

Quad Double Arithmetic in Lisp

Richard Fateman
Computer Science
University of California
Berkeley, CA, USA

March 10, 2006

Abstract

In a numerical calculation, sometimes double-precision floating-point arithmetic is not sufficiently precise. An alternative is to use a software package implementing arbitrary-precision extended floating-point arithmetic, and choose a suitable precision. This choice may be too slow for convenience and there are intermediate possibilities combining extra precision and only modest loss in speed. Sometimes merely doubling the number of bits in a double-floating-point fraction is enough, and so arithmetic on double-double (DD) operands is useful. Another possibility is to go for another doubling to quad-double (QD) arithmetic: instead of the machine double-floats giving about 16 decimal digits of precision, QD supplies about 64 digits. DD and QD as used here provide the same exponent range as ordinary double.

Here we describe how we incorporated QD arithmetic in a Common Lisp system, providing a smooth interface while adding only modest overhead to the run-time costs (compared to C or C++). Most of the lessons from QD can be used for other versions of arithmetic.

1 Introduction

Quad-double (QD) arithmetic is described in Hida, Li, Bailey [1], and programs implementing QD are available free. Although one can think about QD numbers as indivisible 256-bit quantities storing one number, the storage mechanism is that of four double-float numbers in an array.

It is assumed in the documentation [1] that access to the library will be from traditional batch languages such as C++, with wrappers supplied for C and Fortran 90. While these languages may represent the most common numerical environments, this choice inhibits experimentation: programs must be carefully composed and compiled before any numbers can be produced. An interactive shell for prototyping computations can be much more productive in the programming and debugging stage.

Lisp is of course not the only interactive computation system; in fact others that are more widely known in the mathematical computation community include computer algebra systems like Maple, Mathematica, Macsyma, Axiom, or numerical systems like Matlab, Octave, MathCad.

Lisp has the advantage of having a rather simple programming model, and it has a number of packages available for which it can be used as a scripting level. One feature of most Lisp implementations is that an optimizing compiler is always available, and so one can define and test even a one-line function; debugging it interactively, and then using a one-line command to compile it and make it run much faster. In this way a QD arithmetic module may be easily transferred from interactive to compiled (even batch) production mode. In Lisp a compiled module can still be accessed interactively and will cooperate with most debugging tools. In some sense the highly-touted Java notion of JIT or “just in time” compiling has been supported in Lisp for several decades.

2 If you write in Lisp, here's what you do

These instructions pertain specifically to Allegro Common Lisp running on Windows/Intel computers. Other implementations should have similar setups.

If you are not especially concerned about speed, but want the extra precision, or perhaps are just testing an algorithm, obtain our program directory and try this:

```
(load "packs")
(load "ga.fasl")
(load "qd.fasl")
(load "qd.dll")
(load "my-own-system")
(in-package :ms)
```

Here we assume that the package for `my-own-system` is called `:ms`. A template, which can be copied into your own file, for doing this is shown below. The file `my-own-system.lisp` is this:

```
(defpackage :ms
  (:use :ga :cl)
  (:shadowing-import-from :ga
    "+" "-" "/" "*"
    = > <
    sin cos tan
    atan asin acos
    sinh cosh tanh
    asinh acosh atanh
    expt log exp sqrt
    1- 1+ abs
  ))

(in-package :ms)
(eval-when (compile) (load "qd"))

;; Your programs go here. For example,

(defun fact(x)(if (= x 0) 1 (* x (fact (1- x)))))
```

This factorial program can be used for ordinary Lisp numbers or qd numbers. It looks like this:

```
ms(196): (fact 50)
30414093201713378043612608166064768844377641568960512000000000000
ms(197): (fact (qd:into 50)) ;; note, exponent follows the Q
0.30414093201713378043612608166064768844377641568960512Q65
```

If you want extra speed, you can do this:

```
ms(197): (compile 'fact)
```

If you want even more speed, it may pay to put the function definition in a file with the package (etc) header and compile the whole file. That way you can be sure that the environment of the compilation process is correct. (It is entirely possible to run previously-compiled programs in a Lisp system that knows nothing about the equivalent in-line expansions of code; once the in-line expansion macros are loaded, further compilations will use them.) For programs with loops or recursive calls, using `with-temps` or `dsetv`, discussed below, may be helpful in getting the best efficiency. The effect of using these functions on expressions computing QD numbers is that run-time type-testing and most of the storage re-allocation is bypassed.

What is bypassed? Ordinarily some checking is needed to discriminate among the various types that can be fed at run-time into a function like “+”. By assuming all arguments to numeric functions are QD numbers, all functions used in the enclosed expressions pass their arguments directly to the *c*-coded library. If they are passed values that are not QD numbers, or cannot be converted to QD numbers *at compile time*, then the results will probably be some uninformative error message such as “Segmentation violation”.

3 How it is done

We define a “package” called “ga” for generic arithmetic. For a program that runs “using” the ga package, each of the function operators like +, *, sin, cos, etc. is overloaded so that if any of their operands is a QD number, the operation is done in QD arithmetic.

The simplest way of producing a QD number from an ordinary Lisp number is `(qd:into X)`. The original X can be a single- or double-float, an integer of any number of digits (yes, Lisp allows this), or a ratio of two arbitrary-precision integers.

Lisp normally uses a prefix notation so that addition would look like `(+ 1/4 1/4)` to evaluate to `1/2`. In QD, `(+ (qd:into 1/4) 1/4)` returns `0.5Q0`. In this last expression we show both the use of the `qd:into` program to convert an ordinary lisp number to a QD number, as well as the result showing the notation for a QD answer: a string of characters with the exponent following a Q. Naturally it is possible to extract the exponent and fraction part of a result using the same programs used by the `print` routine, in the program listings.

Anticipating that readers unfamiliar with Lisp will object to the prefix notation, let us point out that there are numerous parsers written in Lisp that will allow infix notation, e.g. `a+4*c` or `1/4+into(1/4)` for such expressions. We supply a parser with this collection of generic arithmetic systems. Nevertheless we will continue to use prefix notation for clarity and consistency with ordinary Lisp.

A complete listing of the QD functionality and further examples are available in the directory with the files.

Other overloadings for Lisp, for example MPFR, the (arbitrary-length) floating-point with rounding package associated with the GMP package, can be similarly accessed by `(mpfr:into X)` and are also documented further in that directory.

4 For Lisp fans

A data structure and type for a QD number, which we call `aqd` is defined and QD methods are essentially refactored into the “normal” arithmetic associated with numbers native to Lisp. Then the arithmetic associated with new operations on QD or mixed types are overlaid. These are actually built upon programs with names like `two-arg+` out of which an n-ary “+” is built. For unary functions like sin or cos, the situation is even simpler. The similarities among the Lisp-to-C interfaces for `sin`, `asin`, `sinh`, `asinh` are exploited by repeating a macroexpansion of the piece of code, only changing the name “sin” to “asin” etc. via a parameter.

The Lisp n-ary functions like “+” are redefined as macros which expand forms like `(+ a b c)` into `(two-arg+ a (two-arg+ b c))` so as to allow the compiler to compile specific versions of `two-arg+`

depending on the declared or deduced types of `a`, `b`, or `c`. They need not all be QD numbers. If they are all QD numbers, and the computation is done inside a loop, it should run faster by using `with-temps` mentioned below.

Printing of QD numbers is implemented by conversion of each of the component double-precision numbers to exact (arbitrary-precision) rational numbers, exactly adding them together, and then taking the integer part after multiplication by a power (possibly negative) of the print radix. (The default radix is 10).

Many implementations of Common Lisp use a compacting garbage collector (GC). This means that data structures, once allocated, may be moved in memory at some later time. This has the unwanted side effect of making pointers into those objects (from C programs) invalid, if it is possible for the GC to be run during one of the QD programs. This would happen perhaps if the Lisp were allowed to run another multiprocessing thread during execution of a C program, and that other thread provoked a GC. While this cannot happen in the current Lisp implementation, it seems prudent to be aware of this possibility. The Allegro implementation allows arrays to be allocated in memory such that they are not relocated with a GC, and so as a precaution we have specified that QD arrays are allocated in this “lisp-static” space. We do not believe this is necessary, however.

5 Functional vs. State-Transition

If we chose a uniform functional notation for arithmetic, we would say `add(A,B)`; If we used an “infix” notation, then `A + B`. In Lisp, parenthesized prefix notation `(+ A B)` means “return the memory location of a *new* qd-containing array M containing the sum of A and B.” The numbers A and B are unchanged.

This notation extends so that to compute `A + B * C` write `(+ A (* B C))`

The state-transition version uses a different set of fundamental operations. We would say `add(A,B,M)`. This means “destroy what was in the target memory M, and overwrite it with the sum of A and B.” Thus to compute `A+B*C` we need to identify (or perhaps newly allocate) two temporaries, T1 and T2, and compute `mul(b,c,t1)`, `add(a,t1,t2)`; The answer is in T2.

The obvious clumsiness of the latter approach is a trade-off related to the associated freedom from storage issues. In fact, if the computation is done repeatedly in a loop, the same T1 and T2 can be used, and this will use a fixed amount of storage and perhaps save time too. Compared this to the functional version, where new pieces of memory would be allocated: eventually the functional version requires that something be done to reclaim storage used for variables whose values are no longer used.

We can, however, use either paradigm in Lisp by (essentially) changing the requirement “allocate a new array M” to “find a (perhaps previously used, but no longer needed) array, like T1 or T2.”

Since Lisp is perfectly able to figure out how many temporaries are need for this at compile time, and then re-use them as needed, we allow this:

```
(define foo (A B C) (qd:with-temps (+ A (* B C))))
```

which essentially associates with `foo` two temporaries as needed. One for the product and one for the sum. If `foo` is called twice and the first result is not to be overwritten by the second, it is important to make a fresh copy this way (`copy (foo x y z)`). Otherwise the return value from `foo`, a pointer to one of those temporary locations, will be aliased between the two calls.

Perhaps more often than not in scientific computing you really DO want to overwrite a location with a new value, as for example, updating an array.

Instead of using the Lisp idiom `(setf (aref G i j) (+ ...))` we implemented a (new) idiom `(qd:dsetv (aref G i j) (+ ...))` which destroys the previous value in $G_{i,j}$ rather than creating a new object and pointing $G_{i,j}$ to it. The name `dsetv` comes from Destructively SET the Value. The `qd` package prefix indicates that all components must be QD numbers for this to work properly: the compiler macro-expansion will enforce that assumption on all arguments of its known functions; all explicit constants encountered will be converted to QDs at compile time, and all variables are assumed to be QDs. The expansion of `with-temps`,

which can have any nested arithmetic expressions supported by QD arithmetic as well as `setq` makes extensive use of `dsetv`. As is usual, the Lisp function `macroexpand` can be used to look at the generated code that the compiler actually uses. The programs `dsetv` and `with-temps` can be used together; a good example of this use can be seen in the QD-FFT code. The FFT is computed in-place in the input array, and no additional storage is allocated, even though the program nominally looks like it is using functional programming. While `dsetv` could be elaborated to include more of Lisp, for example, capturing additional control structures, our current examples do not seem to need this. As an example of the code generation, consider the macro-expansion illustrated below.

```
;; macroexpansion of (with-temps (setq (aref x 10) (+ y (sin z))))))

;; We have edited this code for clarity.
;; gg is really a gensym.
;; reg1,reg2, reg3 ... are "registers" .. allocated QD numbers
;; at compile time. They can be overwritten.
;; Type declarations have been removed for brevity.
;; Notice direct calls to qd_sin and qd_add.

(let ((gg
      (let ((a1 (aqd-q y))                ;grab qd value in y
            (a2 (aqd-q (let ((a1 (aqd-q z));grab qd value in z
                            (tt (aqd-q reg3)))
                              (qd_sin a1 tt)    ;put sin(z) in reg3
                              reg3)))
            (tt (aqd-q reg2)))
        (qd_add a1 a2 tt)                ;put y+sin(x) in reg2
        reg2)))                          ;gg holds val in reg2
      (qd_copy_into (aqd-q gg) (aqd-q (aref x 10))) ;copy gg into x[10]
      gg)
```

6 How fast is this?

The amount of overhead in Common Lisp for the wrapping may depend critically on compilation optimization flags and making sure the compiler knows as much as possible about the arguments to functions. Allegro Common Lisp is able to compile a direct function call from the Lisp invocation (`sin x`) to the C language program `c_qd_sin` in about 3 instructions on an Intel X86 processor. This same implementation requires that the floating-point control word be set at the interface; this is a consequence of decisions related to saving state at process switching, and that QD arithmetic depends on rounding to double, not extended. Other lisp systems (and non-X86 versions of Allegro) do not need to reset this control word.

One might be concerned about the overhead of type discrimination needed to determine which of the (perhaps many) overloaded methods will be used for the argument types given to such general programs as “+”. The typechecking and its attendant overhead can be eliminated by `qd:dsetv` or `qd:with-temps`.

As an example, consider evaluating by Horner’s rule a polynomial t with 999 terms, beginning $x^{998} - x^{997} + x^{996} - \dots$. That is, the coefficients are alternating ± 1 . We evaluate this polynomial at 2. The value is about 1.78D300.

We timed it on an Intel 2.5Ghz Pentium 4, using Allegro Common Lisp 7.0. We averaged the times over 10000 trials.

- The time to evaluate exactly as a 300 decimal digit integer is about 1.3ms.

- The time to evaluate approximately as a 64-digit quad-double using Horner's rule written in Lisp, is about 1.6ms.
- The time to evaluate approximately as a 64-digit quad-double written in C++ is about .42 ms.
- The (best) time to evaluate approximately as a 16-digit machine double-float is about 0.017 ms.

The Lisp and C programs look like this:

```
(defun polyeval (qdlst x) ;;qdlst is a list of qd objects, x is a qd object
  (let ((sum (qd:into 0)))
    (dolist (i qdlst sum)
      (setf sum (with-temps (+ i (* x sum)))))) ;; 10% slower without "with-temps"
```

```
(defun polyevald(dlst x) ;; dlst is a list of doubles, x is a double
  (let ((sum 0.0d0))
    (dolist (i dlst sum)
      (setf sum (+ i (* x sum))))))
```

;; To speed this up, we need to insert some declares:

```
(defun polyevald (dlst x)
  (let ((sum 0.0d0))
    (declare (double-float sum x) (optimize (speed 3)(debug 0)))
    (dolist (i dlst sum)
      (declare (double-float i))
      (setf sum (the double-float
                  (cl::+ i (the double-float
                              (cl::* x sum))))))))
```

/* the C++ program, operating on a ‘‘flattened’’ polynomial */

```
void c_qd_polyevalflat( double *c, int n, const double *x, double *result) {
  double *p = c + 4*n; /* pointer to current coefficient */
  fpu_fix_start(NULL); /* set rounding to 'double' not double-extended on x86 */
  qd_real r = qd_real(p);
  qd_real xx = qd_real(x);
  for (int i = n-1; i >= 0; i--)
    { p -= 4;
      r *= xx;
      r += qd_real(p);
    }
  TO_DOUBLE_PTR(r, result); }
```

Observing the final result: If we send the QD numbers flattened out as an array of 3996 double-floats to a C++ program, we remove almost all the overhead attributable to using Lisp. Compare this to the time to execute the same polynomial evaluation making individual Lisp calls to qd-addition and qd-multiplication: There is a factor of 4 slowdown in using Lisp to script the Horner's rule¹.

¹The timings also depend on the computer: the slowdown factor for a Pentium III between Lisp and C++ using QD

On the Pentium 4, then, if we compute with QD numbers but they are not needed—we can use double-precision machine floats instead—then we are paying a penalty of about a factor of 25 (writing in C++) or a factor of 94 (writing in Lisp). Of course double-floats should be used when the extra precision is not needed; considering that multiplying a pair of QDs uses 16 double-float machine multiplications, this seems reasonable.

These figures can vary substantially depending on the compilers, both Lisp and C++ involved. We used optimization setting /O2 for the Microsoft Visual Studio (version 8) 2005 C++ compiler.

A final note on timing of Horner’s rule: this algorithm uses an equal number of QD adds and multiplies. Profiling the QD Horner benchmark suggests that the cost for a QD multiply is about 1.7 times the cost of an addition.

Another test, computing the fast Fourier transform suggests a similar ratio (of about 30) between double-float and QD.

We also programmed an alternative version of the FFT, still written in Lisp, but removing all tags—that is, the QD numbers cannot be distinguished in any way from other length-4 double-float arrays, provides a modest 15-25 percent boost. This version looks essentially like assembly language with lines like `(qd_mul a c t1)`.

7 What about other kinds of arithmetic?

If you like the QD version of overloading of Lisp, the same ideas are available for MPFR [2], a library in which computations can be done for much longer fractions and larger exponents. This arithmetic is several times slower at the QD length, but if QD is insufficiently precise or has inadequate range (which is the same as the double-precision exponent range) then MPFR, built on the GMP library, is an alternative. In fact, the code base described here was, in part, originally written for use with MPFR.

Yet other kinds of arithmetic, including symbolic and polynomial arithmetic, can be supported. For these domains there is much less concern about linkage overhead: the time to multiply two polynomials is usually many times greater than the time needed to determine the types of the inputs and the allocation of the space for the result.

8 Pieces left out, limitations

There are several difficulties in the design of Common Lisp’s arithmetic support that require delicate treatment. These include parameter-count variations and the transition from real to complex numbers.

Two built-in Common Lisp functions of interest to QD arithmetic, namely `atan` and `log` can take one or two arguments. In the case of `atan`, the two-argument version is mapped to `atan2`, a well-known function. This mapping requires some scuffling around at compile time, noticing whether the call is to `one-arg-atanh` or `two-arg-atanh` and expanding to the right form as necessary. It is also necessary to convert one or the other of the arguments to QD form if it not already in the right form. The `log` function of two args is logarithm with respect to a particular base, the second argument. The processing of `log` is provided with a special recognition of `log10` which is available from the qd library. Otherwise `(two-arg-log a b)` is simply `/tt(/ (log a)(log b))`.

There are other one- or two-argument functions, namely `floor` and `ceil` we ignore, at least for now. They merge notions of division (by the second argument) and truncation.

A more serious issue is the transition from real to complex, which is not treated consistently in the QD library with Common Lisp. Consider the `sqrt` function, which among many others in Common Lisp, can

arithmetic is only 1.7, not 4. We expect that cache size is an issue, since the polynomial itself is 1000 qd numbers, each of which is 32 bytes. The Pentium III L1 data cache is only 16k bytes, and so can hold only half the polynomial. The Pentium 4 has 512k bytes, so all the data can fit.

return a complex result even if given a real argument. This is required of ANSI Standard Common Lisp, but the QD library does not support, at the moment, complex numbers. Furthermore, it is problematical in a situation in which types and hence storage formats “matter”. The original QD library exits unceremoniously in such situations, writing a error message in the process. We prefer that it return a “not a number” or `NaN`; perhaps even better (from a Lisp perspective) would be to provide a complex QD type. ‘ Not knowing whether (`sqrt x`) will return a real or complex number, Lisp is still obligated to store whatever is returned. This raises difficulties with respect to computational efficiency: perhaps returning `NaN` or signalling an exception is then preferable. Since a QD `NaN` has quite a few extra bits, more information could be conveyed: the function causing the problem, and (most of) the bits of the argument could be returned. Given either a `NaN` or an exception, it is easy to build a (light-weight) wrapper in Lisp around functions that sometimes return complex values, so that in the case of an otherwise unanticipated non-real result, the function can be recomputed as complex. This would be similar to automatically recomputing `sqrt(-5)` as though it been `complex_sqrt(-5+0*i)`. A Lisp function may not be inconvenienced by such a return value; a Lisp function using carefully declared type assumptions, as with our `with-temps` macro features, would by-pass these wrappers and consider the unexpected introduction of a complex value as an error². As another thought, the QD number format is rather generous in size, and given that only one component, the first one, is needed to display a `NaN`, the other three words can be used as an encoding to inform the consumer of the cause of the `NaN`, including perhaps a tag for the program counter, and one or more of the arguments.

9 Acknowledgments

Yozo Hida helped get `qd` running in a Windows dll for use by Lisp. A newsgroup posting by Ingvar Mattsson provoked us into thinking, once again, about generic arithmetic in Lisp. We especially thank Duane Rettig who patiently advised us how to coax substantially better performance from Allegro Common Lisp’s foreign function interface.

References

- [1] Yozo Hida, Xiaoye S. Li and David H. Bailey. “Quad-Double Arithmetic: Algorithms, Implementation, and Application” October 30, 2000 Report LBL-46996 (see <http://www.cs.berkeley.edu/~yozo> for links and related information)
- [2] Multiple Precision Floating-Point Reliable Library, <http://www.mpfr.org/>.

Appendix

Lisp source code for the QD interface, including comments and examples, can be found in the directory//<http://www.cs.berkeley.edu/~fateman/generic>, as well as a compiled “dll” suitable for windows x86 linkages.

²Note that small errors, perhaps caused by roundings, can change the argument to `sqrt` from zero or a small positive number to some small negative quantity. In such a case it may be prudent to allow square-root of a small negative number to be 0.