

FMMTL: FMM Template Library

Generalized Framework for Kernel Matrices

Cris Cecka, Simon Layton

Institute for Applied Computational Science
School of Engineering and Applied Sciences
Harvard University

November 15, 2013



Dense Matrices

$$\mathbf{K} = \begin{bmatrix} K_{00} & K_{01} & \cdots & K_{0N} \\ K_{10} & K_{11} & \cdots & K_{1N} \\ \vdots & \vdots & \ddots & \vdots \\ K_{N0} & K_{N1} & \cdots & K_{NN} \end{bmatrix} \in \mathbb{K}^{N,N}$$

- Storage: $\mathcal{O}(N^2)$
- Matrix-vector product: $\mathcal{O}(N^2)$ ops.
- Matrix-matrix product: $\mathcal{O}(N^3)$ ops.
- Matrix factorizations (LU, QR, SVD): $\mathcal{O}(N^3)$ ops.

Structured Dense Matrices

- Low-rank Matrices
- Hierarchical Matrices



Kernel Matrices

$$\mathbf{K} = \begin{bmatrix} K(t_0, s_0) & K(t_0, s_1) & \cdots & K(t_0, s_N) \\ K(t_1, s_0) & K(t_1, s_1) & \cdots & K(t_1, s_N) \\ \vdots & \vdots & \ddots & \vdots \\ K(t_N, s_0) & K(t_N, s_1) & \cdots & K(t_N, s_N) \end{bmatrix} = [K_{ij}]$$

and the matrix equation

$$r_i = K(t_i, s_j) c_j$$

where

- $K(\cdot, \cdot)$: The *kernel* generating matrix elements,
- s_j : The *sources* of the kernel,
- c_j : The *charges* of the sources,
- t_i : The *targets* of the kernel (could be = s_j),
- r_j : The *results (potentials)* of the targets.



Kernel Matrices

$$\mathbf{K} = \begin{bmatrix} K(t_0, s_0) & K(t_0, s_1) & \cdots & K(t_0, s_N) \\ K(t_1, s_0) & K(t_1, s_1) & \cdots & K(t_1, s_N) \\ \vdots & \vdots & \ddots & \vdots \\ K(t_N, s_0) & K(t_N, s_1) & \cdots & K(t_N, s_N) \end{bmatrix} = [K_{ij}]$$

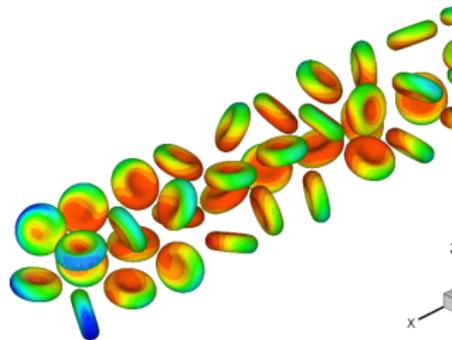
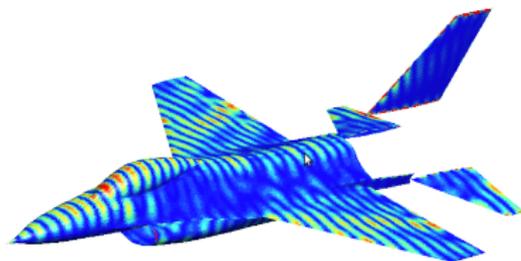
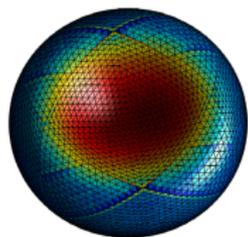
Appear in

Name/Equation	$K(x, y)$	Field
Laplace, Poisson	$1/ x - y $	$\mathbb{R}^3 \rightarrow \mathbb{R}$
Yukawa, Helmholtz	$e^{k x-y }/ x - y $	$\mathbb{R}^3 \rightarrow \mathbb{C}$
Stokes	$\frac{1}{ x-y } \left(\mathbf{I} + \frac{(x-y)(x-y)^T}{ x-y ^2} \right)$	$\mathbb{R}^3 \rightarrow \mathbb{R}^{3,3}$
Gaussian	$e^{-\varepsilon x-y ^2}$	$\mathbb{R}^n \rightarrow \mathbb{R}$

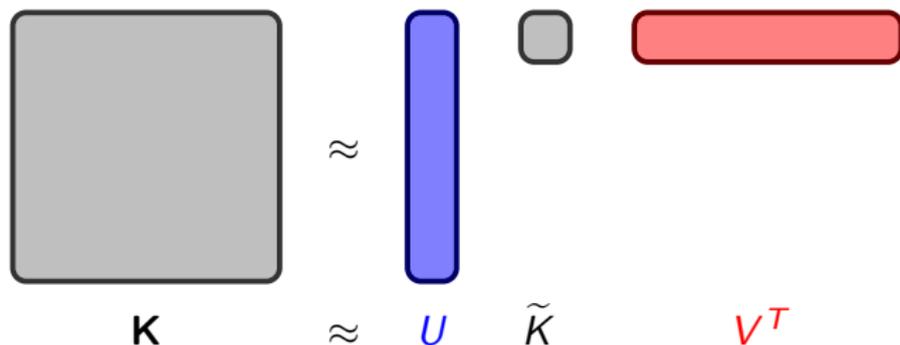
- Other physics:
 - Elastics, Maxwell, Biot-Savart, any Green's function.
- Other machine learning, interpolation
 - (Multi)quadric, Inv (Multi)quadric, splines, any RBF.



Flashy Applications



Low Rank Matrices

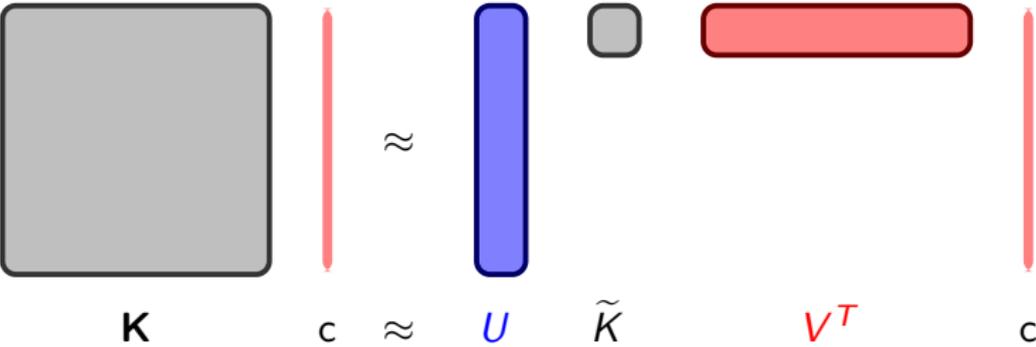
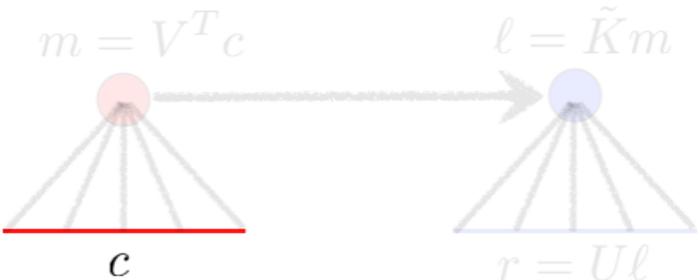


- Storage: $\mathcal{O}(rN)$ where $\tilde{K} \in \mathbb{K}^{r,r'}$
- Matrix-vector product: $\mathcal{O}(rN)$
- Constructed with:

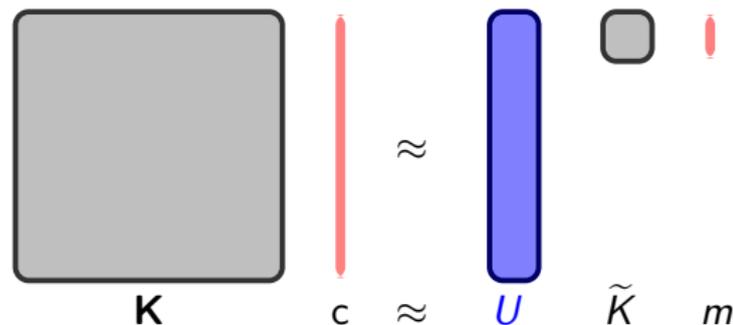
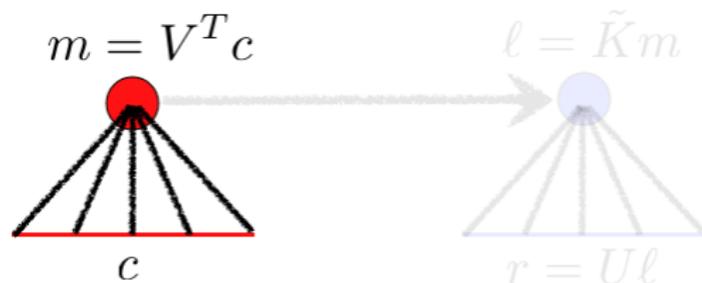
Algebraic techniques	Cost
SVD	$\mathcal{O}(N^3)$
RRLU, RRQR	$\mathcal{O}(rN^2)$
ACA	$\mathcal{O}(r^2N)$
Pseudo-skeletal, CUR	$\mathcal{O}(rN)$

Analytic techniques	Cost
Series expansion	$\mathcal{O}(rN)$
Interpolation	$\mathcal{O}(rN)$

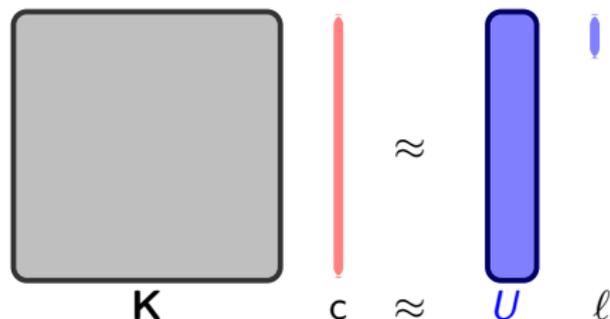
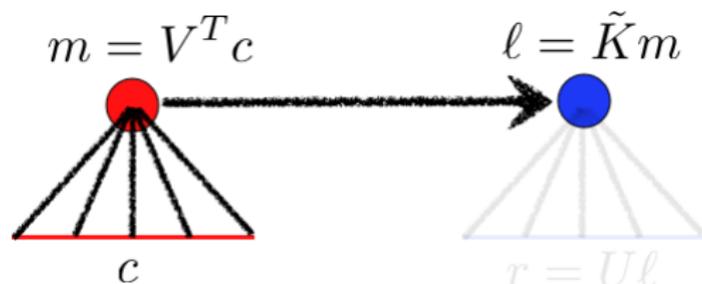
Low Rank Interpretation



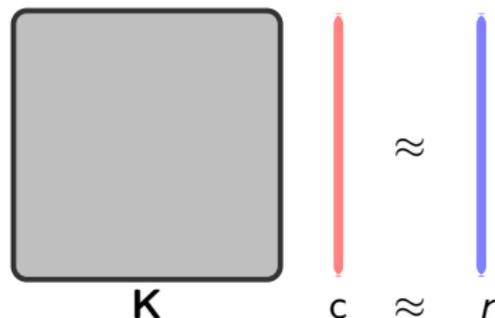
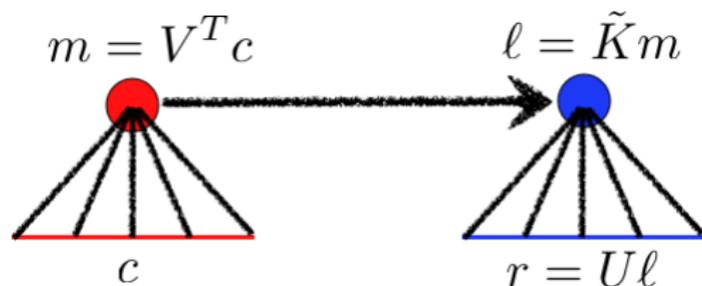
Low Rank Interpretation



Low Rank Interpretation

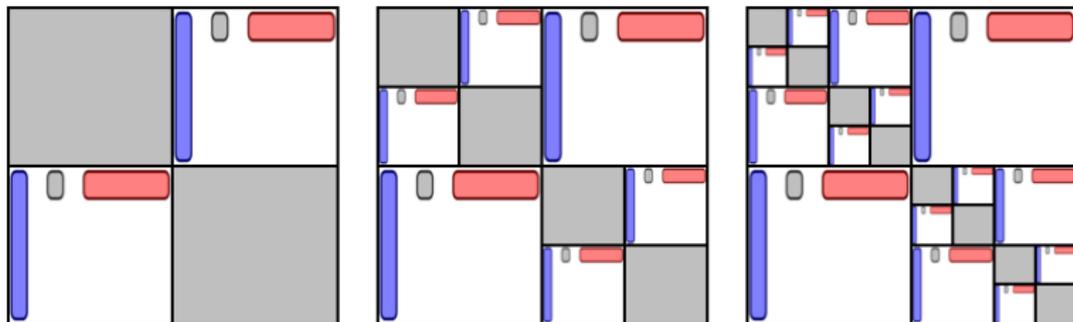


Low Rank Interpretation



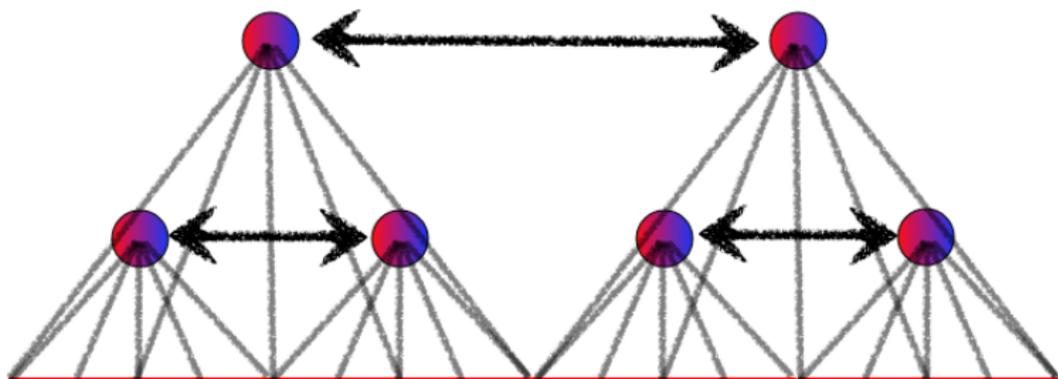
Hierarchically Low Rank

- Fully low rank matrices are scarce.
- Instead, hierarchically off-diagonal low-rank (HODLR):



HODLR Interpretation

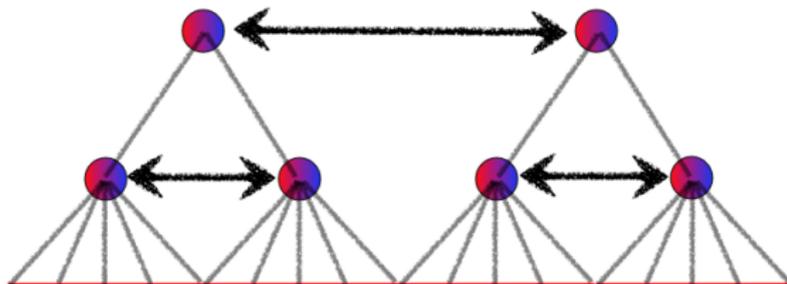
Disjoint “clusters” of sources and targets have low-rank interaction.



Additional Low Rank Operations

- Nested low-rank property.
 - Multipole/Local expansions may be passed up/down the tree.

$$\mathbf{K}_{IJ} = U_{C(I)} \tilde{U}_I \tilde{K}_{IJ} \tilde{V}_J^T V_{C(J)}^T$$



- Called hierarchically semiseparable (HSS) matrices.
- Well-separated condition
 - Stricter than simply disjoint clusters: Need well-separated.
 - Forces off-diagonal dense blocks.
 - Called \mathcal{H} and \mathcal{H}^2 matrices.



Another View

Using a low rank approximation,

$$K(\mathbf{r}_i + \mathbf{r}_0 + \mathbf{r}_j) = \sum_{p=1}^r \sum_{q=1}^{r'} U_p(\mathbf{r}_i) \tilde{K}_{pq}(\mathbf{r}_0) V_q^T(\mathbf{r}_j) + \varepsilon$$

• X_j

y_j •



Another View

Using a low rank approximation,

$$K(\mathbf{r}_i + \mathbf{r}_0 + \mathbf{r}_j) = \sum_{p=1}^r \sum_{q=1}^{r'} U_p(\mathbf{r}_i) \tilde{K}_{pq}(\mathbf{r}_0) V_q^T(\mathbf{r}_j) + \varepsilon$$

• X_j

y_j •

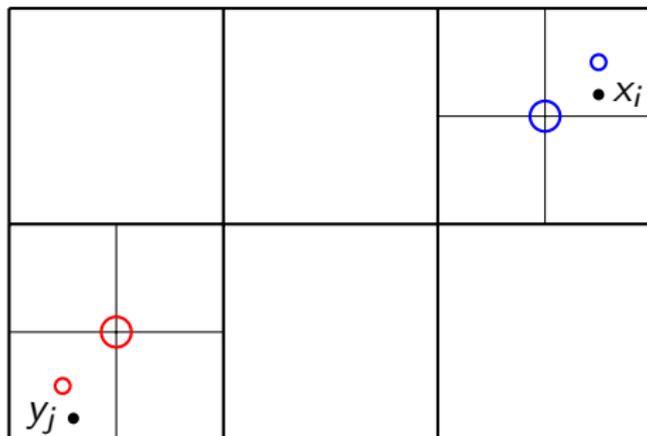


Another View

Using a low rank approximation,

$$K(\mathbf{r}_i + \mathbf{r}_0 + \mathbf{r}_j) = \sum_{p=1}^r \sum_{q=1}^{r'} U_p(\mathbf{r}_i) \tilde{K}_{pq}(\mathbf{r}_0) V_q^T(\mathbf{r}_j) + \varepsilon$$

We define “multipole” and “local” fields at the nodes of a tree:

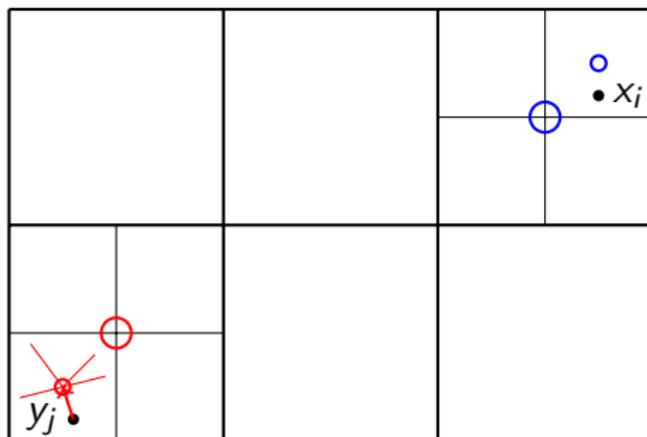


Another View

Using a low rank approximation,

$$K(\mathbf{r}_i + \mathbf{r}_0 + \mathbf{r}_j) = \sum_{p=1}^r \sum_{q=1}^{r'} U_p(\mathbf{r}_i) \tilde{K}_{pq}(\mathbf{r}_0) V_q^T(\mathbf{r}_j) + \varepsilon$$

We define “multipole” and “local” fields at the nodes of a tree:

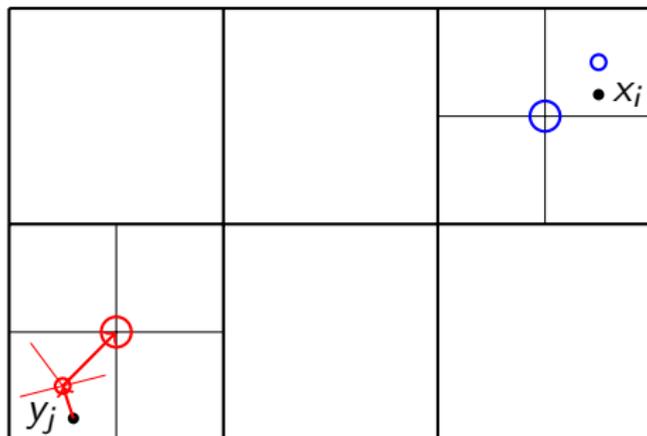


Another View

Using a low rank approximation,

$$K(\mathbf{r}_i + \mathbf{r}_0 + \mathbf{r}_j) = \sum_{p=1}^r \sum_{q=1}^{r'} U_p(\mathbf{r}_i) \tilde{K}_{pq}(\mathbf{r}_0) V_q^T(\mathbf{r}_j) + \varepsilon$$

We define “multipole” and “local” fields at the nodes of a tree:

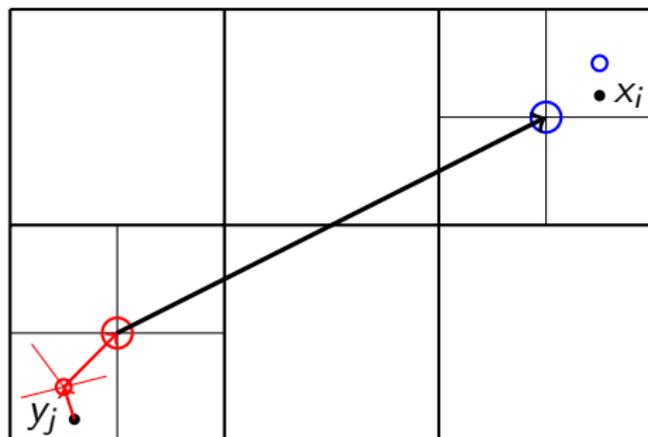


Another View

Using a low rank approximation,

$$K(\mathbf{r}_i + \mathbf{r}_0 + \mathbf{r}_j) = \sum_{p=1}^r \sum_{q=1}^{r'} U_p(\mathbf{r}_i) \tilde{K}_{pq}(\mathbf{r}_0) V_q^T(\mathbf{r}_j) + \varepsilon$$

We define “multipole” and “local” fields at the nodes of a tree:

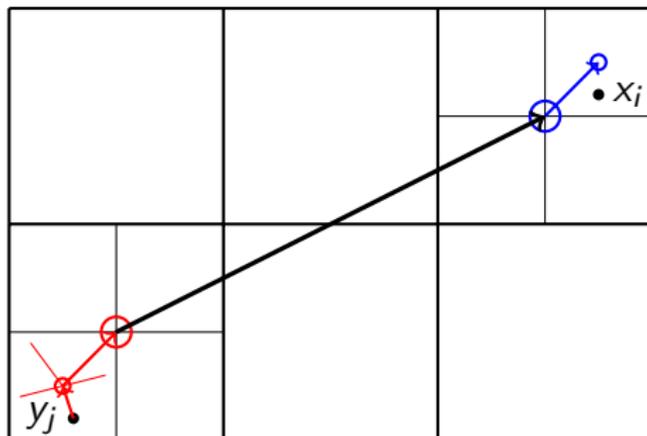


Another View

Using a low rank approximation,

$$K(\mathbf{r}_i + \mathbf{r}_0 + \mathbf{r}_j) = \sum_{p=1}^r \sum_{q=1}^{r'} U_p(\mathbf{r}_i) \tilde{K}_{pq}(\mathbf{r}_0) V_q^T(\mathbf{r}_j) + \varepsilon$$

We define “multipole” and “local” fields at the nodes of a tree:

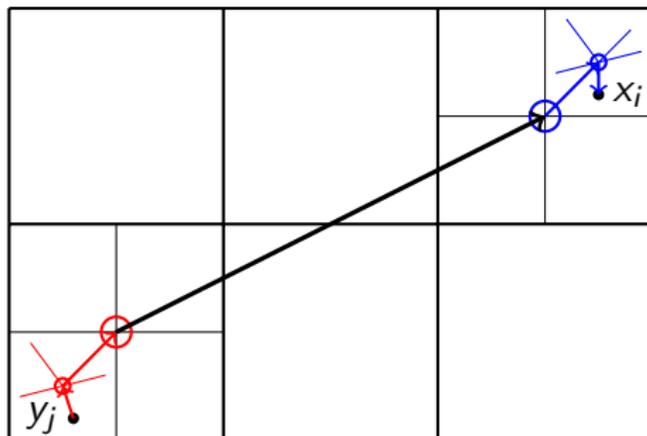


Another View

Using a low rank approximation,

$$K(\mathbf{r}_i + \mathbf{r}_0 + \mathbf{r}_j) = \sum_{p=1}^r \sum_{q=1}^{r'} U_p(\mathbf{r}_i) \tilde{K}_{pq}(\mathbf{r}_0) V_q^T(\mathbf{r}_j) + \varepsilon$$

We define “multipole” and “local” fields at the nodes of a tree:



Operators

$$\mathbf{K}_{IJ} = U_{C(I)} \tilde{U}_I \tilde{K}_{IJ} \tilde{V}_J^T V_{C(J)}^T$$

Abstraction of each operator:

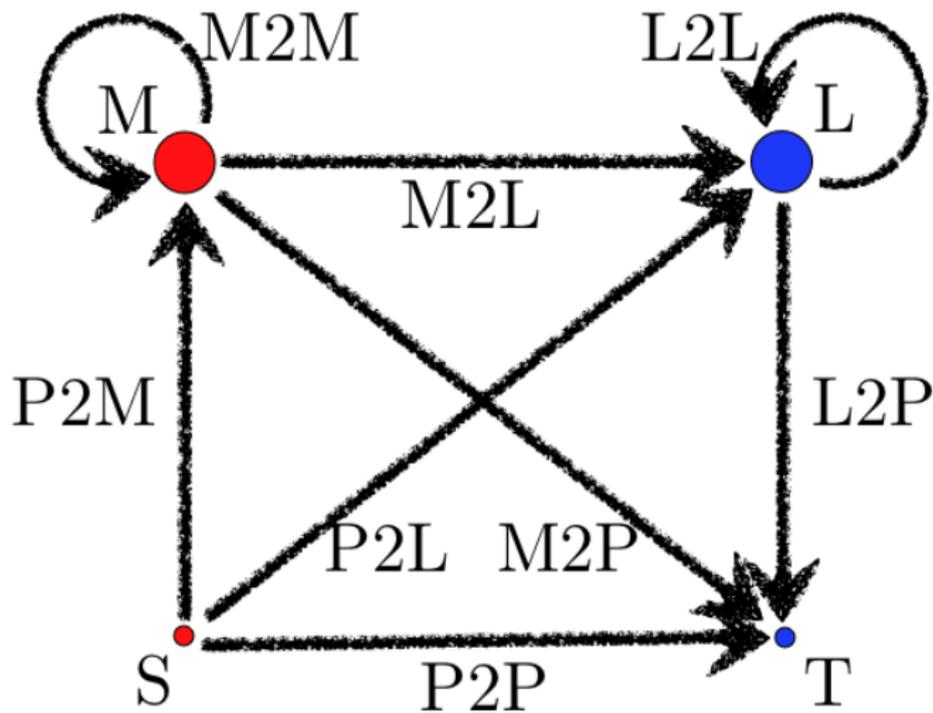
- P2P: $K(t, s)$
- P2M: V^T and/or $\tilde{V}^T V^T$
- M2M: \tilde{V}^T
- M2L: \tilde{K}
- L2L: \tilde{U}
- L2P: U and/or $U\tilde{U}$
- MAC: Boolean function of two clusters
 - Weak - Adjacent clusters are accepted (are low-rank).
 - Strong - Well-separated clusters are accepted (are low-rank).
 - Dynamic - Function of the multipole values.

Useful/Common additions:

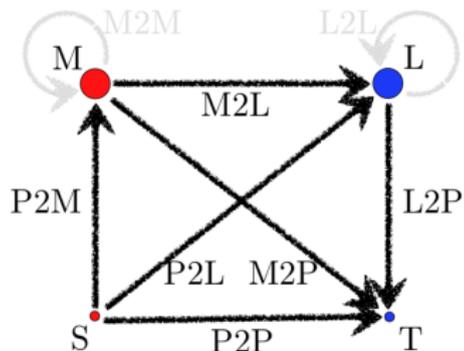
- P2L: $\tilde{K} \tilde{V}^T V^T$
- M2P: $U \tilde{U} \tilde{K}$



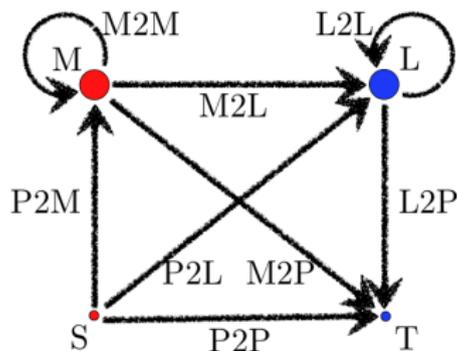
Operator Graph



Special Cases



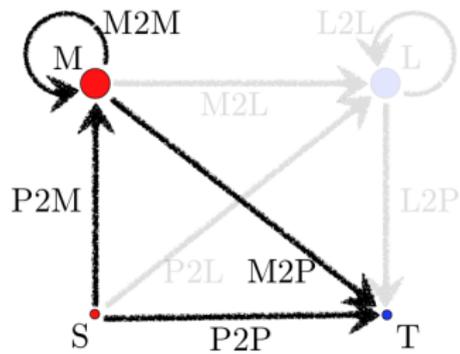
HODLR (Weak MAC)
 \mathcal{H} (Strong MAC)



HSS (Weak MAC)
 \mathcal{H}^2 (Strong MAC)

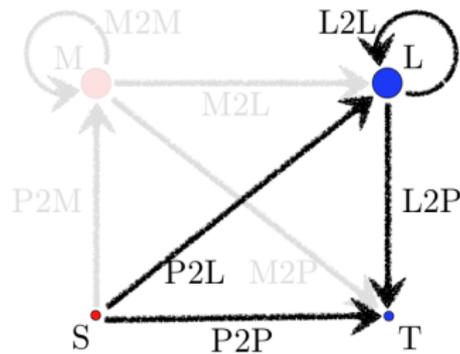


Special Cases



with strong/dynamic MAC.

Cluster-Particle Treecode

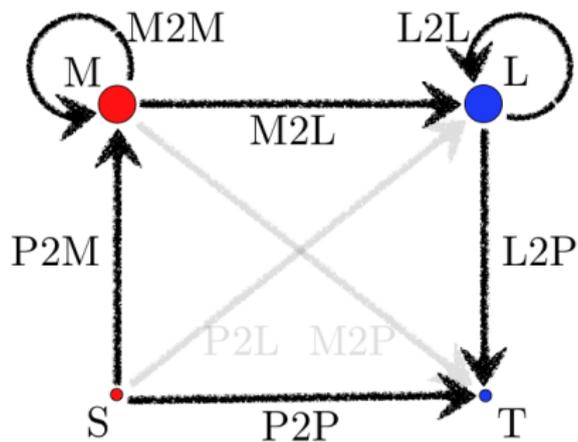


with strong/dynamic MAC.

Particle-Cluster Treecode



Special Cases



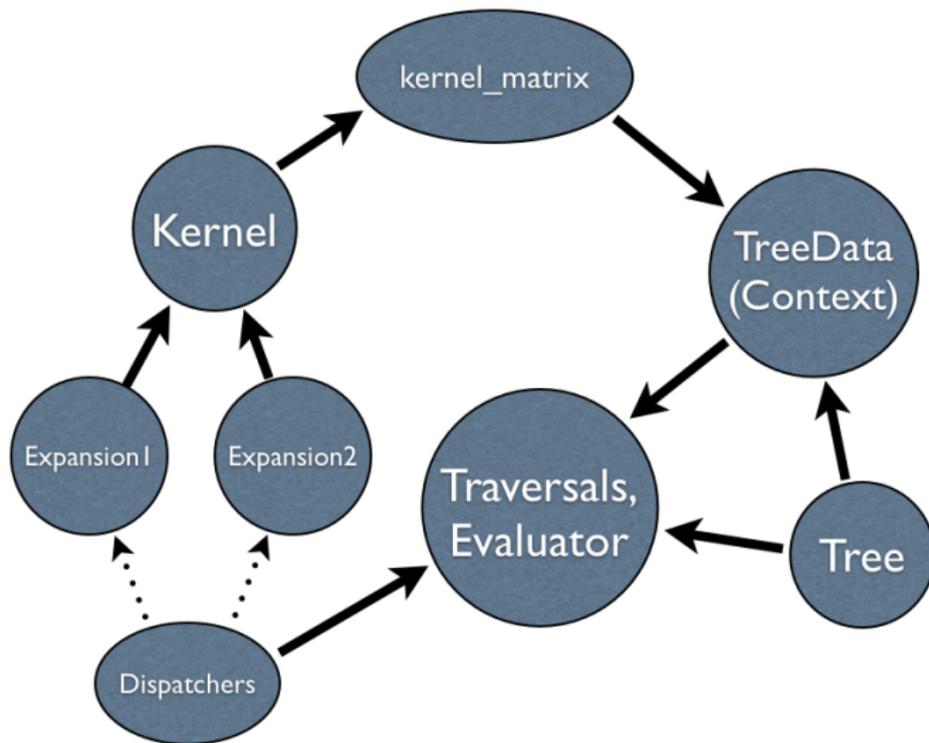
with strong MAC.

Classic FMM.



Code Design

Separate the components of a hierarchical code



MyKernel

```
1 struct MyKernel
2   : public fmmtl::Kernel<MyKernel> {
3     typedef Vec<3,double> source_type;
4     typedef double        charge_type;
5     typedef Vec<3,double> target_type;
6     typedef double        result_type;
7
8     typedef double        kernel_value_type;
9
10    kernel_value_type operator()(const target_type& t,
11                                const source_type& s) const {
12        // Compute  $K(t,s)$ 
13    }
14    /** OPTIONAL! */
15    kernel_value_type transpose(const kernel_value_type& kts) const {
16        // Compute  $K(s,t)$  from  $K(t,s)$  if possible
17    }
18 };
```

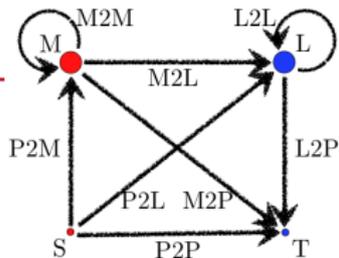


MyExpansion

```
1 struct MyExpansion
2 : public fmmatl::Expansion<MyKernel, MyExpansion> {
3
4     /** For use in defining and building the tree
5     * source_type and target_type must be convertible to a point_type:
6     * static_cast<point_type>(source_type)
7     * static_cast<point_type>(target_type) */
8     static constexpr unsigned dimension = 3;
9     typedef Vec<dimension, double> point_type;
10
11     typedef std::vector<double> multipole_type;
12     typedef std::vector<double> local_type;
13
14     ...
15
16     void P2M(const source_type& s, const charge_type& c,
17             const point_type& center, multipole_type& M) const {
18         // Compute M += V^T(s) * c
19     }
20
21     ...
22 };
```



Kernel/Expansion



- Kernel stored in `.kern` file.
 - May be templated and/or contain state (e.g. κ).
 - Multiple architecture support: CPU/GPU compilation.
 - Optional transpose and vectorized P2P methods.
- Expansion stored in `.hpp` file.
 - Implements pathway(s) through the operator graph.
 - Optional vectorized P2X and X2P methods.
- Statically detect available computational pathways (SFINAE).



Multipole Acceptance

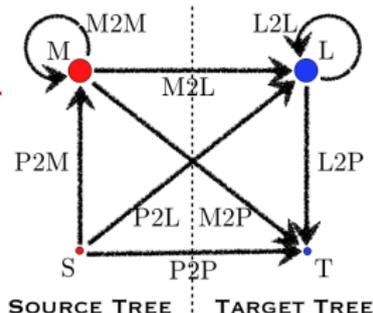
In development...

```
1  /** Strong/Weak cluster interactions - well-separated criteria */
2  bool MAC(const point_type& tbox_extents, const point_type& tbox_center,
3           const point_type& sbox_extents, const point_type& sbox_center) const {
4  }
5
6  /** Dynamic criteria */
7  bool MAC(const multipole_type& M, const local_type& L) const {
8  }
```



Tree(s)

- $\{s_j\} \equiv \{t_i\}$
 - Single Tree
- $\{s_j\} \neq \{t_i\}$
 - Dual Tree



- Whether single tree or dual tree, a TreeContext shall provide

```
1  concept TreeContext {
2      typename source_tree_type;
3      typename target_tree_type;
4
5      [const] source_tree_type& source_tree() [const];
6      [const] target_tree_type& target_tree() [const];
7
8      permuted_iterator source_tree_permute(data_iter di,
9                                          source_body_iter sbi) const;
10     permuted_iterator target_tree_permute(data_iter di,
11                                         target_body_iter sbi) const;
12 };
```

- Traversals and evaluators conceptually work with dual trees.



Tree data

- Data to be associated with (not implemented with!) a Tree.
 - Facilitates optimization (data structures/parallelization).
 - Promotes code independence.
- Source Box data: M
- Target Box data: L
- Source Body data: source, charge
- Target Body data: target, result

- Stored independently, but associated with boxes and bodies:

```
1  concept DataContext {
2      multipole_type& multipole(    const source_box_type& sbox);
3      local_type&     local(       const target_box_type& tbox);
4      ...
5      source_iterator source_begin(const source_box_type& sbox) const;
6      source_iterator source_end(  const source_box_type& sbox) const;
7      ...
8      result_iterator result_begin(const target_box_type& tbox) const;
9      result_iterator result_end(  const target_box_type& tbox) const;
10     ...
11 };
```



e.g.

```
1 #include <fmmtl/KernelMatrix.hpp>
2 #include "MyExpansion.hpp"
3
4 int main() {
5     typedef MyExpansion expansion;
6     expansion K(...);           // Expansion order, error target, etc
7     ...
8     std::vector<source_type> s = ...
9     std::vector<charge_type> c = ...
10
11     std::vector<target_type> t = ...
12
13     fmmtl::KernelMatrix<expansion> M = K(t,s);           // Construct
14     fmmtl::set_options(M, opts);                       // Set options?
15     ...
16     std::vector<result_type> r_aprx = M * c;           // FMM/Treecode
17     std::vector<result_type> r_exact = fmmtl::direct(M * c); //  $O(N^2)$ 
18     ...
19     M.expansion().set_something(...); // Mutate the expansion
20     ...
21 }
```

Examples in FMMTL.



Applications

- `KernelSkeleton.kern` – Documentation
- `UnitKernel.kern`, `ExpKernel.kern` – Testing
- `Laplace.kern` – single layer, double layer, both, BEM
 - `LaplaceSpherical.hpp`
 - `LaplaceCartesian.hpp`
 - `LaplaceSphericalBEM.hpp`
- `Yukawa.kern` – single layer, double layer, both, BEM
 - `YukawaCartesian.hpp`
 - `YukawaSpherical.hpp`
 - `YukawaCartesianBEM.hpp`
- `Stokes.kern` – single layer, double layer, both, BEM
 - `StokesSpherical.hpp`
 - `StokesSphericalBEM.hpp`
- `Helmholtz.kern` – single layer
 - `HelmholtzFourier.hpp`
- Generalized
 - `KIExpansion.hpp` – KIFMM (in progress)
 - `bbExpansion.hpp` – bbFMM (wanted!)



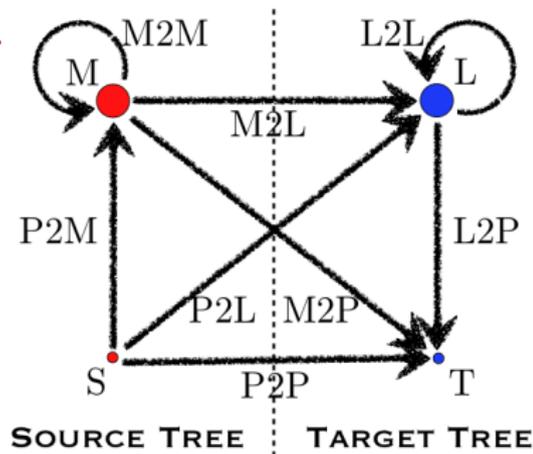
Current Work

- Comparison and research into Expansions and Traversals.
 - Fair comparison between bbFMM and KIFMM.
 - Scheduling and parallelization research.
 - Development/testing of new expansions [Harvard, UMass]
- Collaborations:
 - Hans Johnson, UMass - Development/testing of expansions.
 - Lorena Barba, BU - BEM with inexact GMRES: Laplace, Stokes.
 - Rio Yokota, KAUST - Fast preconditioners using FMM.
 - L Mahadevan, Harvard - Integral polyharmonic PDE solvers.
 - Brian Quaife, UT Austin - Quaternionic FMM expansions.
- Applications
 - Stokes BEM (Boston University)
 - Biharmonic BEM (Harvard University)
 - Laplace NBody (Harvard University)
 - Wanted: more!



Open Problems

- Swapping operators
 - $M2L \Rightarrow M2P$
 - $M2L \Rightarrow P2L$
 - $M2L \Rightarrow P2P$
- Balancing P2P and M2L.
- Scheduling.



Future Work

- Collect more Kernels/Expansions!
 - If you have one (especially if it's odd...) contact me!
- Continue development:
 - Parallelization – MPI, more OpenMP, more GPU.
 - Classic matrix interface
 - Sugar + interop with (e.g.) Eigen, uBlas, ViennaCL, Python
- Million-dollar scheduling/autotuning questions:
 - Given a subset of the operations
 $\{P2P, P2M, P2L, M2M, M2L, M2P, L2L, L2P\}$ (Statically),
(+potentially cost estimates/target error) and a
 MAC (Static/Dynamic),
and a
 $[Dual]$ Tree (Dynamically),
find an optimal dispatch such that all source/target pairs
interact.

