792016 \times 10^{-9}, \\ &4.581289989505891 \times 10^{-9}. \end{aligned}$$

This comes from the fact that $\exp(x)$ can be approximated by $1+x+x^2/2+x^3/6$ in this domain, and even by $1+x+x^2/2$ for smaller values of x . Table 3 gives some other bad cases for $x < 10^{-9}$.

4 Conclusion

Like in binary arithmetic, correct rounding can be guaranteed in decimal arithmetic at a reasonable cost if the upper bound on the necessary precision for

⁵ Available on <http://www.loria.fr/~stehle/>.

Table 3. Some bad cases of the exponential function in the decimal64 format, for $c^+ = 4.999999999999999 \cdot 10^{-16} \leq x < 10^{-9}$. At most two bad cases are given per exponent.

x	$\exp x$
$5.999879998200072 \times 10^{-10}$	1.00000000059998800 ¹⁶ 431...
$6.000119998199928 \times 10^{-10}$	1.00000000060001199 ¹⁶ 567...
$1.01999999994798 \times 10^{-11}$	1.0000000001019999 ¹⁷ 646...
$1.03999999994592 \times 10^{-11}$	1.0000000001039999 ¹⁷ 625...
$1.0999999999395 \times 10^{-12}$	1.000000000109999 ²⁰ 556...
$1.1999999999280 \times 10^{-12}$	1.000000000119999 ²⁰ 423...
$1.1999999999928 \times 10^{-13}$	1.000000000011999 ²³ 423...
$1.3999999999902 \times 10^{-13}$	1.000000000013999 ²³ 085...
$1.9999999999980 \times 10^{-14}$	1.000000000001999 ²⁵ 733...
$2.9999999999955 \times 10^{-14}$	1.000000000002999 ²⁵ 099...
$1.9999999999998 \times 10^{-15}$	1.000000000000199 ²⁸ 733...
$3.9999999999992 \times 10^{-15}$	1.0000000000000399 ²⁷ 786...
$9.9999999999995 \times 10^{-16}$	1.0000000000000099 ²⁹ 666...

the intermediate computations is determined. This requires exhaustive tests on the whole input domain. While some subdomains can easily be handled, a large number of input values need to be tested.

For the 754r decimal32 format, the tests can be carried out with naive algorithms. However, for the 754r decimal64 format, specific algorithms needed to be designed and implemented. Partial results for the exponential function have been given in this paper. The current worst case — if we disregard very small values — is

$$\exp(9.407822313572878 \cdot 10^{-2}) = \underbrace{1.098645682066338}_{16 \text{ digits}} \underbrace{5000000000000000}_{17 \text{ digits}} 278 \dots,$$

which means that a faithful approximation to 34 digits, which corresponds to the decimal128 format, would be enough to guarantee correct rounding for the exponential in the decimal64 format. At the time being, tests are still running and results for positive arguments should be complete in a few weeks or months. Once this is over, the test of the negative arguments will be started. Then, other elementary functions could be tested as well, with the same algorithms. As a consequence, standards could recommend (or even require) correct rounding for these functions in these formats.

At the same time, the implementation is still being improved, to speed up the tests. But for the same reasons as in radix 2, some functions are still out of reach in some domains, like the trigonometric functions for large arguments. In such a case, a standard could recommend correct rounding for such functions only in a limited domain.

Acknowledgements

The writing of this paper was completed while the second author was visiting the University of Sydney, whose hospitality is gratefully acknowledged.

The last author acknowledges the support from the Schloss Dagstuhl International Conference and Research Center for Computer Science, in particular the Dagstuhl Seminar 06021 *Reliable Implementation of Real Number Algorithms: Theory and Practice*, which stimulated the writing of this article.

References

1. IEEE: IEEE Standard for Binary Floating-Point Arithmetic, ANSI/IEEE Standard 754-1985. Institute of Electrical and Electronics Engineers, New York (1985)
2. IEEE: IEEE Standard for Radix-Independent Floating-Point Arithmetic, ANSI/IEEE Standard 854-1987. Institute of Electrical and Electronics Engineers, New York (1987)
3. Cowlishaw, M., Schwarz, E.M., Smith, R.M., Webb, C.F.: A decimal floating-point specification. In Burgess, N., Ciminiera, L., eds.: Proceedings of the 15th IEEE Symposium on Computer Arithmetic, Vail, Colorado, USA, IEEE Computer Society Press, Los Alamitos, CA (2001) 147–154
4. Cowlishaw, M.: Decimal arithmetic encoding strawman 4d, draft version 0.96. Report, IBM UK Laboratories, Hursley, UK (2003)
5. Lefèvre, V., Muller, J.M.: Worst cases for correct rounding of the elementary functions in double precision. In Burgess, N., Ciminiera, L., eds.: Proceedings of the 15th IEEE Symposium on Computer Arithmetic, Vail, Colorado, IEEE Computer Society Press, Los Alamitos, CA (2001) 111–118
6. Stehlé, D., Lefèvre, V., Zimmermann, P.: Searching worst cases of a one-variable function using lattice reduction. *IEEE Transactions on Computers* **54**(3) (2005) 340–346
7. Dunham, C.B.: Feasibility of “perfect” function evaluation. *ACM Sigum Newsletter* **25**(4) (1990) 25–26
8. Gal, S., Bachelis, B.: An accurate elementary mathematical library for the IEEE floating point standard. *ACM Transactions on Mathematical Software* **17**(1) (1991) 26–45
9. Muller, J.M.: *Elementary Functions, Algorithms and Implementation*. Birkhauser, Boston (1997)
10. Fousse, L., Hanrot, G., Lefèvre, V., Pélissier, P., Zimmermann, P.: MPFR: A multiple-precision binary floating-point library with correct rounding. Research report RR-5753, INRIA (2005)
11. Lefèvre, V.: *Moyens arithmétiques pour un calcul fiable*. PhD thesis, École Normale Supérieure de Lyon, Lyon, France (2000)
12. Stehlé, D.: *Algorithmique de la réduction de réseaux et application à la recherche de pires cas pour l’arrondi de fonctions mathématiques*. PhD thesis, Université Henri Poincaré – Nancy 1 (2005)
13. Nguyen, P., Stehlé, D.: Floating-point LLL revisited. In: Proceedings of Eurocrypt 2005. Volume 3494 of Lecture Notes in Computer Science., Springer-Verlag (2005) 215–233