QDP++ Data Parallel Interface for QCD

Version 1.0.2

Robert G. Edwards
for the SciDAC Software Coordinating Committee

March 7, 2003

1 Introduction

This is a user's guide for the C++ binding for the QDP Data Parallel Applications Programmer Interface developed under the auspices of the U.S. Department of Energy Scientific Discovery through Advanced Computing (SciDAC) program.

The QDP Level 2 API has the following features:

- Provides data parallel operations (logically SIMD) on all sites across the lattice or subsets of these sites.
- Operates on lattice objects, which have an implementation-dependent data layout that is not visible above this API.
- Hides details of how the implementation maps onto a given architecture, namely how the logical problem grid (i.e. lattice) is mapped onto the machine architecture.
- Allows asynchronous (non-blocking) shifts of lattice level objects over any permutation map of sites onto sites. However, from the user's view these instructions appear blocking and in fact may be so in some implementation.
- Provides broadcast operations (filling a lattice quantity from a scalar value(s)), global reduction operations, and lattice-wide operations on various data-type primitives, such as matrices, vectors, and tensor products of matrices (propagators).
- Operator syntax that support complex expression constructions.

2 Datatypes

The N_d dimensional lattice consists of all the space-time sites in the problem space. Lattice data are fields on these sites. A data primitive describes data on a single

site. The lattice fields consist of the primitives over all sites. We do not define data types restricted to a subset of the lattice — rather, lattice fields occupy the entire lattice. The primitive types at each site are represented as the (tensor) product space of, for example, a vector space over color components with a vector space over spin components and complex valued elements.

2.1 Type Structure

Generically objects transform under different spaces with a tensor product structure as shown below:

	<i>Lattice</i>		<i>Color</i>		<i>Spin</i>		<i>Complexity</i>
Gauge fields :	Lattice	⊗	Matrix(Nc)	⊗	Scalar	⊗	Complex
Fermions :	Lattice	⊗	Vector(Nc)	⊗	Vector(Ns)	⊗	Complex
Scalars :	Scalar	⊗	Scalar	⊗	Scalar	⊗	Scalar
Propagators :	Lattice	⊗	Matrix(Nc)	⊗	Matrix(Ns)	⊗	Complex
Gamma :	Scalar	⊗	Scalar	⊗	Matrix(Ns)	⊗	Complex

Nd is the number of space-time dimensions

Nc is the dimension of the color vector space

Ns is the dimension of the spin vector space

Gauge fields can left-multiply fermions via color matrix times color vector but is diagonal in spin space (spin scalar times spin vector). A gamma matrix can right-multiply a propagator (spin matrix times spin matrix) but is diagonal in color space (color matrix times color scalar).

Types in the QDP interface are parameterized by a variety of types including:

- *Word type*: int, float, double, bool. Basic machine types.
- *Reality type*: complex or scalar. This is where the idea of a complex number lives.
- *Primitive type*: scalar, vector, matrix, etc. This is where the concept of a gauge or spin field lives. There can be many more types here.
- *Inner grid type*: scalar or lattice. Supports vector style architectures.
- *Outer grid type*: scalar or lattice. Supports super-scalar style architectures. In combination with Inner grid can support a mixed mode like a super-scalar architecture with short length vector instructions.

There are template classes for each of the type variants listed above. The interface relies heavily on templates for composition - there is very little inheritance. The basic objects are constructed (at the users choice) by compositions like the following:

```
typedef OLattice<PScalar<PColorMatrix<RComplex<float>, Nc> > > LatticeColorMatrix
typedef OLattice<PSpinVector<PColorVector<RComplex<float>, Nc>, Ns> > LatticeFermion
```

The classes PScalar, PSpinVector, PColorMatrix, PColorVector are all subtypes of a primitive type. The relative ordering of the classes is important. It is simply a user convention that spin is used as the second index (second level of type composition) and color is the third. The ordering of types can be changed. From looking at the types one can immediately decide what operations among objects makes sense.

2.2 Generic Names

The linear algebra portion of the QDP API is designed to resemble the functionality that is available in the Level 1 QLA API and the C Level QDP API. Thus the datatypes and function naming conventions are similar. Predefined names for some generic lattice field datatypes are listed in the table below. Because the API is based heavily on templates, the possible types allowed is much larger than listed below.

name	description
LatticeReal	real
LatticeComplex	complex
LatticeInt	integer
LatticeColorMatrix	$N_c \times N_c$ complex matrix
LatticeHalfFermion	two-spin, N_c color spinor
LatticeDiracFermion	four-spin, N_c color spinor
LatticeStaggeredFermion	one-spin, N_c color spinor
LatticeDiracPropagator	$4N_c \times 4N_c$ complex matrix
LatticeStaggeredPropagator	$N_c \times N_c$ complex matrix
LatticeSeed	implementation dependent

Single site (lattice wide constant fields) versions of types exist without the Lattice prepended. All types and operations defined for QDP live within a C++ namespace called QDP thus ensuring no type conflicts with other namespaces.

2.3 Specific Types for Color and Precision

According to the chosen color and precision, names for specific floating point types are constructed from names for generic types. Thus LatticeColorMatrix becomes LatticeColorMatrixPC, where the precision P is D or F according to the table below

abbreviation	description
D	double precision
F	single precision

and C is 2, 3, or some arbitrary N , if color is a consideration. Note, the value of N is an arbitrary compile time constant.

If the datatype carries no color, the color label is omitted. Also, if the number of color components is the same as the compile time constant, then the color label can be omitted. Integers also have no precision label. The specification of precision and number of colors is not needed for functions because of overloading.

For example, the type

LatticeDiracFermionF3

describes a lattice quantity of single-precision four-spin, three-color spinor field.

2.4 Color and Precision Uniformity

The only place that the number of color or spin components occur is through instance of the global constant variables N_c , and N_s . These are only directly used in the typedef constructions of user defined types. Nothing restricts a user from constructing types for other number of colors. In fact, the use of N_c in the construction of user defined types is simply a convenience for the user, and as such a user can use any integer that is reasonable. The API merely requires that the types used in operations are conforming.

However, in standard coding practice it is assumed that a user keeps one of the precision, color, and spin options in force throughout the compilation. So as a rule all functions in the interface take operands of the same precision, color, and number of spin components. As with data type names, function names come in generic color-, spin- and precision-specific forms, as described in the next section. Exceptions to this rule are functions that explicitly convert from double to single precision and vice versa. If the user choose to adopt color and precision uniformity, then all variables can be defined with generic types and all functions accessed through generic names. The prevailing color is defined through the compile time constant N_c . The interface automatically translates data type names and function names to the appropriate specific type names through typedefs. With such a scheme and careful coding, changing only the compile time N_c and the QDP library converts code from one color and precision choice to another.

2.5 Breaking Color and Precision Uniformity

It is permissible for a user to mix precision and color choices. This is done by declaring variables with specific type names, using functions with specific names, and making appropriate precision conversions when needed.

3 QDP Functions

The QDP functions are grouped into the following categories:

1. Entry and exit from QDP
2. Layout utilities
3. Data parallel functions
4. Data management utilities
5. Subset definition
6. Shift creation
7. I/O utilities
8. Temporary exit and reentry

3.1 Entry and exit from QDP

QDP must be initialized before any other routine can be used. The initialization is broken into two steps – initializing the underlying hardware and initializing the layout.

Initialization of QDP

Prototype	<code>void QDP_initialize(int *argc, char ***argv)</code>
Purpose	Places the hardware into a known state.
Example	<code>QDP_initialize();</code>

This routine will be responsible for initializing any hardware like the physical layer of the message passing system. For compatibility with QMP, the addresses of the main programs `argc` and `argv` must be passed. They may be modified.

Shutdown of QDP

Prototype	<code>void QDP_finalize()</code>
Purpose	Shutdown QDP.
Example	<code>QDP_finalize();</code>

This call provides for an orderly shutdown of QDP. It is called by all nodes. It concludes all communications, does housekeeping, if needed and performs a barrier wait for all nodes. Then it returns control to the calling process.

Panic exit from QDP

Prototype	<code>void QDP_abort(int status)</code>
Purpose	Panic shutdown of the process.
Example	<code>QDP_abort(1);</code>

This routine may be called by one or more nodes. It sends kill signals to all nodes and exits with exit status `status`.

Entry into QDP

Prototype	<code>void Layout::create()</code>
Purpose	Starts QDP with layout parameters in <code>Layout</code> .
Example	<code>Layout::create();</code>

The routine `Layout::create()` is called once by all nodes and starts QDP operations. It calls the layout routine with the parameters set in the namespace `Layout` specifying the layout. The layout is discussed in Section 3.2.

This step is separated from the `QDP_initialize()` above so layout parameters can be read and broadcasted to the nodes. Otherwise the layout parameters have to be set from the environment or fixed in the compilation.

Exit from QDP

Prototype	<code>void Layout::destroy()</code>
Purpose	Exits QDP.
Example	<code>Layout::destroy();</code>

This call provides for an orderly exit from QDP. It is called by all nodes. It concludes all communications, does housekeeping, if needed and performs a barrier wait for all nodes. The communication layer is not finalized.

3.2 Layout utilities

Routines for constructing the layout are collected in the namespace `Layout`. The *setter* and *getter* routines provide a way to set parameters like the lattice size.

The layout creation function determines which nodes get which lattice sites and in what linear order the sites are stored. The `Layout` namespace has entry points that allow a user to inquire about the lattice layout to facilitate accessing single site data from a QDP lattice field. For code written entirely with other QDP calls, these routines may be ignored by the user, with the exception of the useful routine `latticeCoordinate`. However, if a user removes data from a QDP lattice object (see `expose` or `extract`) and wishes to manipulate the data on a site-by-site basis, the global entry points provided here are needed to locate the site data.

Some implementations may have a built-in tightly constrained layout. In flexible implementations there may be several layout choices, thereby allowing the user the freedom to select one that works best with a given application. Furthermore, such

implementations may allow the user to create a custom layout to replace one of the standard layouts. As long as the custom layout procedure provides the entry points and functionality described here, compatibility with the remainder of the QDP library is assured.

3.2.1 QDP setup

Layout creation

The layout creation routine `Layout::create()` defined in Section 3.1 generates pre-defined lattice subsets for specifying even, odd, and global subsets of the lattice:

Subset even, odd, all

It also creates the nearest-neighbor shifts for each coordinate direction.

Defining the layout

There are set/accessor functions to specify the lattice geometry used in the layout. Generically, the accessors have the form:

Generic	<code>void Layout::set<something>(<param>)</code>
Purpose	Set one of the site data layout configurations.
Example	<code>Layout::setLattSize(size);</code>

The type of input information needed by the layout is as follows:

1. Number of dimensions N_d . Must be the compile time dimensions.
2. Lattice size (e.g., $L_0, L_1, \dots, L_{N_d-1}$)
3. SMP flag

These parameters are accessed and set with the following functions:

Generic	<code>void Layout::setLattSize(const multiid<int>& size)</code>
Purpose	Set the lattice size for the data layout.
Default	No default value. Must always be set.
Example	<code>Layout::setLattSize(size);</code>

Generic	<code>void Layout::setSMPFlag(bool)</code>
Purpose	Turn on using multi-processor/threading
Default	Default value is false - single thread of execution.
Example	<code>Layout::setSMPFlag(true);</code>

Generic	<code>void Layout::setNumProc(int N)</code>
Purpose	In a multi-threaded implementation, use N processors.
Default	Default value is 1 - single thread of execution.
Example	<code>Layout::setNumProc(2);</code>

3.2.2 Generic layout information

The following global entry points are provided in the Layout namespace. They provide generic user information.

Returning the spacetime coordinates

Prototype	<code>LatticeInt Layout::latticeCoordinate(int d)</code>
Purpose	The dth spacetime coordinate.
Example	<code>LatticeInt coord = Layout::latticeCoordinate(2);</code>

The call `Layout::latticeCoordinate(d)` returns an integer lattice field with a value on each site equal to the integer value of the dth space-time coordinate on that site.

Lattice volume

Prototype	<code>int vol()</code>
Purpose	Return the total lattice volume
Example	<code>int vol = Layout::vol();</code>

3.2.3 Entry points specific to the layout

The additional global entry points are provided in the Layout namespace. They reveal some information specific to the implementation.

Node number of site

Prototype	<code>int Layout::nodeNumber(const multiid<int>& x)</code>
Purpose	Returns logical node number containing site x.
Example	<code>node = Layout::nodeNumber(x);</code>

Linear index of site

Prototype	<code>int Layout::linearIndex(const multiid<int>& x)</code>
Purpose	Returns the linearized index for the lattice site x.
Example	<code>int k = Layout::linearIndex(x);</code>

Map node and linear index to coordinate

Prototype	<code>multiid<int> Layout::siteCoords(int node, int index)</code>
Purpose	Returns site coordinate x for the given node and linear index index.
Example	<code>multiid<int> lc = Layout::siteCoords(n, i);</code>

Number of sites on a node

Prototype	<code>int Layout::sitesOnNode()</code>
Purpose	Returns number of sites assigned to a node.
Example	<code>int num = Layout::sitesOnNode();</code>

The linear index returned by `Layout::linearIndex()` ranges from 0 to `Layout::sitesOnNode() - 1`.

3.3 Data Parallel Functions

Data parallel functions are described in detail in Sec. 7. In the C++ API, there are overloaded functions that can be applied to site or lattice wide objects. Arbitrarily complicated expressions can be built from these functions. The design of the API describes that all operations are to be performed site-wise. The only connection between sites is via a map or shift function.

The class of operations are generically described by site-wise operations (the “linear algebra” part of the API), and shift (or map) versions. The latter generically involves communications among processors in a parallel implementation.

The operator style provided by the API thus allows operations like the following:

```
Lattice A, B;
LatticeColorMatrix U;
B = U * A;
```

From the type declarations

```
typedef OLattice<PScalar<PColorMatrix<RComplex<float>, Nc> > > LatticeColorMatrix
typedef OLattice<PSpinVector<PColorVector<RComplex<float>, Nc>, Ns> > LatticeFermion
```

one can see a `OLattice` multiplies a `OLattice`. At each site, the `U` field is a scalar in spin space, thus a `PScalar` multiplies a `PSpinVector` - a vector in spin space. For each spin component, there is a `PColorMatrix` multiplying a `PColorVector`. The multiplications involve complex numbers.

Thus we see that mathematically the expression carries out the product

$$B_{\alpha}^i(x) = U^{ij}(x) * A_{\alpha}^j(x)$$

for all lattice coordinates x belonging to the subset `all`. Here A and B are objects of lattice Dirac fermion fields and U is an object of type lattice gauge field. The superscripts i, j refer to the color indices and the subscript α refers to the spin index. For each spin and color component, the multiplication is over complex types.

This tensor product factorization of types allows for potentially a huge variety of mathematical objects. The operations between the objects is determined by their tensor product structure.

The API allows for operations to be narrowed to a subset of sites. The infix notation does not allow for extra arguments to be passed to an operation, so the subset is fixed via the target. The API mandates that there is in use in even a complex operation, namely the target specifies the subset to use. To narrow an operation to a specific subset, one specifies the subset in the target as follows:

```
chi[even] = u * psi;
```

which will store the result of the multiplication on only the *even* subset.

The C++ API differs from the C API significantly in the name of functions. In C++, there is no need for naming conventions for the functions since one can overload the function name on the types of its arguments. More significantly, the C API uses a functional style where the destination of an operation is part of the arguments for an operation, and all functions return void. The C++ API uses an operator/infix style allowing complex expressions to be built.

3.3.1 Constant Arguments

In some cases it is desirable to keep an argument constant over the entire subset. For example the function

```
Complex z;  
LatticeFermion c, b;  
c[s] = z * b;
```

multiplies a lattice field of color vectors by a complex constant as in

```
c[x] = z*b[x]
```

for x in subset s .

3.3.2 Functions

In the C++ API all operations are functions that act on their argument and most functions return their results. Except for explicit shift functions and global reductions, these functions are point-wise. The C++ API differs from the C API in that there are no combined operations like adjoint with a multiply. Instead, one simply calls the adjoint function. Thus

```
c = adj(u)*b
```

carries out the product

```
c[x] = adj(u[x])*b[x]
```

for all sites x in subset all .

3.3.3 Gamma matrices

Multiplication of spin vectors and matrices by γ -matrices is provided.

```
LatticeDiracFermion c, b;  
int n;  
c = Gamma(n) * b;
```

Where the multiplication is spin matrix times a spin vector. Right multiplication is also supported.

```
LatticePropagator q, r;  
int n;  
q = r * Gamma(n);
```

The `Gamma(n)` provides an enumeration of all possible γ -matrix combinations. See the Section 5 for more details.

3.3.4 Shift

A shift function is a type of map that maps sites from one lattice site to another. In general, maps can be permutation maps but there are nearest neighbor shift functions provided by default. See the discussion of shifts below in Section 3.8. Thus

```
c[s] = shift(b,sign,dir)
```

shifts an object along the direction specified by `dir` and `sign` for all sites `x` in destination subset `s`.

3.4 Creating and destroying lattice fields

The declaration of an object of some type say `LatticeReal` will call a constructor. The implementation guarantees the object is fully created and all memory needed for it is allocated. Thus, there is no need for the user to use `new` to create an object. The use of pointers is discouraged. When an object goes out of scope, a destructor is called which will guarantee all memory associated with the object is released.

There is no aliasing or referencing of two objects with the same internal data storage. Each object a user can construct has its own unique storage.

3.5 Array container objects

For convenience, the API provides array container classes with much limited facility compared to the Standard Template Library. In particular, one, two, three, and four dimensional array container classes are available. The benefit of two and higher dimension classes is that they can be allocated after they are declared. This is in contrast to the STL technique, which builds multi-dimensional arrays out of nested one-dimensional array, and one must allocate a nested array of array classes by looping over the individual elements allocating each one.

An array of container classes is constructed as follows:

```
multi1d<LatticeComplex> r(Nd); // a 1-D array of LatticeComplex
multi2d<Real> foo(2,3);        // a 2-D array of Real with first index slowest
```

3.6 Function objects

Function objects are used in the constructions of Sets/subsets and maps/shifts. The objects created by maps are themselves function objects. They serve the role as functions, but because of their class structure can also carry state.

A function object has a struct/class declaration. The key part is the function call operator. A generic declaration is something like:

```
struct MyFunction
{
    MyFunction(int dir) : mu(dir) {}
    Real operator()(const int& x)
```

```

    {\* operates on x using state held in mu and returns a Real *\}

    int mu;
}

```

A user can then use an object of type MyFunction like a function:

```

MyFunction foo(37); // hold 37 within foo
int x;
Real boo = foo(x); // applies foo via operator()

```

3.7 Subsets

It is sometimes convenient to partition the lattice into multiple disjoint subsets (e.g. time slices or checkerboards). Such subsets are defined through a user-supplied function that returns a range of integers $0, 1, 2, \dots, n-1$, so that if $f(x) = i$, then site x is in partition i . A single subset may also be defined by limiting the range of return values to a single value (i.e. 0). This procedure may be called more than once, and sites may be assigned to more than one subset. Thus, for example an even site may also be assigned to a time slice subset and one of the subsets in a 32-level checkerboard scheme. A subset definition remains valid until its destructor is called.

Defining a set Subsets are first defined through the construction of an object of type Set using a function object. This function object is a derived type of SetFunc. Function objects are described in Section 3.6. Subsets are defined through the data type Subset.

Prototype	<pre> Set::make(const SetFunc& func) int SetFunc::operator()(const multild<int>& x) int SetFunc::numSubsets() </pre>
Purpose	Creates a Set that holds numSubsets subsets based on func.
Requirements	<p>The func is a derived type of SetFunc and maps lattice coordinates to a partition number.</p> <p>The function in func.numSubsets() returns number of partitions.</p>
Example	<pre> Set timeslice; class timesliceFunc : public SetFunc; timeslice.make(timesliceFunc); </pre>

Here is an explicit example for a timeslice:

```

struct TimeSliceFunc : public SetFunc
{
    TimeSliceFunc(int dir): mu(dir) {}

    // Simply return the mu'th coordinate

```

```

int operator()(const multild<int>& coord)
{return coord[mu];}

// The number of subsets is the length of the lattice
// in direction mu
int numSubsets() {return Layout::lattSize()[mu];}

int mu; // state
}

Set timeslice;
timeslice.make(TimeSliceFunc(3)) // makes timeslice in direction 3

```

It is permissible to call `Set.make()` with a function object having only 1 subset. In this case the partition function must return zero if the site is in the subset and nonzero if not. (Note, this is opposite the “true”, “false” convention in C).

Extracting a subset A subset is returned from indexing a `Set` object.

Prototype	Subset Set::operator[] (int i)
Purpose	Returns the i-th subset from a Set object.
Example	Set timeslice; Subset origin = timeslice[0];

The `Set::make()` functions allocates all memory associated with a `Set`. A `Subset` holds a reference info to the original `Set`. A destructor call on a `Set` frees all memory.

Using a subset A subset can be used in an assignment to restrict sites involved in a computation:

```

LatticeComplex r, a, b;
Subset s;
r[s] = 17 * a * b;

```

will multiply $17 * a * b$ onto `r` only on sites in the subset `s`.

3.8 Maps and shifts

Shifts are general communication operations specified by any permutation of sites. Nearest neighbor shifts are a special case. Thus, for example,

```

LatticeHalfFermion a, r;
r[s] = shift(a,sign,dir);

```

shifts the half fermion field `a` along direction `dir`, forward or backward according to `sign`, placing the result in the field `r`. Nearest neighbor shifts are specified by values of `dir` in the range $[0, N_d - 1]$. The sign is +1 for shifts from the positive direction,

and -1 for shifts from the negative direction. That is, for $\text{sign} = +1$ and $\text{dir} = \mu$, $r(x) = a(x + \mu)$. For more general permutations, dir is missing and sign specifies the permutation or its inverse.

The subset restriction applies to the destination field r . Thus a nearest neighbor shift operation specifying the even subset shifts odd site values from the source a and places them on even site values on the destination field r .

Creating shifts for arbitrary permutations The user must first create a function object for use in the map creation as described in Section 3.6. Thus to use the make a map one uses a function object in the map creation:

Prototype	<code>Map::make(const MapFunc& func)</code>
Purpose	Creates a map specified by the permutation map function object <code>func</code> .
Requirements	The <code>func</code> is a derived type of <code>MapFunc</code> and must have a <code>multild<int> operator()(const multild<int>& d)</code> member function that maps a source site to d .
Result	Creates an object of type <code>map</code> which has a function call <code>template<class T> T Map::operator()(const T& a)</code>
Example	<pre>Map naik; LatticeReal r,a; r = naik(a);</pre>

The coordinate map function object `func` above that is handed to the map creation function `Map::make()` maps lattice coordinates of the the destination to the source lattice coordinates. After construction, the function object of type `Map` can be used like any function via the `operator()`. It can be applied to all QDP objects in an expression.

The function object has an operator that given a coordinate will return the source site coordinates. An example is as follows:

```
struct naikfunc : public MapFunc
{
    naik(int dir) : mu(dir) {}
    multild<int> operator()(const multild<int>& x)
        {\/* maps x to x + 3*mu where mu is direction vector *\}

    int mu;
}
```

For convenience, there are predefined `Map` functions named `shift` that can shift by 1 unit backwards or forwards in any lattice direction. They have the form

```
shift(const QDPType& source, int sign, int dir);
```

The construction of a `Map` object allocates all the necessary memory needed for a shift. Similarly, a destructor call on a `Map` object frees memory.

3.9 I/O utilities

[Under development.]

3.10 Temporary entry and exit from QDP

For a variety of reasons it may be necessary to remove data from QDP structures. Conversely, it may be necessary to reinsert data into QDP structures. For example, a highly optimized linear solver may operate outside QDP. The operands would need to be extracted from QDP fields and the eventual solution reinserted. It may also be useful to suspend QDP communications temporarily to gain separate access to the communications layer. For this purpose function calls are provided to put the QDP implementation and/or QDP objects into a known state, extract values, and reinsert them.

Exposing QDP data

Prototype	<i>QLA_Type</i> * QDP_expose(<i>Type</i> & src)
Purpose	Deliver data values from field src.
<i>Type</i>	All numeric types
Example	<code>r = QDP_expose(a);</code>

This function grants direct access to the data values contained in the QDP field `src`. The return value is a pointer to an array of QLA data `dest` of type *QLA_Type*. The order of the data is given by `Layout::linearIndex`. No QDP operations except `QDP_insert` are permitted on exposed data until `QDP_reset` is called. (See below.)

Returning control of QDP data

Prototype	<code>void QDP_reset(<i>Type</i>& field)</code>
Purpose	Returns control of data values to QDP.
<i>Type</i>	All numeric types
Example	<code>QDP_reset(r);</code>

This call signals to QDP that the user is ready to resume QDP operations with the data in the specified field.

Extracting QDP data

Prototype	<code>void QDP_extract(multiid<<i>QLA_Type</i>>& dest, const <i>Type</i>& src, const Subset& s)</code>
Purpose	Copy data values from field src to array dest.
<i>Type</i>	All numeric types
Example	<code>LatticeReal a; multiid<Real> r(Layout::sitesOnNode()); QDP_extract(r,a,even);</code>

The user must allocate the space of size `Layout::sitesOnNode()` for the destination array before calling this function, regardless of the size of the subset.

This function copies the data values contained in the QDP field `src` to the destination field. Only values belonging to the specified subset are copied. Any values in the destination array not associated with the subset are left unmodified. The order of the data is given by `Layout::linearIndex`. Since a copy is made, QDP operations involving the source field may proceed without disruption.

Inserting QDP data

Prototype	<code>void QDP_insert(<i>Type</i>& dest, const <i>QLA_Type</i>& src, const Subset& s)</code>
Purpose	Inserts data values from QLA array <code>src</code> .
Type	All numeric types
Example	<code>multi1d<Fermion> a(Layout::sitesOnNode()); LatticeFermion r; QDP_insert(r,a,odd);</code>

Only data associated with the specified subset are inserted. Other values are unmodified. The data order must conform to `Layout::linearIndex`. This call, analogous to a fill operation, is permitted at any time and does not interfere with QDP operations.

Suspending QDP communications If a user wishes to suspend QDP communications temporarily and carry on communications by other means, it is first necessary to call `QDP_suspend`.

Prototype	<code>void QDP_suspend(void)</code>
Purpose	Suspends QDP communications.
Example	<code>QDP_suspend();</code>

No QDP shifts can then be initiated until `QDP_resume` is called. However QDP linear algebra operations without shifts may proceed.

Resuming QDP communications To resume QDP communications one uses

Prototype	<code>void QDP_resume(void)</code>
Purpose	Restores QDP communications.
Example	<code>QDP_resume();</code>

4 Compilation with QDP

4.1 Generic header and macros

The compilation parameters:

N_d – the number of space-time dimensions

N_c – the dimension of the color vector space

N_s – the dimension of the spin vector space

are defined in `qdp++/include/params.h`. There are macros `ND`, `NC`, `NS` that are used to set the above parameters via

```
const int Nd = ND;
const int Nc = NC;
const int Ns = NS;
```

They are set in the build directories file `include/qdp_config.h` during configuration.

4.2 Nonuniform color and precision

Users wishing to vary color and precision within a single calculation must use specific type names whenever these types and names differ from the prevailing precision and color. Type declarations can be found in `qdp++/include/defs.h`. A convenient definition of a `LatticeColorMatrix` and `LatticeDiracFermion` is as follows:

```
typedef OLattice<PScalar<ColorMatrix<Complex<float>, Nc> > > LatticeColorMatrix
typedef OLattice<SpinVector<ColorVector<Complex<float>, Nc>, Ns> > LatticeFermion
```

However, for the user to choose a specific number of colors:

```
const int NN = 17 // work in SU(17)
typedef OLattice<PScalar<ColorMatrix<Complex<float>, NN> > > LatticeColorMatrix17
```

5 Spin Conventions

The following set of γ -matrices are used in four dimensions:

$$\begin{array}{llll}
 \gamma_0 & \begin{pmatrix} 0 & 0 & 0 & i \\ 0 & 0 & i & 0 \\ 0 & -i & 0 & 0 \\ -i & 0 & 0 & 0 \end{pmatrix} & \begin{pmatrix} 0 & i\sigma^1 \\ -i\sigma^1 & 0 \end{pmatrix} & -\sigma^2 \otimes \sigma^1 \\
 \gamma_1 & \begin{pmatrix} 0 & 0 & 0 & -1 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ -1 & 0 & 0 & 0 \end{pmatrix} & \begin{pmatrix} 0 & -i\sigma^2 \\ i\sigma^2 & 0 \end{pmatrix} & \sigma^2 \otimes \sigma^2 \\
 \gamma_2 & \begin{pmatrix} 0 & 0 & i & 0 \\ 0 & 0 & 0 & -i \\ -i & 0 & 0 & 0 \\ 0 & i & 0 & 0 \end{pmatrix} & \begin{pmatrix} 0 & i\sigma^3 \\ -i\sigma^3 & 0 \end{pmatrix} & -\sigma^2 \otimes \sigma^3 \\
 \gamma_3 & \begin{pmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{pmatrix} & \begin{pmatrix} 0 & \mathbf{1} \\ \mathbf{1} & 0 \end{pmatrix} & \sigma^1 \otimes 1
 \end{array}$$

The basis is chiral. All the possible gamma matrix products are represented via

$$\Gamma(n) = \gamma_0^{n_0} \gamma_1^{n_1} \gamma_2^{n_2} \gamma_3^{n_3}$$

where n_i are single bit fields. Since γ_0 comes first the bit for it must come first. So, $\gamma_5 = \gamma_0 \gamma_1 \gamma_2 \gamma_3$ is represented as 1111b = 15d, and $\gamma_0 \gamma_1 \gamma_3$ is represented as 1011b = 11d (note the ordering). The conventional γ -matrices are

$$\begin{aligned}
 \Gamma(1) &= \gamma_0 \\
 \Gamma(2) &= \gamma_1 \\
 \Gamma(4) &= \gamma_2 \\
 \Gamma(8) &= \gamma_3
 \end{aligned}$$

This enumeration is γ -basis independent.

6 Implementation Details

The following table lists some of the QDP headers.

name	purpose
qdp.h	Master header and QDP utilities
qdptype.h	Main class definition
qdpexpr.h	Expression class definition
primitive.h	Main header for all primitive types
primscale.h	Scalar primitive class and operations
primmatrix.h	Matrix primitive and operations
primvector.h	Vector primitive and operations
primseed.h	Seed (random number) primitive
reality.h	Complex number internal class
simpleword.h	Machine word-type operations

7 Supported Operations

This section describes in some detail the names and functionality for all functions in the interface involving linear algebra with and without shifts.

All QDP objects are of type QDPType, and QDP functions act on objects of this base class type. Unless otherwise indicated, operations occur on all sites in the specified subset of the target, often an assignment statement or object definition. The indexing of a QDPType returns an lvalue suitable for assignment (but not object definition). It is also used to narrow the lattice sites participating in a global reduction since the result of such a reduction is a lattice scalar, hence are independent of lattice sites.

Supported operations are listed below. Convention: prototypes are basically of the form:

```
QDPType unary_function(const QDPType&)
QDPType binary_function(const QDPType&, const QDPType&)
```

7.1 Subsets and Maps

```
Set::make(const SetFunc&) Set construction of ordinality num subsets. func maps
                           coordinates to a coloring in [0,num)
Map::make(const MapFunc&) Construct a map function from source sites to the dest
                           site.
```

7.2 Infix operators

Unary infix (e.g., operator-):

```
- : negation
+ : unary plus
~ : bitwise not
! : boolean not
```

Binary infix (e.g., operator+):

```
+ : addition
- : subtraction
* : multiplication
/ : division
% : mod
& : bitwise and
| : bitwise or
^ : bitwise exclusive or
<< : left-shift
>> : right-shift
```

Comparisons (returning booleans, e.g., operator<):

<, <=, >, >=, ==, !=
&& : and of 2 booleans
|| : or of 2 boolean

Assignments (e.g., operator+=):

=, +=, -=, *=, /=, %=, |=, &=, ^=, <<=, >>=

Trinary:

where(bool, arg1, arg2) : the C trinary "?" operator -> (bool) ? arg1 : arg2

7.3 Functions (standard C math lib)

Unary:

cos, sin, tan, acos, asin, atan, cosh, sinh, tanh,
exp, log, log10, sqrt,
ceil, floor, fabs

Binary:

ldexp, pow, fmod, atan2

7.4 Additional functions (specific to QDP)

Unary:

adj : hermitian conjugate (adjoint)
conj : complex conjugate
transpose : matrix tranpose, on a scalar it is a nop
trace : matrix trace
real : real part
imag : imaginary part
colorTrace : trace over color indices
spinTrace : trace over spin indices
timesI : multiplies argument by imag "i"
localNorm2 : on fibers computes trace(adj(source)*source)

Binary:

cmplx : returns complex object $\text{arg1} + i*\text{arg2}$
localInnerProduct : at each site computes $\text{trace}(\text{adj}(\text{arg1})*\text{arg2})$
outerProduct : at each site constructs $(\text{arg1}_i * \text{arg2}_j)_{ij}$

7.5 In place functions

random(dest) : uniform random numbers - all components
gaussian(dest) : uniform random numbers - all components
copymask(dest,mask,src) : copy src to dest under boolean mask

7.6 Broadcasts

Broadcasts via assignments (via, operator=):

<LHS> = <constant> : globally set conforming LHS to constant
<LHS> = zero : global always set LHS to zero

7.7 Global reductions

sum(arg1) : sum over lattice indices returning object of same fiber type
norm2(arg1) : sum(localNorm2(arg1))
innerProduct(arg1,arg2) : sum(localInnerProduct(arg1,arg2))
sumMulti(arg1,Set) : sum over each subset of Set returning #subset objects of same fiber type

7.8 Accessors

Peeking and poking (accessors) into various component indices of objects.

peekSite(arg1,multiId<int> coords) : return object located at lattice coords
peekColor(arg1,int row,int col) : return color matrix elem row and col
peekColor(arg1,int row) : return color vector elem row
peekSpin(arg1,int row,int col) : return spin matrix elem row and col
peekSpin(arg1,int row) : return spin vector elem row

pokeSite(dest,src,multiId<int> coords) : insert into site given by coords
pokeColor(dest,src,int row,int col) : insert into color matrix elem row and col
pokeColor(dest,src,int row) : insert into color vector elem row
pokeSpin(dest,src,int row,int col) : insert into spin matrix elem row and col
pokeSpin(dest,src,int row) : insert into spin vector elem row

7.9 More exotic functions:

Applies spin projection $(1 + isign * \gamma_\mu) * \psi$ returning a half spin vector or matrix

`spinProject(QDPTyp psi, int dir, int isign)`

Applies spin reconstruction of $(1 + isign * \gamma_\mu) * \psi$ returning a full spin vector or matrix

`spinReconstruct(QDPTyp psi, int dir, int isign)`

7.10 Operations on subtypes

Types in the QDP interface are parameterized by a variety of types, and can look like the following:

```
typedef OLattice<PScalar<PColorMatrix<RComplex<float>, Nc> > > LatticeColorMatrix
typedef OLattice<PSpinVector<PColorVector<RComplex<float>, Nc>, Ns> > LatticeFermion
```

- *Word type*: int, float, double, bool. Basic machine types.
- *Reality type*: RComplex or RScalar.
- *Primitive type*: PScalar, PVector, PMatrix, PSeed.
- *Inner grid type*: IScalar or ILattice.
- *Outer grid type*: OScalar or OLattice.

Supported operations for each type level as follows:

Grid type: OScalar, OLattice, IScalar, ILattice

All operations listed in Sections 7.2–7.9

Primitive type:

PScalar: All operations listed in Sections 7.2–7.9

PMatrix<N>:

<i>Unary:</i>	<code>-(PMatrix), +(PMatrix)</code>
<i>Binary:</i>	<code>-(PMatrix,PMatrix), +(PMatrix,PMatrix), *(PMatrix,PScalar),</code> <code>*(PScalar,PMatrix), *(PMatrix,PMatrix)</code>
<i>Comparisons:</i>	none
<i>Assignments:</i>	<code>=(PMatrix), =(PScalar), ==(PMatrix), +=(PMatrix), *=(PScalar)</code>
<i>Trinary:</i>	<code>where</code>
<i>C-lib funcs:</i>	none
<i>QDP funcs:</i>	all
<i>In place funcs:</i>	all
<i>Reductions:</i>	all

PVector<N>:

Unary: ~(PVector), +(PVector)
Binary: -(PVector,PVector), +(PVector,PVector), *(PVector,PScalar),
 *(PScalar,PVector), *(PMatrix,PVector)
Comparisons: none
Assignments: =(PVector), ==(PVector), +=(PVector), *=(PScalar)
Trinary: where
C-lib funcs: none
QDP funcs: real, imag, timesI, localNorm2, cmplx, localInnerProduct,
 outerProduct
In place funcs: all
Broadcasts: =(Zero)
Reductions: all

PSpinMatrix<N>: Inherits same operations as PMatrix

Unary: spinTrace
Binary: *(PSpinMatrix,Gamma), *(Gamma,PSpinMatrix)
Exotic: peekSpin, pokeSpin, spinProjection, spinReconstruction

PSpin Vector<N>: Inherits same operations as PVector

Binary: *(Gamma,PSpinVector)
Exotic: peekSpin, pokeSpin, spinProjection, spinReconstruction

PColorMatrix<N>: Inherits same operations as PMatrix

Unary: colorTrace
Binary: *(PColorMatrix,Gamma), *(Gamma,PColorMatrix)
Exotic: peekColor, pokeColor

PColor Vector<N>: Inherits same operations as PVector

Binary: *(Gamma,PColorVector)
Exotic: peekColor, pokeColor

Reality: *RScalar, RComplex*

All operations listed in Sections 7.2–7.9

Word: *int, float, double, bool*

All operations listed in Sections 7.2–7.9. Only boolean ops allowed on bool.

8 Detailed function description

The purpose of this section is to show some explicit prototypes and usages for the functions described in Section 7. In that section, all the functions are shown with complete information on which operations and their meaning are supported on some combination of types. The purpose of this section is something like the inverse - namely show all the functions and what are some (selected) usages.

8.1 Unary Operations

Elementary unary functions on reals

Syntax	<i>Type</i> <i>func</i> (const <i>Type</i> & a)
Meaning	$r = \text{func}(a)$
<i>func</i>	cos, sin, tan, acos, asin, atan, sqrt, abs, exp, log, sign
<i>Type</i>	Real, LatticeReal

Elementary unary functions on complex values

Syntax	<i>Type</i> <i>func</i> (const <i>Type</i> & a)
Meaning	$r = \text{func}(a)$
<i>func</i>	exp, sqrt, log
<i>Type</i>	Complex, LatticeComplex

Assignment operations

Syntax	<i>Type</i> operator=(const <i>Type</i> & r, const <i>Type</i> & a)
Meaning	$r = a$
<i>Type</i>	All numeric types

Shifting

Syntax	<i>Type</i> shift(const <i>Type</i> & a, int sign, int dir)
Meaning	$r = a$
<i>Type</i>	All numeric types

Hermitian conjugate

Syntax	<i>Type</i> adj(const <i>Type</i> & a)
Meaning	$r = a^\dagger$
<i>Type</i>	Real, Complex, ColorMatrix, DiracPropagator Also corresponding lattice variants

Transpose

Syntax	<i>Type</i> transpose(const <i>Type</i> & a)
Meaning	$r = \text{transpose}(a)$
<i>Type</i>	Real, Complex, ColorMatrix, DiracPropagator Also corresponding lattice variants

Complex conjugate

Syntax	<i>Type</i> conj(const <i>Type</i> & a)
Meaning	$r = a^*$
<i>Type</i>	Real, Complex, ColorMatrix, DiracFermion, DiracPropagator Also corresponding lattice variants

8.2 Type conversion

Types can be precision converted via a conversion function of the destination class.

Convert integer or float to double

Syntax	<i>Type2</i> <i>Type2</i> (const <i>Type1</i> & a)
Example	LatticeReal a; LatticeRealD r = LatticeRealD(a) LatticeColorMatrix a; LatticeColorMatrixD r = LatticeColorMatrixD(a)
<i>Type1</i>	All single precision numeric types
<i>Type2</i>	All conforming double precision numeric types

Convert double to float

Syntax	<i>Type2</i> <i>Type2</i> (const <i>Type1</i> & a)
Example	LatticeRealD a; LatticeReal r = LatticeReal(a) LatticeColorMatrixD a; LatticeColorMatrix r = LatticeColorMatrix(a)
<i>Type1</i>	All double precision numeric types
<i>Type2</i>	All conforming single precision numeric types

Integer to real

Syntax	<i>Type2</i> <i>Type2</i> (const <i>Type1</i> & a)
Example	LatticeInt a; LatticeReal r = LatticeReal(a)
<i>Type1</i>	All integer precision numeric types
<i>Type2</i>	All conforming real precision numeric types

Real to integer

Syntax	<i>Type2</i> <i>Type2</i> (const <i>Type1</i> & a)
Example	LatticeReal a; LatticeInt r = LatticeInt(a)

Real to float

Syntax	<code>float toFloat(const Real& a)</code>
Meaning	<code>r = float(a);</code>
Example	<code>Real a; float r = toFloat(a);</code>

The QDP type `Real` is not a primitive type, so an explicit conversion is provided.

Double to double

Syntax	<code>double toDouble(const RealD& a)</code>
Meaning	<code>r = double(a);</code>
Example	<code>RealD a; double r = toDouble(a);</code>

The QDP type `RealD` is not a primitive type, so an explicit conversion is provided.

Bool to bool

Syntax	<code>bool toBool(const Boolean& a)</code>
Meaning	<code>r = bool(a);</code>
Example	<code>Boolean a; bool r = toBool(a);</code>

The QDP type `Boolean` is not a primitive type, so an explicit conversion is provided.

8.3 Operations on complex arguments

Convert real and imaginary to complex

Syntax	<code>Type cmplx(const Type1& a, const Type2& b)</code>
Meaning	$\text{Re } r = a, \text{Im } r = b$
<i>Type1</i>	constant, Real, Also corresponding lattice variants
<i>Type2</i>	constant, Real, Also corresponding lattice variants
<i>Type</i>	Complex, Also corresponding lattice variants
Example	<code>Real a;</code> <code>Complex = cmplx(a, 0);</code>

Real part of complex

Syntax	<code>Type real(const Type& a)</code>
Meaning	$r = \text{Re } a$

Imaginary part of complex

Syntax	<code>Type imag(const Type& a)</code>
Meaning	$r = \text{Im } a$

8.4 Component extraction and insertion

Accessing a site object

Syntax	<code>Type peekSite(const LatticeType& a, const multild<int>& c)</code>
Meaning	$r = a[x]$

Accessing a color matrix element

Syntax	<code>LatticeComplex peekColor(const LatticeColorMatrix& a, int i, int j) LatticeSpinMatrix peekColor(const LatticeDiracPropagator& a, int i, int j)</code>
Meaning	$r = a_{i,j}$

Inserting a color matrix element

Syntax	<code>LatticeColorMatrix& pokeColor(LatticeColorMatrix& r, const LatticeComplex& a, int i, int j)</code>
Meaning	$r_{i,j} = a$

Accessing a color vector element

Syntax	<code>LatticeComplex peekColor(const LatticeColorVector& a, int i) LatticeSpinVector peekColor(const LatticeDiracFermion& a, int i)</code>
Meaning	$r = a_i$

This function will extract the desired color component with all the other indices unchanged.

A lattice color vector is another name (typedef) for a LatticeStaggeredFermion. Namely, an object that is vector in color spin and a scalar in spin space. Together with spin accessors, one can build a LatticeDiracFermion.

Inserting a color vector element

Syntax	<code>LatticeColorVector& pokeColor(LatticeColorVector& r, const LatticeComplex& a, int i)</code>
Meaning	$r_i = a$

This function will extract the desired color component with all the other indices unchanged.

A lattice color vector is another name (typedef) for a LatticeStaggeredFermion. Namely, an object that is vector in color spin and a scalar in spin space. Together with spin accessors, one can build a LatticeDiracFermion or a LatticeDiracPropagator.

Accessing a spin matrix element

Syntax	<code>LatticeComplex peekSpin(const LatticeSpinMatrix& a, int i, int j)</code> <code>LatticeColorMatrix peekSpin(const LatticeDiracPropagator& a, int i, int j)</code>
Meaning	$r = a_{i,j}$

Inserting a spin matrix element

Syntax	<code>LatticeSpinMatrix& pokeSpin(LatticeSpinMatrix& r, const LatticeComplex& a, int i, int j)</code>
Meaning	$r_{i,j} = a$

Accessing a spin vector element

Syntax	<code>LatticeComplex peekSpin(const LatticeSpinVector& a, int i)</code> <code>LatticeColorVector peekSpin(const LatticeDiracFermion& a, int i)</code>
Meaning	$r = a_i$

This function will extract the desired spin component with all the other indices unchanged.

A lattice spin vector is an object that is a vector in spin space and a scalar in color space. Together with color accessors, one can build a `LatticeDiracFermion` or a `LatticeDiracPropagator`.

Inserting a spin vector element

Syntax	<code>LatticeSpinVector& pokeSpin(LatticeSpinVector& r, const LatticeComplex& a, int i)</code>
Meaning	$r_i = a$

This function will extract the desired spin component with all the other indices unchanged.

A lattice spin vector is an object that is a vector in spin space and a scalar in color space. Together with color accessors, one can build a `LatticeDiracFermion` or a `LatticeDiracPropagator`.

Trace of matrix

Syntax	<code>Type2 trace(const Type1& a)</code>
Meaning	$r = \text{Tr } a$
Type1	ColorMatrix, DiracPropagator, Also corresponding lattice variants
Type2	Complex, Complex, Also corresponding lattice variants
Example	<code>LatticeColorMatrix a;</code> <code>LatticeComplex r = trace(a);</code>

Traces over all matrix indices. It is an error to trace over a vector index. It will trivially trace a scalar variable.

Color trace of matrix

Syntax	<i>Type2</i> traceColor(const <i>Type1</i> & a)
Meaning	$r = \text{Tr } a$
<i>Type1</i>	SpinMatrix, Also corresponding lattice variants
<i>Type2</i>	Complex, Also corresponding lattice variants
Example	LatticeDiracPropagator a; LatticeSpinMatrix r = traceColor(a);

Traces only over color matrix indices. It is an error to trace over a color vector index. All other indices are left untouched. It will trivially trace a scalar variable.

Spin trace of matrix

Syntax	<i>Type2</i> traceSpin(const <i>Type1</i> & a)
Meaning	$r = \text{Tr } a$
<i>Type1</i>	DiracPropagator, Also corresponding lattice variants
<i>Type2</i>	ColorMatrix, Also corresponding lattice variants
Example	LatticeDiracPropagator a; LatticeColorMatrix r = traceSpin(a);

Traces only over spin matrix indices. It is an error to trace over a spin vector index. All other indices are left untouched. It will trivially trace a scalar variable.

Dirac spin projection

Syntax	<i>Type2</i> spinProject(const <i>Type1</i> & a, int d, int p)
Meaning	$r = (1 + p\gamma_d)a$
<i>Type1</i>	DiracFermion, Also corresponding lattice variants
<i>Type2</i>	HalfFermion, Also corresponding lattice variants

Dirac spin reconstruction

Syntax	<i>Type2</i> spinReconstruct(const <i>Type1</i> & a, int d, int p)
Meaning	$r = \text{recon}(p, d, a)$
<i>Type1</i>	HalfFermion, Also corresponding lattice variants
<i>Type2</i>	DiracFermion, Also corresponding lattice variants

8.5 Binary Operations with Constants

Multiplication by real constant

Syntax	<i>Type</i> operator*(const Real& a, const <i>Type</i> & b) <i>Type</i> operator*(const <i>Type</i> & b, const Real& a)
Meaning	$r = a * b$ (a real, constant)
<i>Type</i>	All floating types

Multiplication by complex constant

Syntax	<i>Type</i> operator*(const Real& a, const <i>Type</i> & b) <i>Type</i> operator*(const <i>Type</i> & b, const Real& a)
Meaning	$r = a * b$ (a complex, constant)
Type	All numeric types

Left multiplication by gamma matrix

Syntax	<i>Type</i> operator*(const <i>Gamma</i> & a, const <i>Type</i> & b)
Meaning	$r = \gamma_d * a$
<i>Gamma</i> <i>Type</i>	Gamma constructed from an explicit integer in $[0, N_s^2 - 1]$ SpinVector, SpinMatrix, HalfFermion, DiracFermion, DiracPropagator, and similar lattice variants
Example	<code>r = Gamma(7) * b;</code>

See Section 5 for details on γ -matrix conventions.

Right multiplication by gamma matrix

Syntax	<i>Type</i> operator*(const <i>Type</i> & a, const <i>Gamma</i> & b)
Meaning	$r = a * \gamma_d$
<i>Gamma</i> <i>Type</i>	Gamma constructed from an explicit integer in $[0, N_s^2 - 1]$ SpinMatrix, DiracPropagator, and similar lattice variants
Example	<code>r = a * Gamma(15);</code>

See Section 5 for details on γ -matrix conventions.

8.6 Binary Operations with Fields

Division of real fields

Syntax	<i>Type</i> operator/(const <i>Type</i> & a, const <i>Type</i> & b)
Meaning	$r = a/b$

Addition

Syntax	<i>Type</i> operator+(const <i>Type</i> & a, const <i>Type</i> & b)
Meaning	$r = a + b$
Type	All numeric types

Subtraction

Syntax	<i>Type</i> operator-(const <i>Type</i> & a, const <i>Type</i> & b)
Meaning	$r = a - b$
Type	All numeric types

Multiplication: uniform types

Syntax	<i>Type</i> operator*(const <i>Type</i> & a, const <i>Type</i> & b)
Meaning	$r = a * b$
<i>Type</i>	constant, Real, Complex, Integer, ColorMatrix, SpinMatrix, DiracPropagator

ColorMatrix matrix from outer product

Syntax	<i>Type</i> outerProduct(const <i>Type1</i> & a, const <i>Type2</i> & b)
Meaning	$r_{i,j} = a_i * b_j^*$
<i>Type1,2</i>	ColorVector, LatticeColorVector
<i>Type</i>	ColorMatrix, LatticeColorMatrix

Left multiplication by gauge matrix

Syntax	<i>Type</i> operator*(const <i>Type1</i> & a, const <i>Type</i> & b)
Meaning	$r = a * b$
<i>Type1</i>	ColorMatrix, LatticeColorMatrix
<i>Type</i>	constant, Complex, ColorMatrix, ColorVector, SpinVector, DiracPropagator, and similar lattice variants

Right multiplication by gauge matrix

Syntax	<i>Type</i> operator*(const <i>Type</i> & a, const <i>Type1</i> & b)
Meaning	$r = a * b$
<i>Type1</i>	ColorMatrix, LatticeColorMatrix
<i>Type</i>	ColorMatrix, SpinMatrix, DiracPropagator, and similar lattice variants

8.7 Boolean and Bit Operations

Comparisons

Syntax	<i>Type2</i> op(const <i>Type</i> & a, const <i>Type1</i> & b)
Meaning	$r = a \text{ op } b$ or $r = \text{op}(a, b)$
op	<, >, !=, <=, >=, ==
<i>Type1</i>	Integer, Real, RealD, and similar lattice variants
<i>Type2</i>	Boolean or LatticeBoolean (result is lattice if any arg is lattice)

Elementary binary operations on integers

Syntax	<i>Type2</i> op(const <i>Type</i> & a, const <i>Type1</i> & b)
Meaning	$r = a \text{ op } b$ or $r = \text{op}(a, b)$
op	<<, >>, & (and), (or), ^ (xor), mod, max, min

Elementary binary operations on reals

Syntax	<i>Type</i> op(const <i>Type1</i> & a, const <i>Type2</i> & b)
Meaning	$r = a \text{ op } b$ or $r = \text{op}(a, b)$
op	mod, max, min
<i>Type</i>	Real, RealD, and similar lattice variants

Boolean Operations

Syntax	<i>Type</i> op(const <i>Type</i> & a, const <i>Type</i> & b)
Meaning	$r = a \text{ op } b$
op	(or), & (and), ^ (xor)
<i>Type</i>	Boolean, LatticeBoolean

Syntax	<i>Type</i> op(const <i>Type</i> & a)
Meaning	$r = \text{not } a$
op	! (not)
<i>Type</i>	Boolean, LatticeBoolean

Copymask

Syntax	void copymask(const <i>Type2</i> & r, const <i>Type1</i> & a, const <i>Type1</i> & b)
Meaning	$r = b$ if a is true
<i>Type</i>	All numeric types

8.8 Reductions

Global reductions sum over all lattice sites in the subset specified by the left hand side of the assignment.

Norms

Syntax	Real norm2(<i>Type</i> & a)
Meaning	$r = \sum a ^2$
<i>Type</i>	All numeric types

Inner products

Syntax	Complex innerProduct(<i>Type</i> & a, const <i>Type</i> & b)
Meaning	$r = \sum a \cdot b$
<i>Type</i>	All numeric types

Global sums

Syntax	<i>Type</i> sum(const Lattice <i>Type</i> & a)
Meaning	$r = \sum a$
<i>Type</i>	All numeric non-lattice scalar types

8.9 Fills

Coordinate function fills

Syntax	<code>LatticeInt Layout::latticeCoordinate(int d)</code>
Meaning	$r = f(d)$ for direction d .
Purpose	Return the lattice coordinates in direction d

The call `Layout::latticeCoordinate(d)` returns an integer lattice field with a value on each site equal to the integer value of the d th space-time coordinate on that site.

Constant fills

Syntax	<code>LatticeType operator=(LatticeType& r, const Type& a)</code>
Meaning	$r = a$ for all sites
Type	All non-lattice objects
Example	<code>Real a = 2.0;</code> <code>LatticeReal r = a;</code>

Constant (or lattice global) fills are always defined for lattice scalar objects broadcasting to all lattice sites. These are broadcasts of a lattice scalar type to a conforming lattice type.

NOTE, one can not fill a `LatticeColorVector` with a `Real`.

Syntax	<code>LatticeType operator=(LatticeType& r, const Type& a)</code>
Meaning	$r = \text{diag}(a, a, \dots)$ (constant a)
Type	Complex, ColorMatrix, SpinMatrix
Example	<code>Real a = 2.0;</code> <code>LatticeColorMatrix r = a;</code>

Only sets the diagonal part of a field to a constant a times the identity.

This fill can only be used on primitive types that are scalars or matrices. E.g., it can not be used for a *vector* field since there is no meaning of diagonal. NOTE, a zero cannot be distinguished from a constant like 1. To initialize to zero the `zero` argument must be used.

Zero fills

Syntax	<code>Type operator=(Type& r, const Zero& zero)</code>
Meaning	$r = 0$
Type	All numeric types
Example	<code>LatticeDiracFermion r = zero;</code>

This is the only way to fill a vector field with a constant (like zero).

Uniform random number fills

Syntax	<code>void random(Type& r)</code>
Meaning	r random, uniform on $[0, 1]$
Type	All floating types

Gaussian random number fills

Syntax	<code>void gaussian(<i>Type</i>& r)</code>
Meaning	<i>r</i> normal Gaussian
<i>Type</i>	All floating types

Seeding the random number generator

Syntax	<code>void RNG::setrn(const Seed& a)</code>
Meaning	Initialize the random number generator with seed state <i>a</i>

For details see the discussion of the corresponding scalar function `random.h`.

Extracting the random number generator seed

Syntax	<code>void RNG::savern(Seed& r)</code>
Meaning	Extract the random number generator into seed state <i>r</i>

For details see the discussion of the corresponding scalar function `random.h`.