

How Computers Handle Numbers

A.k.a. Computer Arithmetic Uncovered

N.M. Maclaren
Computing Service

nmm1@cam.ac.uk

May 2011

Summary of Course

This course covers just the bare essentials: i.e. what every person who uses a computer to handle numbers needs to know. It describes the basic arithmetic model used by the almost all of the hardware and software that you will encounter, and how you should write code that will give the answer you expect. It warns about the most common traps only, and often just raises a flag and says “*Don’t go there*”, “*Don’t trust any implementations*” or “*At this point, call an expert*”. Some of its warnings may appear to be overstated, but most of them are actually **understated**.

There are some foils of explanations for a few of the bald statements available as an extra. Most readers are strongly advised to ignore them, as they are neither friendly nor easy to understand, and to simply follow the advice given here. Note that even they are **not** an expert-level description of those areas, and there are comparable complexities in areas that are not explained. Anyone who is interested should consult me (the author) for more information, as a naïve Web search will merely cause confusion.

Users should contact me if they have a numerical problem, especially in the areas of debugging techniques and writing portable, reliable programs. There could well be follow-on courses, and some possible topics are mentioned at the end; requests would be welcomed.

Non-mathematicians should not panic when mathematical terminology is used, as everything that is introduced will be explained in very simple terms. The explanation of the terms may help when reading books or Web documents.

Some people will notice that not everything said here is strictly true, as some things have been oversimplified. For example, “*all computers*” means “*all computers that most of you will ever come across*” and not “*all computers currently in use for any purpose*”, let alone “*all computers that have ever been built*”.

There is some more information on the Web page:

Arithmetic/

You may also like to look at the course currently given to computer scientists:

<http://www.cl.cam.ac.uk/1011/FPComp/>

There is a good, clear book by *Foreman S. Acton* on numerical coding that describes most of the basic principles of how to avoid precision loss and invalidity:

Real Computing Made Real: Preventing Errors in Scientific and Engineering Calculations

Naturally, it explains only how to prevent **some** forms of error!

Stratospheric Overview

The numbers considered here fall into three main classes: the integers (whole numbers), as introduced in primary school; the real numbers, as used in secondary school science; and the complex numbers. If you are not familiar with complex numbers, you needn’t worry about them, as this course spends very little time on them. It does assume that you are familiar with integer and real arithmetic, as used in GCSE science.

For mathematicians only: the classes are the ring of integers of characteristic zero, \mathbb{Z} , and the fields of real and complex numbers, \mathbb{R} and \mathbb{C} . No other classes of number are covered.

Computer hardware has limited approximations to these; for example, the integers are limited in size, and real numbers are represented by limited precision floating-point ones. The software (compilers, libraries etc.)

often extends the hardware in many ways, such as by providing unbounded integers and building complex numbers out of real ones. The principles described here are largely language-independent, and apply to almost every computer language and application, including Python, Perl, Java, Excel, Matlab, C, C++, Fortran, R, C#, Maple, Mathematica etc.

In mathematics, these classes of number are defined by (and define) their properties, such as associativity – i.e. $(A+B)+C$ and $A+(B+C)$ are equivalent, and so are $(A*B)*C$ and $A*(B*C)$. Almost all of the problems with computer arithmetic arise because the computer's approximation doesn't preserve all of the expected mathematical properties, and the programmer accidentally assumes something that is not true. This is a problem with software facilities as much as hardware ones, and with integers as much as with floating-point numbers.

DON'T PANIC†

Anyone put off by the pure mathematics terminology should not panic when it is used, as everything that is introduced will be explained in very simple terms. The explanation of the terms may help when reading books or Web documents.

This may be a very complicated area, but this course is intended to be a map through the minefield. If you follow the recommendations, and take even moderate care, you can avoid most problems and recognise the ones that you can't avoid. Recognising when you are getting into dangerous areas is the key to reliable programs and debugging, because there is no problem so hard to locate as one you don't even know can happen. Some problems you can do nothing about except be aware that they exist, and to check whether they have caught you out when things do go wrong. And, as we all know, Murphy's Law says that they will.

† *With apologies to the late Douglas Adams.*

Consistency/Sanity Checking

I will now take a brief diversion on what were called good programming techniques 25 years back, and are now called software engineering. One of the best defences against numeric problems is to put in a lot of consistency and sanity checking. This should check not just whether values are valid but also whether they are realistic; one of the first symptoms of numeric failure is often that intermediate values start being physically unrealistic.

There is no need to check everything, everywhere, because the aim is to detect failures early (i.e. fairly close to where they started) and not to identify the actual cause. It is enough to make debugging simpler, not to eliminate it a much harder task), so most tests can be fairly simple. Even the simplest tests often detect a failure cascade before it has destroyed all evidence of where it started. The diagnostics should include at least an identification of which test failed, so that you know where to start debugging. For example:

```
if (speed < 0.0 .or. speed > 3.0e8) call panic('Speed error in my_function')
```

WARNING: this example is unreliable; see later about NaNs.

Ideally, this should be placed at the entry to and exit from every major unit of code, and should check all the arguments and global data that are being used, returned or changed, but one must temper perfection with realism. The best engineered programs include full (if simple) state, argument and result checking on entry and exit to all major sections of code, at least as an option.

In larger, more structured programs, a good rule is to write a sanity checker for every major data structure. This makes it much easier to add the checking but, more importantly, makes it easy to track down the first failure when debugging. For example, LAPACK routine `dposv` solves a solution of linear equations using a positive definite matrix held in a triangular form and a matrix of right hand sides; you might write a sanity checker for an upper triangular matrix and another for a rectangular one. The following checks for sanity on both entry and exit – if the first is fine, and the second isn't, your matrix is probably too ill-conditioned and you then need to track down why.

```

call sanity_upper (n, a, lda)
call sanity_rect (n, nrhs, b, ldb)
call dposv ('u', n, nrhs, a, lda, b, ldb, info)
call sanity_upper (n, a, lda)
call sanity_rect (n, nrhs, b, ldb)

```

You should note that this will not slow your program down significantly, if it is done carefully. Solving linear equations takes $O(n^3)$ operations, and simple check of a matrix takes only $O(n^2)$ operations, so the checking time is small compared to the solution time.

It is very common for this sort of approach to double the time until the program compiles, but halve the time until it starts to work (simultaneously). That is because coding and removing syntax errors is generally only a small part of the effort involved in programming. Also, these techniques apply more generally than to just numeric errors; for more information, see:

Debugging/

Do not be afraid to write your own class, because you do not need to use any more memory and modern compilers will compile them efficiently. You can then check the values systematically. This is especially useful for arithmetics like complex, because you could use simple inline code for most operations and check just multiplication, division and other slow actions. If you do this, do not forget to initialise the objects when they are created.

Integer Arithmetic

Paradoxically, problems with simple integer arithmetic cause as much trouble as those with the more complicated real and complex arithmetic, because most people don't expect trouble. It is mostly trivial, and works just as you would expect. By far the two most important causes of problems are overflow and mixing signed and unsigned integers – the latter is a problem only for languages that have both concepts (mostly C-like), of course.

The only other area that causes significant problems is division and remainder, where at least one of the values is negative. The exact effects are language-dependent, so do check up before relying on them. The value is always truncated towards zero if both values are positive. Division by zero is an error, of course, and so is taking the remainder with respect to zero.

As mentioned above, software can extend the hardware in many ways, and this is good example. It is fairly straightforward to emulate multiple-precision integers (e.g. 128-bit ones using 64-bit hardware), and this can be extended to providing effectively unbounded integers to the program. These usually have no limit on their size, except for the amount of memory needed to store them! They are built into Python, are provided by the BigInt package in Perl, and there are libraries that can be used from C, C++, probably Java and possibly Fortran. Perhaps the most widespread is the GNU Multiple Precision library, GMP. The ancestor of such use is Lisp, and most algebraic packages, such as Maple or Mathematica, use them – plus a few odd programs, like `bc` under Unix.

Simple operations on them can be implemented quite efficiently, but multiplication, division, remaindering and base conversion are slow. In good packages, such as Python, they are as easy to use as bounded integers with the benefit that they eliminate the complexities associated with overflow. Indefinite growth will crash the program, and it is easy to create unreadably large numbers; a good programming technique is to test whether values are within the bounds of reason and stop if not.

Current Integer Hardware

Nowadays, every computer uses binary, twos' complement, integers. Binary means that numbers are expressed in base 2: e.g. using conventional ordering (i.e. most significant bit first) on 8 bits, $01010011 = 2^6 + 2^4 + 2^1 + 2^0 = 83$ (decimal). Twos' complement describes how negative numbers are implemented, and says that the most significant bit is treated as negative, so that $11000101 = -2^7 + 2^6 + 2^2 + 2^0 = -59$. Clean numeric code does not assume any of this, but many languages do expose the integer representation to programs, as described later.

Similarly, all hardware supports both 16- and 32-bit integers, which have ranges of -32768 to 32767 and -2147483648 to 2147483647, respectively), and most support 64-bit ones (c. -9.2×10^{18} to c. 9.2×10^{18}). A very few support 8-bit or 128-bit ones (-128 to 127 and c. -1.7×10^{38} to c. 1.7×10^{38}), but many compilers will emulate 8-bit integers by operating on them in 16 bits, and storing back only the least significant 8 bits. Some will emulate 128-bit integers using 64-bit ones.

Nowadays, every computer will wrap the result on overflow and continue execution (e.g. $2*83 = -90$ and $4*83 = 76$, quietly), but it usually sets a flag which can be tested by the compiler. Some compilers and applications check that, and then fix up the exception. Also, optimise compilers often rearrange code in ways that are mathematically but not computationally equivalent. The result is that an operation in a program's **code** may not wrap even when the **hardware** does; this is described in more detail later.

Problems with Wrapping

Programmers using languages that wrap, quietly, on overflow need to know what it implies. For example, $(M*N)/N$ may not be the same value as M . Even if the code doesn't assume that, optimisation means that the compiler may, which has the consequence that overflow can make a program behave incorrectly in ways that does not look plausible when reading the code. Few people believe this, even when it happens to them, but it is so.

A realistic example (i.e. I have seen it catch several people, including me) is when declaring a multidimensional array and calculating its total length. On a system which can address more memory than the size of an integer (e.g. a 64-bit system with 32-bit integers, which is a **very** common combination nowadays), the array may be definable but the size may wrap. This is a major trap in Fortran, C, C++ and probably other languages. For example, consider code like:

```
parameter (n = 1800)
double precision d(n,n,n)
call init(d,n*n*n)
```

If $n*n*n$ wraps in 32 bits, this is equivalent to calling `init(d,1537032704)`, which means that only part of the array would be initialised, with the consequential chaos. You can test for overflow fairly easily, in several ways, such as checking that you get the same result in double precision real arithmetic, or dividing the total by a factor and checking the result. For example, in this case:

```
ntotal = n*n*n
if (ntotal /= n*dble(n)*n) call panic(...)
if ((ntotal/n)/n /= n) call panic(...)
```

Handling of Overflow

The biggest problem with computer integers is overflow, and at least 90% of all bugs associated with ordinary (signed) integer arithmetic are caused by the program mishandling it. Some languages avoid the problem, and users of those languages need worry only about things like array overflow (e.g. attempting to access element 12345 of a 100-element array), not integer overflow as such.

Some always use floating-point (e.g. Excel and Matlab); some convert large integers to floating-point (e.g. Perl and R); some convert large integers to arbitrary size (e.g. Python); it is now very rare to trap integer overflow and diagnose it, but it used to be more common. All of those are fairly safe options for most uses. Java and C#, however, take the remainder (i.e. wrap) modulo $|2^{bits}|$, which is generally **NOT** what you want, because it means that undetected overflow will often cause your program to do things like access the wrong element of an array.

In other languages (e.g. C, C++ and Fortran), integer overflow is **undefined**, and this is a major cause of wrong answers, crashes etc. It is important to realise that undefined means just that, and it does not mean that the values will wrap modulo $|2^{bits}|$, even when the hardware does. This means that wrapping is **dangerous** because it encourages mistakes. Simple tests for what happens are **usually** misleading, **most** books and Web pages are misleading, and undefined behaviour is **NOT** equivalent to system dependence. This point cannot be stressed too strongly.

The main reason is that optimisation may cause anything to happen, and this includes when the compiler is running in its least optimised, debugging mode. Compilers do not simply map code to hardware operations, but produce sequences that have an equivalent effect, and they assume the mathematical properties of integers when doing that. Almost all language standards say that they can generate any code that is mathematically equivalent, but need do so only for valid operations, and it is the programmer's fault (and problem) if the result overflows. What is more, the effects are likely to change with compiler versions, as well as between systems.

There are sometimes options to trap integer overflow, diagnose and stop the program. NAG Fortran always does and gfortran `-ftrapv` enables it. gcc and g++ `-ftrapv` will trap *some* overflows, but the C and C++ rules do not allow everything to be trapped. The C# `checked` keyword raises an exception on integer overflow. When these are available, they are the best solutions, and are strongly recommended.

This is not the course to go into details, but a compiler might well recognise code including the following and make the indicated assumptions:

- If $D = (A*B)/C$, then $D/A = B/C$
- If $A > 0$ and $B > 0$, then $A*B > 0$

But neither is true if $A*B$ overflows, so the code will assume something that is not true, and chaos ensues. As mentioned above, few people ever believe this, even when faced with the evidence.

Let us give an over-simplified (and not very convincing) example. We have code like:

```
B = C = D = 5000
A = B*C*D      [ = 445948416 ]      Wrong
E = A/D        [ = 89189 ]          Wrong
print E        [ = 89189 ]          Wrong but consistent
```

This fairly often actually compiles into code that executes vaguely like:

```
E = A/D        [ = 89189 ]          Wrong
A = B*C*D      [ = 445948416 ]      Wrong but consistent
    print E    =>    print B*C
print E        [ = 25000000 ]        But this is E!
```

Integer Formatted I/O

Integers can be written (whether for input or output) in any base: binary 01010011 is decimal 83 is octal (base 8) 133 is hexadecimal (base 16) A3. Sometimes this is explicit, using more-or-less well-understood conventions, as `2r01010011` or `0xa3`, but not always. C and C++ have a rather nasty convention where a leading zero indicates octal, in some contexts only, as in `0133`. Unix sometimes uses unqualified octal, which means that you need to know the context to decide what value 136 (or, for that matter, 0136) represents. Generally, most users will use decimal for formatted I/O, but it is quite easy to use any other base.

Consequences of Binary

As mentioned above, the actual format is not usually important, and clean numeric code will work unchanged on systems where the actual base used for integers is not 2 – there have been computers that used base 10 for integers. But there are a few things you need to be aware of. The first is that the terms *base* and *radix* are interchangeable in this context – some authors use one and some the other!

In binary, bitwise AND, OR, NOT etc. make sense, so integers can be used as arrays of flag bits. This is a little unclean, but most languages support it nowadays; some of the specialist ones, like R. You are strongly advised to do this only with positive numbers; each language's rules on how negative ones are handled are different, they may be undefined or defined to have bizarre effects.

Similarly, left and right shift by N places are multiplication and division by 2^N . However, it is important to note that shifting is safe **only** if all of the arguments and results are positive and the number of places to shift is **smaller** than the number of bits in an integer. For example, using 8-bit twos' complement, don't trust `2r11000101 = -59` shifted right two places to be `2r11110001 = -14` or `2r00110001 = 49`; similarly, `1234`

shifted right 200 places will probably **not** be zero (4 is more likely, but it could be anything). This lunacy is a relic of 1950s electronic constraints.

Unsigned Arithmetic

C, C++, Java, Perl and perhaps other languages support **unsigned** integer arithmetic. This is exactly like twos' complement, except that the top bit is **not** negated, so negative integers cannot be represented and we now have (in 8 bits) $11000101 = 2^7 + 2^6 + 2^2 + 2^0 = 197$. Most of the basic arithmetic is performed modulo $|2^{bits}|$ (i.e. overflow wraps modulo $|2^{bits}|$). Because there are no negative numbers, bitwise logical operations and shifts are defined for all values, but still use only shifts smaller than the number of bits in a word.

*For mathematicians: this is **not** $GF(2^N)$, nor is it true modular arithmetic, despite the occasional claim that it is one or the other by people who ought to know better.*

Because numbers are always non-negative, there are some strange effects, such as subtracting 5 from 3 gives a large positive number. Division and remaindering (A modulo B) are not performed modulo $|2^{bits}|$ (*this statement isn't likely to make much sense to non-mathematicians*), which has some subtle consequences.

More importantly, the interactions between unsigned integer types and both signed integer types and floating-point types are foul. A few details are given later, especially for C and C++, which also has the trap that the basic *char* type may be either signed or unsigned, but most people should merely note that this area is a minefield and keep signed numbers (both integer and floating-point) and unsigned integers as separate as possible.

Fixed-Point Arithmetic

This has a fixed number of digits after the decimal point, where that number of digits (the precision) is a part of variable's type and does not depend on the value being stored. Despite appearances, this has more in common with integer arithmetic than floating-point, and indeed is usually implemented as scaled integers. It is heavily used for financial calculations, but is rare in scientific computing, except for `bc`, `dc` etc. Several algebraic packages support it, and there are often extensions or libraries that provide it.

It is generally as easy to use as integer arithmetic, though no easier (e.g. overflow is a similar mess), except for a few problem areas. The rounding of results after multiplication or division (remainder is easy, surprisingly) is often not what is expected. Both mixing fixed-point numbers with different precisions and converting to and from other types and formats are definitely tricky. Using mathematical functions (square root, logarithm etc.) is even harder, but few people do that with fixed-point.

Scaled fixed-point is fixed-point with a separate scale factor. This was commonly used in the 1950s, but was replaced in the 1960s by true floating-point, which is almost always better. C# has resuscitated it with its `decimal` type, though it may be implemented using IEEE 754 decimal floating-point (usually emulated). Scaled fixed-point is almost always a complete waste of effort, because one or other of true fixed-point or true floating-point are better (depending on the requirement). It is closely related to unnormalised floating-point, which is covered later under IEEE 754 decimal floating-point. There is also a proposed DEC64 format, which is not covered further.

Rational Arithmetic

This is the rational arithmetic we were taught at school, and is a main mode in Mathematica, with unlimited size integers. Other than the problems cause by exponential explosion in the number of digits (e.g. running out of space or time or printing incomprehensible results), it does exactly what you would expect.

Rational arithmetic with fixed-size integers does exist, and has some advantages and uses, but is so esoteric that I may be the only person in Cambridge ever to have used it! It had a brief flurry of interest under the name "fixed-slash" arithmetic some 30 years ago.

Basics of Floating-Point

It may help if we go through the basic principles of floating-point. When used for display purposes, this is often called scientific notation, and the form we shall consider here uses a leading zero before the decimal point (i.e. not the most common scientific format). Numbers are of the following form:

$$\text{sign} \times \text{mantissa} \times \text{base}^{\text{exponent}}$$

The mantissa is treated as a fixed-point number strictly less than 1 and greater than or equal to the inverse of the base (i.e. the first digit after the decimal point is non-zero). Such numbers are called *normalised*; if the first digit after the decimal point is zero, they are called *unnormalised* or *denormalised*. It is then multiplied by the sign and scaled by the exponent to give the value; for example, $+0.12345 \times 10^2 = 12.345$. Zero is represented specially. Here, we shall consider only formats that use a fixed precision, P (the number of digits after the decimal point); these are also called the significant digits. Such numbers have a relative accuracy of $\times(1 \pm \text{base}^{-P})$.

A easier way of thinking of this is to regard the sign and mantissa as a fixed-point number between -1 and 1, exclusive (i.e. neither -1 nor 1 is possible), scaled by the exponent. This ignores the precision, but can be useful when designing programs, as this simpler model is generally enough to analyse the overflow behaviour.

Floating-Point versus Reals

Floating-point is not mathematical real arithmetic and the differences are important. The easiest (and usually the best) way to think of floating-point arithmetic is that it is effectively non-deterministic to within the accuracy limits – i.e. the results are guaranteed to be within those limits, but are otherwise unpredictable – which means that relying on the precise value is almost always a mistake. However, all of the differences lead to either a result very close to the expected one (i.e. within $\times(1 \pm \text{base}^{-P})$), or arise only for numbers so near zero that they can't be represented properly.

Aside: the above statement was uncontroversial 25 years back, but is now regarded as heresy, because the current dogma is to confuse consistency and accuracy, so you will read precisely the converse in many or most modern books.

The reason that the difference is important is that fixed-point numbers break many of the rules of real arithmetic, and floating-point ones break even more, so ordinary mathematical notation cannot simply be mapped to program code. Wrong assumptions cause wrong answers, and a regrettably high proportion of published papers are based on results that are completely meaningless, for this reason alone. The key is to think in terms of floating-point rules, not the ones of real arithmetic – this is easier to do than it sounds, as 50 years of Fortran experience shows! Write what you mean, and don't assume that something that is mathematically equivalent is necessarily numerically equivalent; with a bit of practice, you will avoid most problems without having to think about them.

Invariants

Invariants are properties that always hold, and are very important and useful in mathematics and software engineering. Let's start with the properties that both floating-point and real numbers share.

- Both are commutative: i.e. $A+B = B+A$, and $A*B = B*A$.
- Both have zero and unity: i.e. $A+0.0 = A$, $A*0.0 = 0.0$, and $A*1.0 = A$.
- Both have negatives (additive inverses): i.e. each A has a $B = -A$, such that $A+B = 0.0$.
- Both are fully ordered: i.e. $A \geq B$ and $B \geq C$ means that $A \geq C$, and $A \geq B$ is equivalent to $B < A$.

Unfortunately, the differences are rather more extensive.

- Floating-point numbers are neither associative nor distributive: i.e. $(A+B)+C$ may not be $A+(B+C)$, $(A*B)*C$ may not be $A*(B*C)$, $(A+B)-B$ may not be A , $(A*B)/B$ may not be A , and $A+A+A$ may not be $3.0*A$.
- They are not continuous (for any of addition, subtraction, multiplication or division): e.g., for addition, $B > 0.0$ may not mean $A+B > A$, $A > B$ and $C > D$ may not mean $A+C > B+D$, and $A > 0.0$ may not mean $0.5*A > 0.0$.

- They have no multiplicative inverse: i.e. not all A have a $B = 1.0/A$, such that $A*B = 1.0$. And $1.0/A$ can fail for very small non-zero A .

Now, the above is true for all fixed-size floating-point, whether it is implemented on a computer or written by hand, but I will bet that most of it wasn't mentioned at school. Some other differences will be mentioned later, especially with regard to exceptional (error) values, which are very different.

Current Floating-Point Hardware

The current floating-point standard is called IEEE 754 (IEEE 854 is a minor variation), and is also known as ISO/IEC 10559, and almost all computers use it nowadays. See:

<http://754r.ucbtest.org/standards/754.pdf>

This is binary (i.e. the base or radix is 2) and signed magnitude (i.e. the sign bit is just that, and it is not twos' complement). The details are messy and so will be omitted here. It has 32- and 64-bit formats, with sometimes 80- or 128-bit extended formats (see later). The accuracies and ranges of the standard formats are:

32-bit, 4 byte, 'single precision': the accuracy is 1.2×10^{-7} (23 bits), and the range is 1.2×10^{-38} to 3.4×10^{38}

64-bit, 8 byte, 'double precision': the accuracy is 2.2×10^{-16} (52 bits), and the range is 2.2×10^{-308} to 1.8×10^{308}

The format includes representations of exceptional values, which indicate errors and other odd states, and they are discussed later.

There is a new version of the IEEE 754 standard, which includes (very complicated) decimal floating-point arithmetic, but it won't have significant impact before 2010, if at all. It is possible that the new IEEE 754 may not be accepted, or will make this whole area much worse. Even a partial explanation of this statement is beyond the scope of this course, but some aspects of it will be remarked on.

Other Sizes and Bases of Floating-point

The use of other sizes of floating-point is not recommended, and you should ask if you need such a facility. The reason is not that it is a major problem in principle, but there are serious portability, accuracy and performance problems. For example, many of the operations may be emulated by trapping the instruction, which can be very slow, or the mathematical functions may not be as accurate as they should be. In addition to the Intel arithmetic described below, some systems have 128-bit IEEE floating-point in several different variations; the new version of the IEEE 754 standard may specify a recommended one, which may help.

Exact floating-point arithmetic is usually futile, because the number of digits required increases exponentially with time in the majority of codes; it has its uses in the hands of an expert. Interval arithmetic is trendy, has some uses, but isn't much easier to use. Arbitrary precision arithmetic is easy, but is very much out of fashion, because IEEE 754 dominates people's thinking, though Mathematica has it.

The most common extended format is the native arithmetic of the Intel and AMD systems, and has an accuracy and range of:

80-bit: the accuracy is 1.1×10^{-19} (63 bits), and the range is 3.4×10^{-4932} to 1.2×10^{4932}

This is typically stored in 12 or 16 bytes (96 or 128 bits), and should be avoided completely if at all possible. It is becoming much less commonly used by compilers, though they (and some applications) often use it internally to increase the accuracy of complicated 64-bit operations. Whether and how they do is one of the causes of differences in results between compilers and systems. See:

http://www.intel.com/design/pentium4/manuals/index_new.htm

Decimal Floating-Point

This was added to IEEE 754R at IBM's instigation, with support from Intel, and both stated that they would deliver it in hardware. Python has a module that supports their specification (inefficiently), and the

C standard is intending to specify it as a Technical Report (optional extra). It is now available on IBM POWER6 and later using IBM's C and C++ compilers (or DB2 - IBM's database!) Amusingly, some of the early computers used it, but it was abandoned because binary was better.

Despite the claims of its proponents, it is **not** a panacea for problems with floating-point – but that doesn't mean that it is any worse than binary, just different. It is frequently claimed to deliver exact arithmetic (e.g. Python does), but those claims are just propaganda. Expressions like π , 1.0/3.0, 1.01²⁵ and almost all results in scientific code are no more exact than they are in binary.

It doesn't make providing decimal fixed-point any easier, despite claims that it does, but few (if any) scientific programmers have any interest in that, anyway.

It solves a few problems but adds others, of which the worst is that the usual formulae for the average of two values do not necessarily lie between the two values. In binary floating-point, if $a \leq b$, then $a \leq a/2+b/2 \leq b$ and $a \leq (a+b)/2 \leq b$, but that is not necessarily true in decimal floating-point. Consider for example, (0.61,0.63) and (0.47,0.47) with 2 decimal places in the mantissa.

There are quite a few more “gotchas”, but they are very arcane and you will probably never notice them. It isn't hard to write code that works well with both binary and decimal floating-point, and all NAG code was like that up until about 1990.

But Intel have abandoned plans to put it into hardware, and IBM has no plans to add it to their Fortran, so one can ask if it will ever be relevant to scientists – that currently seems unlikely.

Very Small Numbers and Underflow

This is a complicated and messy area, but causes relatively few problems. The reason is that it generally doesn't matter what value is returned for a negligible result as long as it **is** negligible. It can cause performance problems and occasional program misbehaviour, but rarely causes wrong answers.

Denormalised numbers were introduced by IEEE 754, but are not always implemented on nominally “IEEE 754” systems. They are just numbers with the minimum possible exponent but zeroes after the decimal point – for example, in an artificial decimal representation, 0.00123×10^{-308} . Whether or not they are supported, numbers that are too small to represent in the chosen format are quietly replaced by zero (this is called underflow); underflow is never trapped nowadays, though it used to be, because so many codes fail if it is.

There are some numeric advantages to allowing denormalised numbers (“soft” underflow), and some disadvantages, but the claims that one approach is numerically better than the other are pure dogma. They can be very slow, because they are often not implemented in hardware, and the system takes an interrupt to fix them up in software, so there is often option to simply replace by zero (“hard” underflow). Few programs notice the difference.

A good rule is to use double precision to minimise the chances of falling into a trap – 10^{-308} is much less likely than 10^{-38} – and, if you do have trouble, it is almost always safe to replace very small numbers by zero. Perhaps the worst trap is that most programs and applications compiled for hard underflow go wrong if presented with a denormalised number, so it is important not to import one into such programs when using binary I/O. And, of course, don't mix code compiled for hard and soft underflow in the same program.

Examples of the sort of thing that can cause trouble when numbers approach zero are:

- A small reduction in the size of a number may lose accuracy or turn it into a zero. For example, $(A/2.0)*2.0$ may not be exactly A, and $A > 0.0$ does not mean $2.0*A > 1.5*A$.
- Two numbers may be different but their difference may be zero. For example, $B > C$ does not mean $B-C > 0.0$.

Error Handling and Exceptions

Old maps are supposed to have marked uncharted waters where travellers had disappeared without trace with “*Here be dragons*” – this area needs the same annotation.

The following is what the ordinary user of floating-point **needs** to know; I wish that it were not so, and it used not to be the case, but it is now. Most of the details have been omitted but we may return to a few

aspects later. **Please** contact me if you hit a problem in this area.

IEEE 754 formats include both positive and negative zeroes, positive and negative infinities, and NaNs (Not-A-Number). Except as described below, it is possible to ignore the signs of zero, and users are strongly recommended to do just that; at best, the sign of zero is grossly unreliable.

The infinities represent values that have overflowed. In simple cases, they imply numbers that are too large to represent, but this is not a reliable assumption. For example, `exp(1000.0)` will overflow and deliver an infinity, and taking the logarithm of plus infinity will also deliver plus infinity. But `log(exp(1000.0))` really represents the value 1000.0, which isn't very large.

NaNs represent the result of an erroneous operation, typically because it is mathematically invalid. They aren't as simple as that, but explaining their complexity is vastly beyond the scope of this course. Treat them as simple indications of an invalid calculation, representing a value that is numerical nonsense.

In theory, both propagate appropriately, and you can tell if a complicated calculation failed because the result is a NaN but, in practice, the error state is not reliable. The reasons for this are fiendishly complicated and tied up with the politics of international standards (IEEE 754, C99 and others), but the executive summary is that this area is a minefield. This course attempts to explain how to traverse it with only a moderate risk of getting your foot blown off.

What Can Be Done?

I shall give a partial solution to the problem before describing the problem in detail. The pedantic answer to "*What Can Be Done?*" about this situation is "*Not a lot*", because the problem is caused by broken language standards and specifications. However, there are some programming techniques that can help reduce the damage.

The first is to put in consistency and sanity checking, as described above – because one bogus value breeds others, such checks often work very well. The second is to use double precision to reduce the likelihood of overflow – because 10^{308} is so much larger than 10^{38} , this often helps a great deal. Note that this can often speed up a program, by avoiding underflow and denormalised numbers, which often interfere with pipelining even when they don't cause interrupts.

But don't assume that the first time your checking catches a bogus value is the first time that one has appeared, or that a failure to catch any means that there were none. No matter how careful you are, some will always slip through the net, so always cross-check your results whenever possible. Note that these remarks apply to **all** classes of error, including array bound overflow and pointer problems.

Overflow, Divide by Zero, Infinities etc.

Mathematica uses a fancy format and rarely overflows, Excel delivers "NUM!", NAG Fortran always traps, and some compilers (for several languages) have an option to trap. All others deliver an *infinity* of the right sign. Python's numpy by default delivers an infinity, but with a warning – it also has an option trap it. In itself, returning an infinity would be perfectly reasonable and safe – i.e. it is just using the affine extension of the real numbers.

At school, you were taught that dividing by zero is an error, and NAG Fortran, Python, Perl, Excel and Matlab trap that and issue a diagnostic. There are no problems with those.

C, C++ and other Fortrans rarely trap, and C99 in "IEEE" mode, Java and R are specified not to, because IEEE 754 states that both overflow and divide-by-zero give an infinity. This makes sense when implementing interval arithmetic, but is a very serious trap for ordinary programmers. To implement this, zero is signed in IEEE 754 and the sign is regarded as meaningful, but this ignores the facts. Many operations deliver an indeterminately signed zero (e.g. `X-X`), and those language standards permit a compiler to rearrange code and remove obviously irrelevant operations in a way that will reverse the sign. For example, if we have:

```
B = A-A;   C = -B;   D = C+0.0
```

then:

```
B = C = D = 0.0
```

but:

$$1.0/B \neq 1.0/C \text{ and } 1.0/C \neq 1.0/D$$

Compilers will often optimise code like $1.0/(A+0.0)$ to $1.0/A$; after all, they **are** mathematically equivalent. The result of this is that you should **never** trust the sign of an infinity, and it would be a much better specification if the languages returned a NaN when dividing by zero. But, because they don't, this is a trap for the unwary, and the sign of infinities is essentially meaningless.

If anyone doubts this, try the following Fortran program or its equivalent in your favourite language. The output suddenly switches from “-infinity” to “+infinity”.

```
program fred
double precision :: x = -1.0d-300
do k = 1,25
x = x-0.9d0*x
print *, 1.0d0/x
end do
end program fred
```

Because the root cause of this problem is that zero's sign is treated as meaningful, but it generally isn't meaningful, exactly the same problem applies to functions that test sign bits, such as Fortran SIGN and C copysign. Those are even worse, because they will extract the sign from NaNs, and IEEE 754 specifies explicitly that the sign of a NaN is undefined.

If you can, trap errors, diagnose them and stop the program. In IEEE 754 terms, the serious errors are *overflow*, *divide by zero* and *invalid*. NAG Fortran always traps those. With gfortran, gcc and g++, use `-ffpe-trap=invalid,zero,overflow`. With Intel compilers (ifort, icc and icpc), use `-fpe0`. With Python numpy, call `seterr(over='raise',divide='raise',invalid='raise')`, or you can use 'call' rather than 'raise'.

NaNs and Error Handling

The general principle of IEEE 754 is that invalid operations should result in a NaN, and that operations on NaN arguments should return NaNs. For example, $0.0/0.0$, ∞/∞ and $\infty-\infty$ all return a NaN. With luck, taking the square root of a negative number will, too. This is numerically respectable, and corresponds to the behaviour of the 1960s and 1970s statistical and engineering applications that pioneered the use of error values. Unfortunately, it does not clearly separate its numeric primitives from its semi-numeric ones (e.g. those that are needed to operate on NaNs specially), and using the latter makes it very easy to lose the NaN state by accident.

C99 and, to a lesser extent, Java and other C-derived languages have taken the approach that, if a non-NaN result can possibly be justified, then it should be preferred. This means both that it does not return a NaN when it should (e.g. `atan2(0.0,0.0)` returns $-\pi$, -0.0 , $+0.0$ or $+\pi$, depending on the signs of the zeroes) and provides a great many ways in which the NaN state is lost, quietly. See:

<http://www.cs.berkeley.edu/~wkahan/JAVAhurt.pdf>

Aside: be careful when reading that. Professor Kahan may be the world's leading numerical computationalist, but many people believe that his designs (including IEEE 754) are sharp-edged tools for numerical experts only, and not safe tools for use by the inexperienced.

He is unquestionably correct that the lack of either traps or flags, compiled with IEEE 754 value semantics, makes Java numerically very dangerous – and the same is true of C99 (and increasingly C++). The result of all this is that you should not trust compilers and applications to detect their errors, and should not even trust them to return or display a NaN when they have performed an invalid operation. Here are a few examples of the many traps for the unwary:

- `int(NaN)` is often 0
- `max(NaN,1.23)` may be 1.23
- Comparisons on NaNs usually deliver false, quietly

I once tried writing some numerical code using the last convention, seeing if I could do it reliably. When I analysed it, I had got the NaN handling of 20% of the tests wrong. As I somewhat arrogantly regard

myself as one of the best numerical coders around, I do not have any confidence that other people will do better. The comparison issue has another problem – take the first example of a sanity checking function:

```
if (speed < 0.0 .or. speed > 3.0e8) call panic('Speed error in my_function')
```

That test will **not** work if `speed` is a NaN, because both comparisons will probably yield false. Now, in theory, we could change it to:

```
if (speed != speed .or. speed < 0.0 .or. speed > 3.0e8) call panic(...)
```

That is perhaps the cleanest solution, but many compilers will optimise that test out, so it isn't reliable. Or we could test for validity and negate the result using `.not.`, but probably the safest approach is to write the ungainly code:

```
if (speed >= 0.0 .and. speed <= 3.0e8) then
  continue
else
  call panic(...)
endif
```

However, cautious users won't trust even that, especially with high levels of optimisation. I am afraid that there really is no good solution – and, despite claims, C99 does not help.

Complex Numbers

Generally, these are simple to use, but that does **not** apply to C99, for reasons that are too complicated to go into here. They are invariably implemented in software as (real,imaginary) pairs of floating-point numbers. All of Python, Fortran, C++, Matlab and R have them built-in, and C99 more-or-less does. There are optional packages for Perl and Java.

I/O is usually done on the raw floating-point numbers (which is a bit unclean), it is easy to lose the imaginary part by accident, and special functions can be slow and unreliable, but those are relatively minor problems. The biggest one is that exception handling is uniformly untrustworthy – often catastrophically so. If you even approach the limits, the usual result is that you will get mathematically wrong answers, quietly.

For mathematicians only: the complex numbers are compactified by adding a single point at infinity, turning them into the Riemann sphere, but this conflicts with using the Cartesian representation.

An analysis of this is far too foul for this part of the course, but interested users should try to write a function to perform complex division that gets correct answers right up to the limits. It isn't a pretty problem. But you will discover that it is easy to produce code that works very well for almost all values, and will go wrong only for ones that are almost overflowing (within a factor of two of it). If you can avoid producing such numbers, then you will have no problems.

So the recommendation is not to trust the exception handling an inch, and to put sanity checks into your code, but otherwise not to worry.

Mixed Type Expressions

Converting signed integer to floating-point to complex is usually not a problem, provided that the precisions are adequate. When converting N-bit integers to N-bit floating-point, and the values are such that they cannot be held precisely (above 8.3×10^6 in 32-bit or 4.5×10^{15} in 64-bit), you may come across some extremely weird rounding. That is about the worst problem. Converting unsigned integers to floating-point usually works, but a few implementations are buggy when the top bit is set (i.e. when the number would be negative if it were signed).

Going the other way is trickier. Floating-point to integer truncates towards zero, in almost every modern language; that is the easy bit. Overflow is undefined in C, C++ and Fortran; the effect is completely unpredictable, and can include plausible but weirdly wrong integers. Java defines what happens, but it is a very dangerous specification in that it will often give incorrect results, quietly. Converting complex to

floating-point takes the real part in all of those languages, quietly, thus providing an opportunity to lose the imaginary part by accident (a **very** common error, especially in languages with implicit declaration).

Converting infinities and NaNs between any of integer, floating-point and complex is effectively undefined in C, C++ and Fortran, in any direction, and there is no such thing as typical behaviour. Java's behaviour is similar to that for overflow – i.e. defined but very dangerous. You do not want to know what C99 says about complex infinities and NaNs, but the executive summary is that they are more broken than you would believe possible. Fortran says nothing about them, which is a vast improvement.

So my recommendation is the same for exception handling in complex arithmetic, though with extra care when converting floating-point to integer.

Other Arithmetics

Let us use Hamiltonian quaternions as an example. They are not covered in this course, because they are rarely used and very specialist. Very few languages have them built-in, but you can get add-on packages for most of them; in languages that have appropriate type extension facilities, they can be made to look very like built-in types.

There are almost no extra problems over using complex numbers, and almost all of the warnings about using those apply. Indeed, the main difference is that they are not commutative, so care must be taken with expression order. Similar remarks to most other advanced and specialised arithmetics, though the details differ, including true Galois fields and more.

Formatted Input/Output

Formatted output is generally safe and, in this context, output includes all conversions of numbers to text (including string types). The conversion accuracy of very large or small numbers is not always good, though this rarely causes trouble. Infinities and NaNs are usually handled correctly nowadays, but not always in the way you might expect.

The most common trap is that some people don't realise that values like 0.1 are not exact in binary – decimal 0.1 is binary 0.0001100110011001... (i.e. binary 0.0 0011 recurring). Because of the base mismatch, only 15 digits are guaranteed reliable for 64-bit floating-point (6 for 32-bit), but 18 (9 for 32-bit) digits must be printed if the number is to be read in again unchanged. So don't trust $10.0*0.1$ to be exactly 1.0, and remember that how many digits you print depends on what you want to do with the output.

And, of course, you should check on infinities, NaNs and denormalised numbers before relying on them; if the implementation is poor, it will fail with those.

Formatted input (including all conversions from text to numbers) is far more of a problem than output, because it needs to deal with bad input. Format errors and overflow are often undefined, and the application or compiler will often fail to detect either, or handle them sanely even when it does. The behaviour can be very weird indeed. Python and Matlab seem pretty safe, but I have not checked most of the others.

Infinities, NaNs and denormalised numbers are always unreliable, and you should never trust the implementation without checking. It is good programming practice to include a minimal cross-check yourself on all data that you read in, to pick up typing mistakes, but it will also help pick up conversion errors. You are recommended to check both the values for sanity, and that the count of numbers is what it should be.

Binary (Unformatted) I/O

This just writes the bits in their internal format to file and reads them back again; it is fast, easy and preserves the value precisely. Because it uses the internal format, don't use it between systems without checking, as you will be caught out sooner or later. Unfortunately, it often depends on the compiler and its options or the application, so it isn't entirely safe even within the same system except for a single program.

Different languages use different methods, and Fortran is very different from C and C++; solutions do exist to transfer between them, but ask for help. Derived and other complicated types may add extra problems; assuming that those don't arise, we can now give an almost complete checklist for how to use this safely.

The systems, applications or languages must use same formats, the same sizes and the same endianness. The formats are defined by the application or language and perhaps its options. The sizes are defined by those, and whether the system is 32- or 64-bit whether it is Unix, a Microsoft Windows system or MacOS and, again, perhaps the system's options – for example, you can create 32-bit programs under most 64-bit systems.

Endianness is whether the bits of integers are stored least significant first (“little endian”) or most significant first (“big endian”). Floating-point numbers are usually stored in the same order as integers, and complex numbers are always stored in the order real, imaginary. Common little endian systems include Intel x86 and clones (including AMD, AMD64 and Intel EM64T) and Alpha; common big endian ones include Sun SPARC, SGI MIPS, HP PA-RISC and IBM PowerPC. Intel Itanium can be either, just to confuse the issue, and the mixed endian systems are probably all now defunct.

And, to reiterate, do check on the denormalised number, infinity, and NaN handling before relying on it, and don't import them into programs that are not expecting them, or anything may happen.

One of the best ways of doing this is to create a file with the ‘numbers’ $0.0, \pm 10^k$ ($k = -323 \dots + 308$), $\pm \text{inf}$ and NaN, read them in, and calculate all 8 million combinations of the basic operations, add, subtract, multiply, divide and compare of. You should get the same results on both systems! The crudest method is to print them out to (say) 12 digits in a fixed format and compare them using *diff*, but this can be done more intelligently.

Single Precision (32-bit)

Do **not** use single precision for serious calculations unless you have analysed your problem carefully. The first reason is that, even if you need less than 6 digits of precision, that cancellation, error accumulation and ill-conditioning mean that it is quite likely that your results may not even have their first digit correct. The second is that you are much more likely to trip across exceptions and, as described above, that is Bad News. Let us look at a trivial but realistic example for the accuracy point:

The quadratic $x^2 + 10^4 \times x + 1$ has roots of c. 10000 and 0.001; now if we apply the usual formula $(-b \pm \sqrt{b^2 - 4ac})/(2a)$ in 32-bit, we get roots of c. 10000 and zero. This may not matter, but it might lead to a division by zero, giving infinity and not c. 1000, which could then make subsequent results wildly wrong.

Another problem is that modern machines have lots of memory, and this allows for big problems; even the most stable of big problems need more accuracy. For example, the error in the solution of linear equations or eigensystems is often of the form:

$$\textit{precision} \times \textit{matrix_dimension} \times \textit{condition_number}$$

In single precision, the first is 1.2×10^{-7} , the second could easily be 1.0×10^4 and even well-behaved matrixes may have a condition numbers of 1.0×10^3 . Oops. Even Fourier transforms are susceptible to this problem.

Numerical Analysis

This analyses the effects of approximate calculations, both those caused by floating-point arithmetic and the way that variations build up in the underlying mathematics, and is **not** covered in this course. The Department of Applied Mathematics and Theoretical Physics has three courses on it. In the days when it was taught more widely, it was the subject of year-long diplomas, and even the scientists' introductions were several afternoons long; unfortunately, it is almost never taught nowadays. You are recommended to use a reasonably trustworthy package or library, because it is more reliable than doing your own thing.

The NAG library is most general reliable library, by a very large margin; the Numerical Algorithms Group is the leading numerical software house in the world. There are some very good open-source libraries (e.g. LAPACK), but many others are seriously unreliable or worse. For reasons that I am not going into now, do **not** trust Numerical Recipes, most especially not the first edition. The same applies to the Web, which has not been called The Net of a Million Lies for no reason. In both cases, some parts are good, and some are ghastly. For a very brief introduction to numerical analysis, see:

[Arithmetic/na1.pdf](#)

and, if you are really interested:

Arithmetic/na2.pdf

I shall now give a two minute introduction to numerical analysis, but the term “over simplified” is understating the case!

You can assume that your results will almost never be better than your input – the old acronym GIGO stands for “*garbage in, garbage out*” and is entirely correct. But you should **not** assume that your result is accurate to even the minimum of your input accuracy and the machine precision, as errors can often build-up rapidly, even exponentially or worse. If you fall into such a trap, even converting single to double precision may not help and your only solution is to improve the algorithm. Again, let’s give a trivial and semi-realistic example:

1:	1	1.0 to 1.0	1.1×10^{-14}
2:	2	2.0 to 2.0	6.7×10^{-12}
3:	6	6.0 to 6.0	7.4×10^{-10}
4:	24	24.0 to 24.0	6.3×10^{-08}
5:	120	120.0 to 120.0	3.6×10^{-06}
6:	720	719.6 to 720.4	5.1×10^{-04}
7:	5040	4914.1 to 5192.6	2.8×10^{-02}
8:	40320	-21032.1 to 102033.9	$1.5 \times 10^{+00}$
9:	362880	-18299033.4 to 16980266.1	$4.9 \times 10^{+01}$

The K ’th differences of a sequence give an approximation to the K ’th derivative, and the K ’th derivative of x^K should be K factorial. Well, as numerical analysts know, that is a classically ill-conditioned problem. But did you? In double precision, here are the true values, the range of the calculated values and the relative errors of the K ’th differences of x^K calculated from 0.5 to 2.0 in steps of 0.01. They give 1 significant figure for $K=7$, and nonsense for $K=8$ and above. Oops.

This example is only semi-realistic, though I have seen a user fall into this very trap. More realistic examples include the solutions of linear equations, determinants, eigensystems, polynomials, regression and parameter fitting (linear and non-linear), analysis of variance, ordinary and partial differential equations, finite elements, graph theoretic work (e.g. Dirichlet tessellation a.k.a. Voronoi diagrams) and so on. In fact, it is quite rare to have a non-trivial problem which doesn’t show this effect to some extent.

At the lowest level, the cause of most such loss of accuracy is cancellation, which occurs when you subtract two nearly-equal values. Obviously, this includes adding two values which have different signs, but it also includes dividing two nearly-equal values and multiplying one by something that is nearly its inverse. If we assume that numbers have P digits of precision, and X and Y have Q leading digits in common, then $X-Y$ and $X/Y-1.0$ will have $P-Q$ digits of precision.

Restructuring expressions can help a great deal with this for little cost. Where it matters, consider changes like the following:

$$\begin{aligned} (X+D)**2-X**2 &\Rightarrow (2*X+D)*D \\ x**5-y**5 &\Rightarrow (x**4+x**3*y+x**2*y**2+x*y**3+y**4)*(x-y) \\ \sin(x+d)-\sin(x) &\Rightarrow \sin(x)*(cos(d)-1.0)+cos(x)*sin(d) \end{aligned}$$

Watch out for large summations, too, because they are a classic place for this problem to occur, and consider *Kahan summation* which often (but not always) increase the accuracy for very little cost. Unfortunately, it may be implicit in the algorithm, and is common with ones that use numerical derivatives; the only solution is to use a more stable algorithm, if there is one.

Unfortunately, fixing the lowest level does not always help, because cancellation is often an *emergent property* of an algorithm. The other problem is that almost any method will give good results in simple, small cases – you can even solve well-conditioned 3×3 sets of equations using Cramer’s rule – but the poor methods will fail or give wildly inaccurate answers in complex, larger cases, such as the ones you really want the answer for and can’t check by hand! There are a few things that can be done which help.

- Put consistency checks into your program, such as cross-checking that the solution of an equation returns reasonable values for the parameters and that it really does satisfy the equation to within reasonable accuracy bounds! This also helps to speed up debugging.

- Use high-quality algorithms and libraries, as recommended above. Some of them will give error estimates, and it is sometimes worth trying a slower but more accurate method if your answers seem weird.
- Try perturbing your input – not just the data, but the parameters you have copied from the literature – and see if that changes the results unreasonably. If so, you have a possible problem. Unfortunately, if there is no change, it does **not** mean that there is no problem – see paragraph 5 in:

How Futile are Mindless Assessments of Roundoff in Floating-Point Computation?

<http://www.cs.berkeley.edu/~wkahan/Mindless.pdf>

- As always, seek expert advice, which does not mean the most self-confident of your colleagues. Until you know which are reliable and which are unreliable, be careful. Students should ask their supervisors or Directors of Studies to put them onto someone suitable, when necessary.

Topics Not Covered

There is a huge amount that has not been covered, including the details of any of the above topics. Indeed, there are too many to list, but the following are some subjects that could be covered in future courses; please request any that would interest you.

- Parameterisation in C, C++, Fortran etc.
- Interval arithmetic and its uses
- Introduction to numerical analysis
- C99 and its consequences
- Older or rarer arithmetic issues
- Number handling in external protocols
- Model, use and analysis of IEEE 754
- IEEE 754R and decimal floating-point
- Interactions with operating systems
- Implementation techniques and implications
- Mathematical models of computer arithmetic
- And so on ...

Reminder – Trapping Options

NAG Fortran traps everything by default. For gfortran, gcc and g++, use `-ftrapv -ffpe-trap=invalid,zero,overflow`. With Intel compilers (ifort, icc and icpc), use `-fpe0`. For Python numpy, call `seterr(over='raise',divide='raise',invalid='raise')`, or you can use `call` rather than `'raise'`. For C#, use the `checked` keyword or option.

This Is Now “Later”

There are no notes on the extra explanations, largely because they were written primarily because I am expecting at least some people to claim that things aren't as bad as I make out. I don't recommend most users even to look at them, but some masochists may want to look at them and recoil in horror. The main reference is the C99 standard, on which the UK voted “no”, but they also refer to C++11 (which has just been standardised), Fortran 2003 and IEEE 754R.