

QDP++: Low Level Plumbing for Chroma

Bálint Joó (bjoo@jlab.org)

Jefferson Lab, Newport News, VA

given at

HackLatt'06

NeSC, Edinburgh

March 29, 2006

Basic Plumbing: QDP++ Types

- ♦ In QDP++ we try to capture the index structure of our lattice fields:

	Lattice	Spin	Colour	Reality	BaseType
Real	Scalar	Scalar	Scalar	Real	REAL
LatticeColorMatrix	Lattice	Scalar	Matrix(Nc,Nc)	Complex	REAL
LatticePropagator	Lattice	Matrix(Ns,Ns)	Matrix(Nc,Nc)	Complex	REAL
LatticeFermionF	Lattice	Vector(Ns)	Vector(Nc)	Complex	REAL32
DComplex	Scalar	Scalar	Scalar	Complex	REAL64

- ♦ To do this we use C++ templated types:

```
typedef OScalar < PScalar < PScalar< RScalar <REAL> > > > Real;
typedef OLattice< PScalar < PColorMatrix< RComplex<REAL>, Nc> > > LatticeColorMatrix;
typedef OLattice< PSpinMatrix< PColorMatrix< RComplex<REAL>, Nc>, Ns> > LatticePropagator;
```

- ♦ QDP++ and Portable Expression Template Engine:
 - ♦ Provide expressions and recursion through type structure

QDP++ Expressions: The power of PETE

- Use "convenient" mathematical expressions:

```
LatticeFermion x,y,z;  
Real a = Real(1);
```

Wowee! No indices or for loops!

```
gaussian(x);  
gaussian(y);
```

multi1d<T> =
1 dimensional array
of T

```
z = a*x + y;
```

```
int mu, nu;
```

Lattice Wide Shifts
FORWARD =
from $x+mu$

```
multi1d<LatticeColorMatrix> u(Nd);
```

```
Double tmp = sum( real( trace( u[mu]
```

*shift(u[nu],FORWARD,mu)
* adj(shift(u[mu],FORWARD,nu))
*adj(u[nu])
)


Functions

)
);

but array indices are explicit

QDP++ Subsets, Ordered Subsets

- ◆ Allow for partitioning the lattice
- ◆ Can be predefined: eg rb is “red-black” colouring
- ◆ Can be contiguous (`OrderedSubset`) or not (`Subset`)
- ◆ Same subset may be contiguous in one layout and not in another



QDP++ Sets, Subsets and Layouts

- ♦ In expressions, subset index is on the target
 - ♦ $\text{bar}[\text{rb}[1]] = \text{foo}$ (Copy foo's rb[1] subset to bar's)
- ♦ Users can define new sets
- ♦ Layout chosen at compile time with configure switch
 - ♦ Default is CB2 (2 colour checkerboard) layout
- ♦ Layout needs to be initialized on entry to QDP++

```
mult1d<int> nrow( 4 );
nrow[ 0 ]=nrow[ 1 ]=nrow[ 2 ]=4; nrow[ 3 ]=8;
Layout::setLattSize(nrow);
Layout::create();
```

QDP++: Reading and Writing XML

- No “static data binding” in QDP++
- Treat documents as “amorphous” (contain anything)
- Interrogate documents using XPath.
- Root of Path expression is “context node”

root node → <?xml version="1.0" encoding="UTF-8"?>
From root: /foo/bar/fred] → <fred>6</fred>
from /bar: ./fred <jim>7 8 9</jim>
 </bar>
 </foo>

QDP++: Reading and Writing XML

```
XMLReader r("filename");

Double y;
multild<Int> int_array;
multild<Complex> cmp_array;

try {  
    read(r, "/foo/cmp_array", cmp_array);  
    Absolute Path Read  
    XMLReader new_r(r, "/foo/b&ar");  
    New Context  
    read(new_r, "./int_array", int_array);  
    Relative Paths  
    read(new_r, "./double", y);  
}  
catch( const std::string& e) {  
    QDPIO::cerr << "Caught exception: "  
        << e << endl;  
    QDP_abort(1);  
} QDP++ error  
"stream"
```

```
<?xml version="1.0"  
      encoding="UTF-8"?>  
<foo>  
    <cmp_array>  
      <elem>  
        <re>1</re>  
        <im>-2.0</im>  
      </elem>  
      <elem>  
        <re>2</re>  
        <im>3</im>  
      </elem>  
    </cmp_array>  
    <bar>  
      <int_array>2 3 4 5</int_array>  
      <double>1.0e-7</double>  
    </bar>  
</foo>
```

Array markup of non-simple types

Array markup for simple types

QDP++: Reading and Writing XML

```
// Write to file
```

```
XMLFileWriter foo("./out.xml");
push(out, "rootTag");
int x=5;
Real y=Real(2.0e-7);
write(foo, "xTag", x);
write(foo, "yTag", y);
pop(out);
```

Open XML tag group

Write to a file

<?xml version="1.0"?>

<rootTag>

<xTag>5</xTag>

<yTag>2.0e-7</yTag>

</rootTag>

Close tag group

Write to Buffer

```
XMLBufferWriter foo_buf;
push(foo_buf, "rootTag");
int x = 5;
Real y = Real(2.0e-7);
write(foo_buf, "xTag", x);
write(foo_buf, "yTag", y);
pop(foo_buf);
QDPIO::cout << "Buffer contains" << foo_buf.str()
    << endl;
```

write tag with
content

get buffer content

QDP++ and Custom Memory Management

- Occasionally need to allocate/free memory explicitly - e.g. to provide memory to external library.
- Memory may need custom attributes (eg fast)
- Memory may need to be suitably aligned.
- May want to monitor memory usage

```
pointer=QDP::Allocator::theQDPAlocator::Instance().allocate( size,  
          QDP::Allocator::FAST );
```

Namespace

```
QDP::Allocator::theQDPAlocator::Instance()::free(pointer);
```

Get reference to allocator
(see singleton pattern later)

Allocate memory
from desired pool
if possible, with
alignment suitable
to pool

MemoryPoolHint (attribute)

QDP++: moveToFastMemoryHint

- QCDOC/SP inspired construct (copy to CRAM?)
- moves/copies data to/from fast memory (eg EDRAM)
- NOP on machines where there is no fast memory
 - It really is a hint!!!

```
LatticeFermion x;
```

```
moveToFastMemoryHint(x);
```


```
// Do some fast computation here
```

```
revertFromFastMemoryHint(x, true);
```

Accelerate x!
Do not copy contents.
Contents of x lost

Bring x back to slow
memory. Copy contents


Interface to QIO



- QIO works with record oriented LIME files
- File is composed of
 - File XML records
 - Record XML records
 - Record Binary data
- SciDAC mandates checksum records
- ILDG mandates certain records

QIO Interface

- Write with QDPFileWriter
- Must supply user file and user record XML as XMLBufferWriter-s
- Read with QDPFileReader
- User File XML and User Record XML returned in XML Readers
- Checksum/ILDG details checked internally



QIO and ILDG

- ♦ The underlying QIO support layer can handle ILDG format.
- ♦ Specialization of `write()` function and type traits:
 - ♦ `multi1d<LatticeColorMatrix>` always written in ILDG format.
 - ♦ ILDG `dataLFN` is optional argument to the `QDPFileWriter` constructor.
- ♦ `nersc2ildg` program provided in `examples/` directory
- ♦ `lime_contents` and `lime_extract_record` programs from QIO automatically built and installed

QDP++ and Efficiency

- ❖ PETE eliminates Lattice sized temporaries
- ❖ But still there are temporaries at the site level
 - ❖ This really hurts performance
- ❖ Workaround: Cliché Expressions Optimized
 - ❖ eg: $a^*x + y$, $a^*x + P y$, $\text{norm2}()$, $\text{innerProduct}()$ etc.
 - ❖ Optimizations in C, SSE and bagel_qdp library
- ❖ Non optimized expression still slow:
 - ❖ eg: SU(3) matrix * matrix, matrix * vector, a^*x+y+z
- ❖ Next round of optimization to address this issue.

QDP++ Summary

- QDP++ is an abstraction of a data parallel lattice computer with:
 - Data Parallel Expressions for Mathematics (PETE)
 - Sophisticated I/O Facilities (XML, Binary, QIO)
 - Multiple Memory Pool Memory management
- QDP++ actually has documentation (doc directory)
- Solid bedrock for Chroma and other physics apps
- Fast when using the Optimized Cliché Expressions
- Native speed to improve in SciDAC 2.