

using namespace cln;

CLN: A Class Library for Numbers

Richard B. Kreckel

Castro Urdiales, September 1 2006

Overview

- History of the library
- Feature Overview
- The Type System
bringing mathematical types and OO together
- Selected Implementation Aspects
 - Implementing `class cl_I : public cl_RA ...`
 - Fixnums and Bignums
- Small Example
printing the largest known perfect number
- Finally. . .
Applications

CLN History

late 1980s-1995: arbitrary precision types within CLisp (Bruno Haible et al.)
<http://clisp.cons.org/> (1987–today)
implementation languages: C and Assembler

1995: spin-off from CLisp
implementation languages: C++ and Assembler
purpose: make the arbitrary precision numbers of CLisp available to a broader public

1996: option to base low-level routines on more efficient GMP routines
<http://www.swox.com/gmp/> (MPN level only)
computation of 1 000 000 decimal digits of $\zeta(3)$

1999-01-12: release of CLN version 1.0

2000: maintainer change, since Bruno Haible was busy doing CLisp, Linux I18N, Unicode Support (libiconv), Glibc and several other free software projects

CLN History

late 1980s-1995: arbitrary precision types within CLisp (Bruno Haible et al.)
<http://clisp.cons.org/> (1987–today)
implementation languages: C and Assembler

1995: spin-off from CLisp
implementation languages: C++ and Assembler
purpose: make the arbitrary precision numbers of CLisp available to a broader public

1996: option to base low-level routines on more efficient GMP routines
<http://www.swox.com/gmp/> (MPN level only)
computation of 1 000 000 decimal digits of $\zeta(3)$

1999-01-12: release of CLN version 1.0

2000: maintainer change, since Bruno Haible was busy doing CLisp, Linux I18N, Unicode Support (libiconv), Glibc and several other free software projects

2006/2007: release of CLN 1.2 with support for huge numbers (>4GB each)

CLN Features

- rich set of number classes with unlimited precision
integers, rational numbers, floats, complex numbers, modular integers, even univariate polynomials
- natural mathematical syntax / type system
algebraic syntax through operator overloading ($z=x+y$ instead of `add(x,y,&z)`)
natural injections like $\mathbb{Z} \rightarrow \mathbb{Q}$ modeled with types
- speed efficiency
C++ compiles to good machine code, usage of assembler for critical parts and common CPUs ('i386', 'x86_64', 'alpha', . . .), asymptotically ideal algorithms (Schönhage-Strassen multiplication, binary splitting, etc.)
- memory efficiency
representation of small numbers as immediate values instead of as pointers to heap allocated storage
object sharing: $x+0$ returns x without copying it, etc.

CLN Type System

Natural injections in an OO environment: $\mathbb{Z} \rightarrow \mathbb{Q}$, $\mathbb{Q} \rightarrow \mathbb{R}$, $\mathbb{R} \rightarrow \mathbb{C}$, etc.

CLN types for those fields:

Integers \mathbb{Z} : `c1_I`

Rationals \mathbb{Q} : `c1_RA`

Reals \mathbb{R} : `c1_R`

Complex numbers \mathbb{C} : `c1_N`

CLN Type System

Natural injections in an OO environment: $\mathbb{Z} \rightarrow \mathbb{Q}$, $\mathbb{Q} \rightarrow \mathbb{R}$, $\mathbb{R} \rightarrow \mathbb{C}$, etc.


CLN types for those fields:

Integers \mathbb{Z} : `cl_I`

Rationals \mathbb{Q} : `cl_RA`

Reals \mathbb{R} : `cl_R`

Complex numbers \mathbb{C} : `cl_N`



CLN Type System

Natural injections in an OO environment: $\mathbb{Z} \rightarrow \mathbb{Q}$, $\mathbb{Q} \rightarrow \mathbb{R}$, $\mathbb{R} \rightarrow \mathbb{C}$, etc.


CLN types for those fields:

Integers \mathbb{Z} : `cl_I`

Rationals \mathbb{Q} : `cl_RA`

Reals \mathbb{R} : `cl_R`

Complex numbers \mathbb{C} : `cl_N`



Short-Floats `cl_SF`: sign, 17 mantissa bits, 8 exponent bits


Single-Floats `cl_FF`: sign, 24 mantissa bits, 8 exponent bits
(IEEE 754 single-precision floating point number type)

Double-Floats `cl_DF`: sign, 53 mantissa bits, 11 exponent bits
(IEEE 754 double-precision floating point number type)

Long-Floats: `cl_LF` sign, arbitrary number of mantissa bits, 32 (or 64) exponent bits


Implementation Aspects: Implementing `class cl_I : public cl_RA`

Problem: Mathematically, `cl_I` must be a specialization of `cl_RA`. But doesn't a rational number carry more data than an integer? Isn't this anti-OO?



Implementation Aspects: Implementing `class cl_I : public cl_RA`

Problem: Mathematically, `cl_I` must be a specialization of `cl_RA`. But doesn't a rational number carry more data than an integer? Isn't this anti-OO?




Solution: implementation follows the „bridge“ design pattern:

```
1 class cl_number {
2     void* pointer_to_hidden_implementation;
3     // no other data fields
4 };
5 ...
6 class cl_RA : public cl_R {
7     // no new data fields...
8 };
9 class cl_I : public cl_RA {
10    // no new data fields...
11 };
```

Implementation Aspects: Implementing `class cl_I : public cl_RA`

Problem: Mathematically, `cl_I` must be a specialization of `cl_RA`. But doesn't a rational number carry more data than an integer? Isn't this anti-OO?



Solution: implementation follows the „bridge“ design pattern:

```
1 class cl_number {
2     void* pointer_to_hidden_implementation;
3     // no other data fields
4 };
5 ...
6 class cl_RA : public cl_R {
7     // no new data fields...
8 };
9 class cl_I : public cl_RA {
10    // no new data fields...
11 };
```

Consequence: `sizeof(cl_number) = ... = sizeof(cl_I) = 4 (or 8)`

Implementation Aspects: Opportunities of the Bridge Design Pattern

So, a user-accessible object is really just a pointer disguised as a type.

- Intrusive reference counting of heap-allocated memory:
efficient, non-interruptive garbage collection
- Declare x , y , z of type `cl_RA` (i.e. $\in \mathbb{Q}$)
let $x = 3/2$, $y = 1/2$, and $z = x+y$
integrality test can be implemented efficiently. User code:

```
if (instanceof(z, cl_I_ring)) {  
    // will be true, even though typeof(z) is cl_RA!}
```
- Object sharing: $x+0$ returns x without copying it, etc.
- Small integers and short floats are *immediate*, not heap allocated

Implementation Aspects: Fixnums and Bignums



Implementation Aspects: Fixnums and Bignums

Chunks of memory on the free store are always aligned.

Return values of `malloc(3)` and friends are multiples of 4 (or 8).

On a typical 32-bit system, such an address is:

$b_{00} b_{01} b_{02} b_{03} b_{04} b_{05} b_{06} b_{07} b_{08} b_{09} b_{10} b_{11} b_{12} b_{13} b_{14} b_{15} b_{16} b_{17} b_{18} b_{19} b_{20} b_{21} b_{22} b_{23} b_{24} b_{25} b_{26} b_{27} b_{28} b_{29} 0 0$

Implementation Aspects: Fixnums and Bignums

Chunks of memory on the free store are always aligned.

Return values of `malloc(3)` and friends are multiples of 4 (or 8).

On a typical 32-bit system, such an address is:

$b_{00} b_{01} b_{02} b_{03} b_{04} b_{05} b_{06} b_{07} b_{08} b_{09} b_{10} b_{11} b_{12} b_{13} b_{14} b_{15} b_{16} b_{17} b_{18} b_{19} b_{20} b_{21} b_{22} b_{23} b_{24} b_{25} b_{26} b_{27} b_{28} b_{29} 0 0$

That leaves 2 (or 3) unused bits which are always zero.

If any of these bits is non-zero, let's interpret the 2 (or 3) bits as a „*type-tag*“ and the remaining bits b_{00} - b_{29} as immediate data.

Implementation Aspects: Fixnums and Bignums

Chunks of memory on the free store are always aligned.

Return values of `malloc(3)` and friends are multiples of 4 (or 8).

On a typical 32-bit system, such an address is:

$b_{00} b_{01} b_{02} b_{03} b_{04} b_{05} b_{06} b_{07} b_{08} b_{09} b_{10} b_{11} b_{12} b_{13} b_{14} b_{15} b_{16} b_{17} b_{18} b_{19} b_{20} b_{21} b_{22} b_{23} b_{24} b_{25} b_{26} b_{27} b_{28} b_{29} 0 0$

That leaves 2 (or 3) unused bits which are always zero.

If any of these bits is non-zero, let's interpret the 2 (or 3) bits as a „*type-tag*“ and the remaining bits b_{00} - b_{29} as immediate data.

Two examples:

- $b_{00} \dots b_{29}$ represent a signed integer in two's complement notation and constant tags $b_{30} = 0$, $b_{31} = 1$ (immediate integers $-2^{29} \dots 2^{29} - 1$)
- b_{00} represents a sign, $b_{01} \dots b_{08}$ an exponent, $b_{09} \dots b_{24}$ a mantissa, and constant tags $b_{30} = 1$, $b_{31} = 0$ (immediate short float type `c1_SF`)

Implementation Aspects: Fixnums and Bignums

Chunks of memory on the free store are always aligned.

Return values of `malloc(3)` and friends are multiples of 4 (or 8).

On a typical 32-bit system, such an address is:

$b_{00} b_{01} b_{02} b_{03} b_{04} b_{05} b_{06} b_{07} b_{08} b_{09} b_{10} b_{11} b_{12} b_{13} b_{14} b_{15} b_{16} b_{17} b_{18} b_{19} b_{20} b_{21} b_{22} b_{23} b_{24} b_{25} b_{26} b_{27} b_{28} b_{29} 0 0$

That leaves 2 (or 3) unused bits which are always zero.

If any of these bits is non-zero, let's interpret the 2 (or 3) bits as a „*type-tag*“ and the remaining bits b_{00} - b_{29} as immediate data.

Two examples:

- $b_{00} \dots b_{29}$ represent a signed integer in two's complement notation and constant tags $b_{30} = 0$, $b_{31} = 1$ (immediate integers $-2^{29} \dots 2^{29} - 1$)
- b_{00} represents a sign, $b_{01} \dots b_{08}$ an exponent, $b_{09} \dots b_{24}$ a mantissa, and constant tags $b_{30} = 1$, $b_{31} = 0$ (immediate short float type `c1_SF`)

⇒ no heap allocation for small values ⇒ efficiency

all this is completely transparent for the user of the library

Implementation Aspects: Fixnums and Bignums

Chunks of memory on the free store are always aligned.

Return values of `malloc(3)` and friends are multiples of 4 (or 8).

On a typical 32-bit system, such an address is:

$b_{00} b_{01} b_{02} b_{03} b_{04} b_{05} b_{06} b_{07} b_{08} b_{09} b_{10} b_{11} b_{12} b_{13} b_{14} b_{15} b_{16} b_{17} b_{18} b_{19} b_{20} b_{21} b_{22} b_{23} b_{24} b_{25} b_{26} b_{27} b_{28} b_{29} 0 0$

That leaves 2 (or 3) unused bits which are always zero.

If any of these bits is non-zero, let's interpret the 2 (or 3) bits as a „*type-tag*“ and the remaining bits b_{00} - b_{29} as immediate data.

Two examples:

- $b_{00} \dots b_{29}$ represent a signed integer in two's complement notation and constant tags $b_{30} = 0$, $b_{31} = 1$ (immediate integers $-2^{29} \dots 2^{29} - 1$)
- b_{00} represents a sign, $b_{01} \dots b_{08}$ an exponent, $b_{09} \dots b_{24}$ a mantissa, and constant tags $b_{30} = 1$, $b_{31} = 0$ (immediate short float type `c1_SF`)

⇒ no heap allocation for small values ⇒ efficiency

all this is completely transparent for the user of the library

Small CLN Example

For the largest known Mersenne prime p , compute $(2^p - 1)2^{p-1}$.
This is the largest known perfect number:

```
1  #include <iostream>
2  #include <cln/cln.h>
3  using namespace std;
4  using namespace cln;
5
6  int main()
7  {
8      int p = 30402457;
9      cl_I x = ((cl_I(1) << p) - 1) << (p-1);
10     cout << x << endl;
11 }
```

⇒ printing 18304103 decimal digits takes ca. 1 minute

Finally. . .

CLN has been around and stable for a very long time
licensed under GPL and available from <http://www.ginac.de/CLN/>

pre-packaged in  **debian**,  **SUSE** and  **FreeBSD**

current focus is stability and evolution in small steps

Some projects using it:

- GiNaC <http://www.ginac.de/>
symbolic system in C++ for use within C++
- Qalculate! <http://qalculate.sourceforge.net/>
GUI desktop calculator on steroids
- RPN-Calculator-Py <http://sourceforge.net/projects/calcrpnpny/>
reverse polish notation interpreter for use as an interactive calculator in conjunction with the Python interactive interpreter
- gTybalt <http://wwwthep.physik.uni-mainz.de/~stefanw/gtybalt/>
combination of several C++ packages (CLN, GiNaC, NTL) under CERN's Root framework