



Grid : Data Parallel QCD Library

Peter A Boyle, Azusa Yamaguchi

Intel Parallel Computing Centre @ Higgs Centre for Theoretical Physics, University of Edinburgh

Guido Cossu

KEK, Tsukuba

We present progress on a new C++ data parallel QCD library. It enables the description of cartesian fields of arbitrary tensor mathematical types.

Ddata parallel interface, conformable array syntax with Cshift and masked operation (c.f. QDP++, cmfortran or HPF).

Three distinct forms of parallelism are transparently used underneath the single simple interface:

- MPI task parallelism
- OpenMP thread parallelism
- SIMD vector parallelism.

The SIMD vector parallelism achieves nearly 100% SIMD efficiency due to the adoption of a virtual node layout transformation, similar to those in the Connection Machine.

This ensures identical and independent work lies in adjacent SIMD lanes. SSE, AVX, AVX2, AVX512 and Arm Neon SIMD targets are implemented.

The library is under development. Solvers for Wilson, Domain, and multiple 5d chiral fermions (Cayley, Continued fraction, partial fraction) are implemented.

Features differing from QDP++:

- Shift by arbitrary distance
- Storage:
 - Checkerboarded fields are half length
 - 5d fields are same type as 4d with different Grid.
- Multiple grids and Projection/Promotion support
 - blockProject, blockPromote, blockInnerProduct
- Stencil object
 - encapsulates geometry of operation
 - Performs halo exchange
 - Simple to write kernel for Dirac operators.
- C++11 : expression template engine < 200 LoC

GRID parallel library

- Geometrically decompose cartesian arrays across nodes (MPI)
- Subdivide node volume into smaller virtual nodes
- Spread virtual nodes across SIMD lanes
- Use OpenMP+MPI+SIMD to process conformable array operations
- Same instructions executed on many nodes, each node operates on four virtual nodes

Over decompose the subgrid

Interleave overdecomposed sub-nodes in SIMD lanes

Grid for single overdecomposed sub-nodes with for vector data type

Division of sub-nodes to provide with 100% SIMD efficiency

GRID data parallel CSHIFT details

- Crossing between SIMD lanes restricted to during shifts between virtual nodes
- Code for N-virtual nodes is identical to scalar code for one, except datum is N fold bigger

$$\begin{matrix} (A, B, C, D) & (E, F, G, H) & \rightarrow & (AE, BF, CG, DH) \\ \text{virtual subnode} & \text{virtual subnode} & & \text{Packed SIMD} \end{matrix}$$

- CSHIFT involves a CSHIFT of SIMD, and a permute only on the surface

$$(AE, BF, CG, DH) \rightarrow (BF, CG, DH, AE)$$

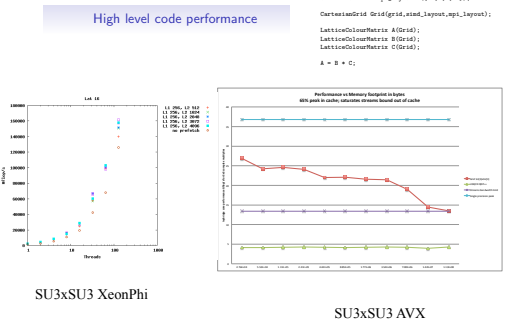
csift bulk

$$(BF, CG, DH, EA) \rightarrow (B, C, D, E) \quad (F, G, H, A)$$

permute face

virtual subnode

- Shuffle overhead is suppressed by surface to volume ratio



vSIMD performance portable SIMD library


```

Define performant classes vRealF, vRealD, vComplexF, vComplexD.
...

```

- Conclusion: Modify data layout to align data parallel operations to SIMD hardware
- Conformable array operations simple and vectorise perfectly

- Fermion actions
 - Wilson
 - Chiral fermions
 - (Cayley/ContFrac/PartFrac)
 - (Zolotarev/Tanh)
 - (Wilson/Shamir/Mobius)
 - Redblack 4d and unpreconditioned
 - CG, MCR, Multishift CG, GCR
 - Zolotarev, Chebyshev, Remez approx
- Gauge actions
 - Quenched Wilson/Symanzik/Iwasaki
- Multigrid coarse operators
- Nersc file I/O
- Roadmap:
 - (R)HMC, gauge and fermion force terms
 - FFT, measurements
 - Link smearing



GRID data parallel template library

Ordering	Layout	Vectorisation	Data Reuse
Microprocessor	Array-of-Structs (AoS)	Hard	Maximised
Vector	Struct-of-Array (SoA)	Easy	Minimised
Bagel	Array-of-structs-of-short-vectors (AoSoSvV)	Easy	Maximised

- Opaque C++11 containers hide layout from user
- Automatically transform layout of arrays of mathematical objects using vSIMD scalar, vector, matrix, higher rank tensors.

Stencil support eases the pain of optimised matrix multiplies

Pass the stencil a list of directions and displacements

```

int npoint;
std::vector<int> directions ;
std::vector<int> displacements;
CartesianStencil Stencil(&CoarseGrid,npoint,Even,directions,displacements)

```

Coarse grid operator in Grid

```


void M(const CoarseVector &in, CoarseVector &out){
conformable(_grid,in,_grid);
conformable(in,_grid,out,_grid);
SimpleCompressor::siteVectors compressor;
Stencil::HaloExchange(in,com_buf,compressor);
PARALLEL_FOR_LOOP
for(int i=0; i<=Grid()->nSites();i++){
siteVector res = zero;
Stencil::nbrz
int offset,local,perm,ptype;
for(int point=0;point<geom.npoint;point++){
offset = Stencil::offsets [point][i];
local = Stencil::is_local [point][i];
perm = Stencil::permute [point][i];
ptype = Stencil::permute_type [point];
if (local & ptype) {
permute(nbr, in._odata[offset], ptype);
} else if (local) {
nbr = in._odata[offset];
} else {
nbr = com_buf[offset];
}
res = res + A[point]->_odata[ss] * nbr;
}
vstream(out._odata[ss], res);
return nom2(out);
}

```

Stencil organises halo exchange for any vector type; compressor can do spin proj for Wilson fermions.

Stencil provides index of each neighbour (knows the geometry)

User dictates how to treat the internal indices in operator



What is the best SIMD strategy?

SIMD most efficient for independent but identical work e.g. apply N small dense matrix-vector multiplies in parallel:

```

for(int i=0; i<N; i++){
for(int j=0; j<R; j++){
fma(y[i*N+j], v[j], r[i]);
}
}

```

Vector = Matrix x Vector

Reduction of vector sum is bottleneck for small N

Many vectors = many matrices x many vectors

No reduction or SIMD lane crossing operations.

SIMD interleave

www.github.com/paboyle/Grid