

Implementing LQCD applications using SciDAC and Chroma

D. Pleiter

DESY

27. November 2006, LAP06, Zeuthen



Agenda

Introduction

Communication

QLA

QDP++

QIO

Chroma

Conclusions and References

Agenda

Introduction

Communication

QLA

QDP++

QIO

Chroma

Conclusions and References

Disclaimer and Credits

Disclaimer: Not the speciality of the cook!

Credits: Talk about work done by others
Here: USQCD collaboration
<http://www.usqcd.org>



Introduction

Challenges for efficient implementations

- ▶ Complex algorithms (see last talks)
- ▶ Complex computer architectures (see next talks)

Nice to face challenges, but: **Time-to-solution may become large!**

Could community efforts help?

- ▶ Interfaces to hide algorithmic details
- ▶ Interfaces to hide machine details (maschine abstraction)

USQCD started to go down this road. **What did they achieve?**

Overview on SciDAC Software

SciDAC software modules (level 1)

- ▶ QCD Message Passing (QMP)
- ▶ QCD Linear Algebra (QLA)

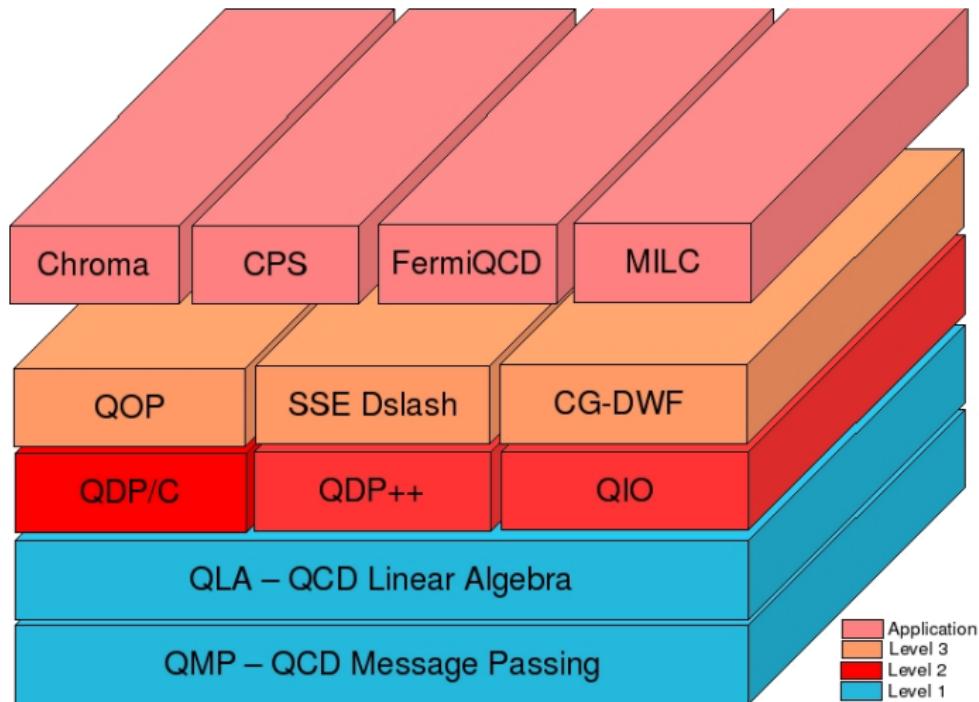
SciDAC software modules (level 2)

- ▶ QCD Data Parallel (QDP)
- ▶ QCD Input Output (QIO)

Applications

- ▶ e.g. Chroma

Overview on SciDAC Software (2)



Agenda

Introduction

Communication

QLA

QDP++

QIO

Chroma

Conclusions and References

MPI: Advantages and Drawbacks

► **Advantages:**

- ▶ Mature and well understood
- ▶ Widely supported formal standard (since 1992)
- ▶ User interface: simple and efficient
- ▶ Available on (almost) all platforms

► **Disadvantages:**

- ▶ Many features are not required for LQCD applications
- ▶ These features may nevertheless affect on performance
- ▶ Possibly difficult to use all hardware features
 - ▶ Architectures may be suitable (BG/L) or less suitable (apeNEXT, QCDOC) for MPI

Design Goals

- ▶ Provide **lean, standard** communication layer **optimized for LQCD**
 - ▶ Simple communication patterns in LQCD
 - ▶ **Homogenous** communication patterns
 - ▶ **Nearest neighbour** communication
 - ▶ **Torus** topology
 - ▶ Collective operations: globus sum, broadcast
 - ▶ Any-to-any communication not required
- ▶ Suitable for very **efficient implementation** on a **variety of interconnects**
- ▶ **Performance requirements:**
 - ▶ Overlap of communication and computation
 - ▶ Multiple outstanding communication requests
 - ▶ Request for send in several directions in one call (QCDOC optimisation)

QMP Functionality

- ▶ Initialisation and finalisation, e.g.
 - ▶ `QMP_init_msg_passing`, `QMP_finalize_msg_passing`
- ▶ Discover machine topology, e.g.
 - ▶ `QMP_get_node_number`
 - ▶ `QMP_get_allocated_number_of_dimensions`
 - ▶ `QMP_get_allocated_dimensions`
- ▶ Logical machine
 - ▶ View of the allocated machine
 - ▶ `QMP_declare_logical_topology(const int *d, int n)`
 - ▶ `QMP_layout_grid (dynamic layout)`
- ▶ Communication
 - ▶ `QMP_declare_send_to, ...`
 - ▶ `QMP_declare_send_relative, ...`
 - ▶ `QMP_declare_multiple`
 - ▶ `QMP_start`, `QMP_is_complete`, `QMP_wait`
- ▶ Collective operations

Example (1)

```
1 int main(int argc, char* argv[])
{
3 ...
4     QMP_status_t status;
5     QMP_thread_level_t req, prv;
7     /* Initialise */
8     req = QMP_THREAD_SINGLE;
9     status = QMP_init_msg_passing(&argc, &argv, req, &prv);
11 ...
13     /* Cleanup */
14     QMP_finalize_msg_passing();
15
16     return 0;
17 }
```

Example (2)

```
1 #define NDIM 4
2
3 int main(int argc, char* argv[])
4 {
5     int l[NDIM] = {4, 4, 4, 8};
6     QMP_status_t status;
7
8     ...
9
10    status = QMP_layout_grid(l, NDIM);
11
12    print_topology();
13    ...
14 }
```

Example (3)

```
1 int main(int argc, char* argv[])
2 {
3     int l[NDIM] = {4, 4, 4, 8};
4     int g[NDIM] = {2, 1, 1, 1};
5     QMP_status_t status;
6
7     ...
8
9     status = QMP_declare_logical_topology(g, NDIM);
10    status = QMP_layout_grid(l, NDIM);
11
12    print_topology();
13
14    ...
15 }
```

Example (4)

```
1 #define NDIM 4
2
3 void print_topology(void) {
4     int nnodes, idim;
5
6     if (QMP_get_node_number() != 0) return;
7
8     nnodes = QMP_get_number_of_nodes();
9     printf("#nodes = %d\n", nnodes);
10    {
11        const int *s = QMP_get_subgrid_dimensions();
12        printf("subgrid =");
13        for (idim = 0; idim < NDIM; idim++)
14            printf(" %d", s[idim]);
15        printf("\n");
16    }
17 }
```

QMP

Review

- ▶ Restriction to suitable subset of MPI
- ▶ Performance issues due to requirements wrt machine architecture

Agenda

Introduction

Communication

QLA

QDP++

QIO

Chroma

Conclusions and References

Introduction

- ▶ Linear algebra operations on single node
- ▶ C binding
- ▶ Data objects:

- ▶ Primitives:

QLA_Real

R real

QLA_Complex

C complex

QLA_ColorMatrix

M $N_c \times N_c$ complex matrix

QLA_DiracFermion

D four-spin, N_c complex matrix

...

- ▶ Array of primitives

Example: Left multiplication by colour matrix

$$\phi_{\mu,i}^a = U_{\nu,i}^{a,b} \psi_{\mu,i}^b$$

```
1 #define N 1000
2 #define xmalloc(t, n) (t *) malloc(sizeof(t) * n)
3
4 int main() {
5     QLA_ColorMatrix    *u;
6     QLA_DiracFermion  *psi, *phi;
7
8     u    = xmalloc(QLA_ColorMatrix, N);
9     psi = xmalloc(QLA_DiracFermion, N);
10    phi = xmalloc(QLA_DiracFermion, N);
11
12    QLA_D_veq_M_times_D(phi, u, psi, N);
13 }
```

Example: Left multiplication by colour matrix (2)

- ▶ Variations: `QLA_T_eqop_M_times_T()`

<code>T =</code>	<code>V</code>	(1-spin, N_c colour spinor)
	<code>H</code>	(2-spin, N_c colour spinor)
	<code>D</code>	(4-spin, N_c colour spinor)
	<code>M</code>	($N_c \times N_c$ complex matrix)
	<code>P</code>	($4N_c \times 4N_c$ complex matrix)
<code>eqop =</code>	<code>[v] eq</code>	(=)
	<code>[v] peq</code>	(+=)
	<code>[v] meq</code>	(-=)
	<code>[v] eqm</code>	(=-)

- ▶ Adjoint colour matrix: `QLA_T_eqop_Ma_times_T()`

Example: Reductions

$$r = \sum |\phi|^2$$

```
1 #define N 1000
2 #define xmalloc(t, n) (t *) malloc(sizeof(t) * n)
3
4 int main() {
5     QLA_Real             r;
6     QLA_DiracFermion   *psi;
7
8     psi = xmalloc(QLA_DiracFermion, N);
9
10    QLA_r_veq_norm2_D(&r, psi, N);
11 }
```

Other Functions

- ▶ Functions with complex arguments (e.g. \sqrt{z})
- ▶ Random numbers (uniform and Gaussian distribution)
- ▶ Spin/colour traces
- ▶ Multiplication with gamma matrices
- ▶ Ternary operations, e.g.
 - ▶ Addition with scalar multiplication $a \cdot b + c$
- ▶ ...

Review

- ▶ Quite exhaustive set of operations
- ▶ Various routines optimised for SSE, SSE2, PPC 440, PPC 440d

Agenda

Introduction

Communication

QLA

QDP++

QIO

Chroma

Conclusions and References

Introduction

- ▶ Data-parallel programming environment
- ▶ Uses QMP for communication
- ▶ C++ binding (alternative: QDP/C)

Data Objects

► Concept: Object types categories

- Basic machine types (e.g. int, float, double)
- Real world types (e.g. complex)
- “Primitive types” (e.g. scalar, vector, matrix)
- Grid type

Construct object types using C++ templates

```
typedef OLattice<  
    PScalar<PColorMatrix< RComplex<float>, Nc>>>  
    LatticeColorMatrix
```

- Grid type: `OLattice`
- Primitive types: `PScalar`, `PColorMatrix`
- Real world types: `Rcomplex`
- Basic machine types: `float`
- (Constants: `Nc`)

Data Objects (2)

► Examples for available objects:

ColorMatrix

$N_c \times N_c$ complex matrix

LatticeColorMatrix

array of colour matrices

LatticeFermion

N_s spin N_c colour spinor

Data Layout

- ▶ **Information relevant for controlling data layout:**

- ▶ Number of dimensions
- ▶ Lattice size
- ▶ Number of threads

- ▶ **Functions to control this:**

- ▶ `Layout::setLattSize(const mult1d<int>& size)`
- ▶ `Layout::setSMPFlag({true|false})`
- ▶ `Layout::setNumProc(n)`

- ▶ **Get layout dependent information:**

- ▶ `Layout::sitesOnNode()`
- ▶ `Layout::nodeNumber(x)`

Operations (1)

- ▶ Typical operation in QCD applications:

$$\psi_{\alpha i} = U_{ij}(\gamma_0)_{\alpha\beta}\phi_{\beta j}$$

- ▶ Implementation in QDP++:

```
1 LatticeFermion      psi
  LatticeFermion      phi
 3 LatticeColorMatrix u
5 psi = u * Gamma(1) * phi
```

- ▶ Operation may involve Hermitian conjugate, e.g. U_{ij}^\dagger :

```
1 psi = adj(u) * Gamma(1) * phi
```

Operations (2)

- ▶ Shift operation: move data to neighbouring lattice site
 - ▶ May include communication
- ▶ Example: calculation of staples:

$$U_\nu(x + \mu) * U_\mu(x + nu)^\dagger * U_\nu(x)^\dagger$$

- ▶ Implementation in QDP++:

```
1 #define NDIM 4

3 int mu, nu;
4 multi1d<LatticeColorMatrix> u(NDIM);
5 LatticeColorMatrix s;

7 s = shift(u[nu], FORWARD, mu) *
8     adj(shift(u[mu], FORWARD, nu)) *
9     adj(u[nu]);
```

Example

```
1 int main(int argc, char *argv[]) {
2     QDP_initialize(&argc, &argv);
3
4     const int xl[] = {4,4,4,4};
5     QDP::multi1d<int> l(QDP::Nd);
6     l = xl;
7     QDP::Layout::setLattSize(l);
8     QDP::Layout::create();
9
10    QDP::multi1d<LatticeColorMatrix> u(Nd);
11
12    ...
13
14    QDP::QDP_finalize();
15
16    return 0;
17 }
```

Agenda

Introduction

Communication

QLA

QDP++

QIO

Chroma

Conclusions and References

Introduction

- ▶ Input/Output Applications Programmer Interface
- ▶ Different file formats supported:
 - ▶ “Single file format”
 - ▶ I/O through single master node
 - ▶ “Partition file format”
 - ▶ I/O through a set of I/O nodes
 - ▶ Serve all nodes of corresponding partition
 - ▶ “Multi file format”
 - ▶ Write of temporary files to local disks
- ▶ QMP is used for data transport
- ▶ Support for structured file format: **LIME**

LIME

- ▶ LIME = **Lattice QCD Interchange Message Encapsulation**
- ▶ LIME file consists of one or more **messages**
- ▶ Each message consists of one or more **records**
- ▶ A record contains ASCII or binary data
- ▶ 144-byte header:
 - ▶ control word (8 byte)
 - ▶ record length (8 byte)
 - ▶ type string (128 byte)

Example

```
1 #include <stdio.h>
2 #include "lime.h"
3
4 main( int argc , char* argv[] )
5 {
6     FILE             *fp ;
7     LimeReader      *r ;
8     unsigned char   *buf ;
9
10    fp = fopen(argv[1] , "r" );
11    r = limeCreateReader(fp );
12
13    limeReaderNextRecord(r );
14    printf("First header = %s\n" , limeReaderType(r ) );
15 }
```

LIME Interface Functionality

- ▶ Read operations (selection)
 - ▶ Create reader: limeCreateReader, limeDestroyReader
 - ▶ Header accessors: limeReaderType, limeReaderBytes
 - ▶ Pay-load read: limeReaderReadData
 - ▶ Skip to next record: limeReaderNextRecord
 - ▶ Seek within record: limeReaderSeek
- ▶ Write operations (selection)
 - ▶ Create writer: limeCreateWriter, limeDestroyWriter
 - ▶ Write operation: limeWriteRecordHeader,
limeWriteRecordData, limeWriterCloseRecord
- ▶ Header creation

Review

- ▶ Efficient read operations possible
- ▶ Support for >2 GBytes file size
- ▶ Standard adopted by ILDG

Agenda

Introduction

Communication

QLA

QDP++

QIO

Chroma

Conclusions and References

Introduction

- ▶ Application level software package
- ▶ Written in C++
- ▶ Dependencies: QDP++, QMP, libXML
- ▶ Optimised for various platforms (PC-Clusters, QCDOC, BGL)
- ▶ Applications
 - ▶ Pure gauge heatbath
 - ▶ HMC
 - ▶ Measurement

Chroma Modules

- ▶ **I/O module:**
 - ▶ Support for different gauge configuration formats (incl. ILDG)
- ▶ **Measurements module:**
 - ▶ Generation of colour sources
 - ▶ Propagator smearing
 - ▶ Calculation of correlation functions
- ▶ **Inverters:** CG, BiCGStab, SUMR, MINRES, ...
- ▶ **Eigenvalues**
- ▶ **Various fermion and gauge actions**

Step 1: Generate SU(3) Configuration

- ▶ Create driver/parameter XML file, say `su3-in.xml`
- ▶ Execute application:

```
purgaug -i su3-in.xml -o su3-out.xml
```

XML Driver File (1)

Start with cold configuration:

```
1<?xml version="1.0"?>
2<purgaug>
3  <Cfg>
4    <cfg_type>UNIT</cfg_type>
5    <cfg_file></cfg_file>
6  </Cfg>
7  ...
8</purgaug>
```

XML Driver File (2)

Define update parameters:

```
<MCControl>
2   <RNG> ... </RNG>

4   <StartUpdateNum>0</StartUpdateNum>
5   <NWarmUpUpdates>2</NWarmUpUpdates>
6   <NProductionUpdates>1000</NProductionUpdates>
7   <NUpdatesThisRun>10</NUpdatesThisRun>
8   <SaveInterval>10</SaveInterval>
9   <SavePrefix>result</SavePrefix>
10 ...
</MCControl>
```

XML Driver File (3)

Define physical parameters

```
1<HBIter>
2  <GaugeAction>
3    <Name>WILSON_GAUGEACT</Name>
4    <beta>6.0</beta>
5    <GaugeState>
6      <Name>SIMPLE_GAUGE_STATE</Name>
7      <GaugeBC>
8        <Name>PERIODIC_GAUGEBC</Name>
9        </GaugeBC>
10      </GaugeState>
11    </GaugeAction>
12    ...
13  <nrow>4 4 4 8</nrow>
14</HBIter>
```

XML Driver File (4)

Define measurements

```
<InlineMeasurements>
2   <elem>
4     <Name>POLYAKOV_LOOP</Name>
4     <Frequency>1</Frequency>
5     <Param>
6       <version>2</version>
7     </Param>
8     <NamedObject>
9       <gauge_id>default_gauge_field</gauge_id>
10    </NamedObject>
11  </elem>
12 </InlineMeasurements>
```

Step 2: Measurement of Hadron Observables

- ▶ Typical measurement application tasks:
 - ▶ Read a gauge configuration
 - ▶ Generate a colour source
 - ▶ Compute quark propagator (“invert” fermion matrix)
 - ▶ Calculate correlation functions
- ▶ Similar procedure as before
 - ▶ Create driver/parameter file, say `hadron-in.xml`
 - ▶ Execute application:

```
chroma -i hadron-in.xml -o hadron-out.xml
```

XML Driver File (1)

Read configuration:

```
<?xml version="1.0"?>
2 <chroma>
  <Cfg>
4    <cfg_type>SCIDAC</cfg_type>
      <cfg_file>su3.lime1</cfg_file>
6    </Cfg>
...
8 </chroma>
```

XML Driver File (2)

Define tasks:

```
2   <Param>
3     <InlineMeasurements>
4       <elem>
5         <Name> . . . </Name>
6         <Frequency> . . . </Frequency>
7         <Param>
8           ...
9         </Param>
10        <NamedObject>
11          <gauge_id>default_gauge_field</gauge_id>
12          <source_id>src</source_id>
13        </NamedObject>
14        ...
15      ...
```

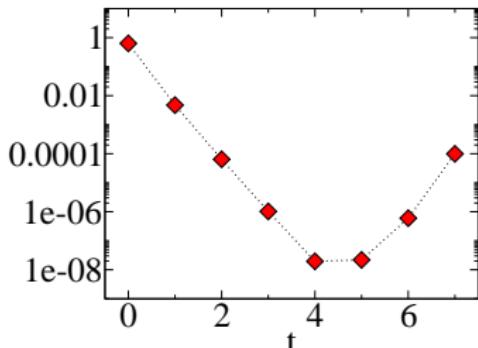
XML Driver File (3)

Calculate propagator:

```
1   <elem>
2     <Name>PROPAGATOR</Name> ...
3     <Param> ...
4       <FermionAction>
5         <FermAct>WILSON</FermAct>
6         <Kappa>0.11</Kappa> ...
7       </FermionAction>
8       <InvertParam>...</InvertParam>
9     </Param>
10    <NamedObject>
11      <gauge_id>default_gauge_field</gauge_id>
12      <source_id>src</source_id>
13      <prop_id>prop</prop_id>
14    </NamedObject>
15  </elem>
```

Results

Nucleon correlator:



Review

- ▶ Some very nice concepts
- ▶ Toolbox for quick experiments
- ▶ Complex software, difficult to adapt to own needs
- ▶ Compiler dependencies
- ▶ Lack of documentation
- ▶ Tutorial: <http://www.ph.ed.ac.uk/~bj/HackLatt06>

Agenda

Introduction

Communication

QLA

QDP++

QIO

Chroma

Conclusions and References

Conclusions

► **Advantages of community codes:**

- ▶ Avoid re-implementation of standard tasks
- ▶ Share porting and optimisation efforts
 - ▶ Highly optimised assembly code
- ▶ Common lower level routines and data formats
 - ▶ Simpler to share code and data

► **Disadvantages of community codes:**

- ▶ A lot of effort required to setup such an infrastructure
- ▶ Some efforts may be needed to obtain good performance
 - ▶ Human resources vs. machine resources

► **SciDAC example:**

- ▶ Very useful software infrastructure → Basis for applications software of most US-collaborations
- ▶ Successful efforts to keep software
 - ▶ portable
 - ▶ (relatively) easy to install
- ▶ Application level software more difficult to re-use

References

- ▶ Software installed in
`/afs/ifh.de/group/ape/lap06/scidac`
- ▶ Copy of sources in
`/afs/ifh.de/group/ape/lap06/scidac/src`
- ▶ For documentation see
 - ▶ <http://www.usqcd.org/usqcd-software>

Hints for Building Software

- ▶ Check file `INSTALL`
- ▶ Execute script `configure`
- ▶ When lucky (typically true): execute `make`
- ▶ Most software modules include examples