Sandia
National
Laboratories

MueLu User's Guide

Luc Berger-Vergiat, Christian A. Glusa, Jonathan J. Hu, Matthias Mayr, Andrey Prokopenko, Christopher M. Siefert, Raymond S. Tuminaro, Tobias A. Wiesner
Approved for public release; further dissemination unlimited.

MueLu User's Guide

Luc Berger-Vergiat
Computational Mathematics
Sandia National Laboratories
Mailstop 1320, P.O. Box 5800
Albuquerque, NM 87185-1322
lberge@sandia.gov

Christian A. Glusa
Scalable Algorithms
Sandia National Laboratories
Mailstop 1318, P.O. Box 5800
Albuquerque, NM 87185-1322
caglusa@sandia.gov

Jonathan J. Hu
Scalable Algorithms
Sandia National Laboratories
Mailstop 9159, P.O. Box 0969
Livermore, CA 94551-0969
jhu@sandia.gov

Matthias Mayr
Institute for Mathematics
& Computer-Based Simulation
University of the Bundeswehr Munich
Werner-Heisenberg-Weg 39
85577 Neubiberg, Germany
matthias.mayr@unibw.de

Andrey Prokopenko
Oak Ridge National Laboratory
P.O. Box 2008
Bldg 5700, MS 6164
Oak Ridge, TN 37831

Christopher M. Siefert
Scalable Algorithms
Sandia National Laboratories
Mailstop 1322, P.O. Box 5800
Albuquerque, NM 87185-1322
csiefer@sandia.gov

Raymond S. Tuminaro
Computational Mathematics
Sandia National Laboratories
Mailstop 9159, P.O. Box 0969
Livermore, CA 94551-0969
rstumin@sandia.gov

Tobias Wiesner
Leica Geosystems AG
Heinrich-Wild-Strasse 201
9435 Heerbrugg, Switzerland
tobias.wiesner@leica-geosystems.com

Abstract

This is the official user guide for MUELU multigrid library in Trilinos version 12.17 (Dev) . This guide provides an overview of MUELU, its capabilities, and instructions for new users who want to start using MUELU with a minimum of effort. Detailed information is given on how to drive MUELU through its XML interface. Links to more advanced use cases are given. This guide gives information on how to achieve good parallel performance, as well as how to introduce new algorithms. Finally, readers will find a comprehensive listing of available MUELU options. *Any options not documented in this manual should be considered strictly experimental.*

Acknowledgment

Contents

Appendix

List of Figures

List of Tables

Chapter 1

Introduction

This guide gives an overview of MUELU's capabilities. If you are looking for a tutorial, please refer to the MUELU tutorial in `muelu/doc/Tutorial` (see also [19]). New users should start with §3. It strives to give the new user all the information he/she might need to begin using MUELU quickly. Users interested in performance, especially in parallel context, should refer to §4. Users looking for a particular option should consult §5, containing a complete set of supported options in MUELU.

If you find any errors or omissions in this guide, have comments or suggestions, or would like to contribute to MUELU development, please contact the MUELU users list, or developers list or open an issue on the Trilinos github repository.

Chapter 2

Multigrid background

Here we provide a brief multigrid introduction (see [6] or [16] for more information). A multigrid solver tries to approximate the original problem of interest with a sequence of smaller (*coarser*) problems. The solutions from the coarser problems are combined in order to accelerate convergence of the original (*fine*) problem on the finest grid. A simple multilevel iteration is illustrated in Algorithm 1.

Algorithm 1 V-cycle multigrid with N levels to solve $Ax = b$.

$A_0 = A$
function MULTILEVEL(A_k, b, u, k)
    // Solve A_k u = b (k is current grid level)
    $u = S_m^1(A_k, b, u)$
    if $(k \neq \mathbf{N} - \mathbf{1})$ **then**
        $P_k =$ determine_interpolant(A_k)
        $R_k =$ determine_restrictor(A_k)
        $\widehat{r}_{k+1} = R_k(b - A_k u)$
        $A_{k+1} = R_k A_k P_k$
        $v = 0$
        MULTILEVEL($\widehat{A}_{k+1}$, $\widehat{r}_{k+1}$, v, $k+1$)
        $u = u + P_k v$
        $u = S_m^2(A_k, b, u)$
    end if
end function

In the multigrid iteration in Algorithm 1, the $S_m^1()$'s and $S_m^2()$'s are called *pre-smoothers* and *post-smoothers*. They are approximate solvers (e.g., symmetric Gauss-Seidel), with the subscript m denoting the number of applications of the approximate solution method. The purpose of a smoother is to quickly reduce certain error modes in the approximate solution on a level k. For symmetric problems, the pre- and post-smoothers should be chosen to maintain symmetry (e.g., forward Gauss-Seidel for the pre-smoother and backward Gauss-Seidel for the post-smoother). The P_k's are *interpolation* matrices that transfer solutions from coarse levels to finer levels. The R_k's are *restriction* matrices that restrict a fine level residual to a coarser level. In a geometric multigrid, P_k's and R_k's are determined by the application, whereas in an algebraic multigrid they are automatically generated. For symmetric problems, typically $R_k = P_k^T$. For nonsymmetric problems, this is not necessarily true. The A_k's are the coarse level problems, and are generated through

a Galerkin (triple matrix) product.

Please note that the algebraic multigrid algorithms implemented in MUELU generate the grid transfers P_k automatically and the coarse problems A_k via a sparse triple matrix product. TRILINOS provides a wide selection of smoothers and direct solvers for use in MUELU through the IFPACK, IFPACK2, AMESOS, and AMESOS2 packages (see §5).

Chapter 3

Getting Started

This section is meant to get you using MUELU as quickly as possible. §3.1 gives a summary of MUELU's design. §3.2 lists MUELU's dependencies on other TRILINOS libraries and provides a sample cmake configuration line. Finally, code examples using the XML interface are given in §3.3.

3.1 Overview of MUELU

MUELU is an extensible algebraic multigrid (AMG) library that is part of the TRILINOS project. MUELU works with EPETRA (32-bit version) and TPETRA matrix types. The library is written in C++ and allows for different ordinal (index) and scalar types. MUELU is designed to be efficient on many different computer architectures, from workstations to supercomputers, relying on "MPI+X" principle, where "X" can be threading, CUDA, or any other back-end provided by the KOKKOS package.

MUELU provides a number of different multigrid algorithms:

1. smoothed aggregation AMG (for Poisson-like and elasticity problems);

2. Petrov-Galerkin aggregation AMG (for convection-diffusion problems);

3. energy-minimizing AMG;

4. aggregation-based AMG for problems arising from the eddy current formulation of Maxwell's equations.

MUELU's software design allows for the rapid introduction of new multigrid algorithms. The most important features of MUELU can be summarized as:

Easy-to-use interface

MUELU has a user-friendly parameter input deck which covers most important use cases. Reasonable defaults are provided for common problem types (see Table 5.2).

Modern object-oriented software architecture

MUELU is written completely in C++ as a modular object-oriented multigrid framework, which provides flexibility to combine and reuse existing components to develop novel multigrid methods.

Extensibility

Due to its flexible design, MUELU is an excellent toolkit for research on novel multigrid concepts. Experienced multigrid users have full access to the underlying framework through an advanced XML based interface. Expert users may use and extend the C++ API directly.

Integration with TRILINOS library

As a package of TRILINOS, MUELU is well integrated into the TRILINOS environment. MUELU can be used with either the TPETRA or EPETRA (32-bit) linear algebra stack. It is templated on the local index, global index, scalar, and compute node types. This makes MUELU ready for future developments.

Broad range of supported platforms

MUELU runs on wide variety of architectures, from desktop workstations to parallel Linux clusters and supercomputers ([10]).

Open source

MUELU is freely available through a simplified BSD license (see Appendix A).

3.2 Configuration and Build

MUELU has been compiled successfully under Linux with the following C++ compilers: GNU versions 4.1 and later, Intel versions 12.1/13.1, and clang versions 3.2 and later. In the future, we recommend using compilers supporting C++11 standard.

3.2.1 Dependencies

Required Dependencies

MUELU requires that TEUCHOS and either EPETRA/IFPACK or TPETRA/IFPACK2 are enabled.

Recommended Dependencies

We strongly recommend that you enable the following TRILINOS libraries along with MUELU:

- EPETRA stack: AZTECOO, EPETRA, AMESOS, IFPACK, ISORROPIA, GALERI, ZOLTAN;

- TPETRA stack: AMESOS2, BELOS, GALERI, IFPACK2, TPETRA, ZOLTAN2.

Tutorial Dependencies

In order to run the MUELU Tutorial [19] located in `muelu/doc/Tutorial`, MUELU must be configured with the following dependencies enabled:

AZTECOO, AMESOS, AMESOS2, BELOS, EPETRA, IFPACK, IFPACK2, ISORROPIA, GALERI, TPETRA, ZOLTAN, ZOLTAN2.

☛ Note that the MUELU tutorial [19] comes with a VirtualBox image with a pre-installed Linux and TRILINOS. In this way, a user can immediately begin experimenting with MUELU without having to install the TRILINOS libraries. Therefore, it is an ideal starting point to get in touch with MUELU.

Complete List of Direct Dependencies

Dependency	Required		Optional	
	Library	Testing	Library	Testing
AMESOS			×	×
AMESOS2			×	×
AZTECOO				×
BELOS				×
EPETRA			×	×
IFPACK			×	×
IFPACK2			×	×
ISORROPIA			×	×
GALERI				×
KOKKOSCLASSIC			×	
TEUCHOS	×	×		
TPETRA			×	×
XPETRA	×	×		
ZOLTAN			×	×
ZOLTAN2			×	×
Boost			×	
BLAS	×	×		
LAPACK	×	×		
MPI			×	×

Table 3.1. MUELU's required and optional dependencies, subdivided by whether a dependency is that of the MUELU library itself (*Library*), or of some MUELU test (*Testing*).

Table 3.1 lists the dependencies of MUELU, both required and optional. If an optional dependency is not present, the tests requiring that dependency are not built.

☛ AMESOS/AMESOS2 are necessary if one wants to use a sparse direct solve on the coarsest level. ZOLTAN/ZOLTAN2 are necessary if one wants to use matrix rebalancing in parallel runs (see §4). AZTECOO/BELOS are necessary if one wants to test MUELU as a preconditioner instead of a solver.

☛ MUELU has also been successfully tested with SuperLU 4.1 and SuperLU 4.2.

☛ Some packages that MUELU depends on may come with additional requirements for third party libraries, which are not listed here as explicit dependencies of MUELU. It is highly recommended to install ParMetis 3.1.1 or newer for ZOLTAN, ParMetis 4.0.x for ZOLTAN2, and SuperLU 4.1 or newer for AMESOS/AMESOS2.

3.2.2 Configuration

The preferred way to configure and build MUELU is to do that outside of the source directory. Here we provide a sample configure script that will enable MUELU and all of its optional dependencies:

```
export TRILINOS_HOME=/path/to/your/Trilinos/source/directory
cmake \
    -D BUILD_SHARED_LIBS:BOOL=ON \
    -D CMAKE_BUILD_TYPE:STRING="RELEASE" \
    -D CMAKE_CXX_FLAGS:STRING="-g" \
    -D Trilinos_ENABLE_EXPLICIT_INSTANTIATION:BOOL=ON \
    -D Trilinos_ENABLE_TESTS:BOOL=OFF \
    -D Trilinos_ENABLE_EXAMPLES:BOOL=OFF \
    -D Trilinos_ENABLE_MueLu:BOOL=ON \
    -D  MueLu_ENABLE_TESTS:STRING=ON \
    -D  MueLu_ENABLE_EXAMPLES:STRING=ON \
    -D  MueLu_ENABLE_Kokkos_Refactor:STRING=ON \
    -D  MueLu_ENABLE_Kokkos_Refactor_Use_By_Default:STRING=ON \
    -D TPL_ENABLE_BLAS:BOOL=ON \
    -D TPL_ENABLE_MPI:BOOL=ON \
  ${TRILINOS_HOME}
```

More configure examples can be found in `Trilinos/sampleScripts`. For more information on configuring, see the TRILINOS CMake Quickstart guide [1].

3.3 Examples in code

The most commonly used scenario involving MUELU is using a multigrid preconditioner inside an iterative linear solver. In TRILINOS, a user has a choice between EPETRA and TPETRA

for the underlying linear algebra library. Important Krylov subspace methods (such as preconditioned CG and GMRES) are provided in TRILINOS packages AZTECOO (EPETRA) and BELOS (EPETRA/TPETRA).

At this point, we assume that the reader is comfortable with TEUCHOS referenced-counted pointers (RCPs) for memory management (an introduction to RCPs can be found in [3]) and the `Teuchos::ParameterList` class [15].

3.3.1 MUELU as a preconditioner within BELOS

The following code shows the basic steps of how to use a MUELU multigrid preconditioner with TPETRA linear algebra library and with a linear solver from BELOS. To keep the example short and clear, we skip the template parameters and focus on the algorithmic outline of setting up a linear solver. For further details, a user may refer to the `examples` and `test` directories.

First, we create the MUELU multigrid preconditioner. It can be done in a few ways. For instance, multigrid parameters can be read from an XML file (e.g., *mueluOptions.xml* in the example below).

```
Teuchos::RCP<Tpetra::CrsMatrix<> > A;
// create A here ...
std::string optionsFile = "mueluOptions.xml";
Teuchos::RCP<MueLu::TpetraOperator> mueLuPreconditioner =
   MueLu::CreateTpetraPreconditioner(A, optionsFile);
```

The XML file contains multigrid options. A typical file with MUELU parameters looks like the following.

```
<ParameterList name="MueLu">

  <Parameter name="verbosity" type="string" value="low"/>

  <Parameter name="max levels" type="int" value="3"/>
  <Parameter name="coarse: max size" type="int" value="10"/>

  <Parameter name="multigrid algorithm" type="string" value="sa"/>

  <!-- Damped Jacobi smoothing -->
  <Parameter name="smoother: type" type="string" value="RELAXATION"/>
  <ParameterList name="smoother: params">
    <Parameter name="relaxation: type" type="string" value="Jacobi"/>
    <Parameter name="relaxation: sweeps" type="int" value="1"/>
    <Parameter name="relaxation: damping factor" type="double" value="0.9"/>
  </ParameterList>

  <!-- Aggregation -->
```

```xml
  <Parameter name="aggregation: type" type="string" value="uncoupled"/>
  <Parameter name="aggregation: min agg size" type="int" value="3"/>
  <Parameter name="aggregation: max agg size" type="int" value="9"/>

</ParameterList>
```

It defines a three level smoothed aggregation multigrid algorithm. The aggregation size is between three and nine(2D)/27(3D) nodes. One sweep with a damped Jacobi method is used as a level smoother. By default, a direct solver is applied on the coarsest level. A complete list of available parameters and valid parameter choices is given in §5 of this User's Guide.

Users can also construct a multigrid preconditioner using a provided `ParameterList` without accessing any files in the following manner.

```cpp
Teuchos::RCP<Tpetra::CrsMatrix<> > A;
// create A here ...
Teuchos::ParameterList paramList;
paramList.set("verbosity", "low");
paramList.set("max levels", 3);
paramList.set("coarse: max size", 10);
paramList.set("multigrid algorithm", "sa");
// ...
Teuchos::RCP<MueLu::TpetraOperator> mueLuPreconditioner =
   MueLu::CreateTpetraPreconditioner(A, paramList);
```

Besides the linear operator *A*, we also need an initial guess vector for the solution *X* and a right hand side vector *B* for solving a linear system.

```cpp
Teuchos::RCP<const Tpetra::Map<> > map = A->getDomainMap();

// Create initial vectors
Teuchos::RCP<Tpetra::MultiVector<> > B, X;
X = Teuchos::rcp( new Tpetra::MultiVector<>(map,numrhs) );
Belos::MultiVecTraits<>::MvRandom( *X );
B = Teuchos::rcp( new Tpetra::MultiVector<>(map,numrhs) );
Belos::OperatorTraits<>::Apply( *A, *X, *B );
Belos::MultiVecTraits<>::MvInit( *X, 0.0 );
```

To generate a dummy example, the above code first declares two vectors. Then, a right hand side vector is calculated as the matrix-vector product of a random vector with the operator *A*. Finally, an initial guess is initialized with zeros.

Then, one can define a `Belos::LinearProblem` object where the `mueLuPreconditioner` is used for left preconditioning

```cpp
Belos::LinearProblem<> problem( A, X, B );
problem->setLeftPrec(mueLuPreconditioner);
```

```
bool set = problem.setProblem();
```

Next, we set up a BELOS solver using some basic parameters

```
Teuchos::ParameterList belosList;
belosList.set( "Block Size", 1 );
belosList.set( "Use Single Reduction", true );
belosList.set( "Maximum Iterations", 100 );
belosList.set( "Convergence Tolerance", 1e-10 );
belosList.set( "Output Frequency", 1 );
belosList.set( "Verbosity", Belos::TimingDetails + Belos::FinalSummary );


Belos::BlockCGSolMgr<> solver( rcp(&problem,false), rcp(&belosList,false) );
```

Finally, we solve the system.

```
Belos::ReturnType ret = solver.solve();
```

3.3.2   MUELU as a preconditioner for AZTECOO

For EPETRA, users have two library options: BELOS (recommended) and AZTECOO. AZTECOO and BELOS both provide fast and mature implementations of common iterative Krylov linear solvers. BELOS has additional capabilities, such as Krylov subspace recycling and "tall skinny QR".

Constructing a MUELU preconditioner for Epetra operators is done in a similar manner to Tpetra.

```
Teuchos::RCP<Epetra_CrsMatrix> A;
// create A here ...
Teuchos::RCP<MueLu::EpetraOperator> mueLuPreconditioner;
std::string optionsFile = "mueluOptions.xml";
mueLuPreconditioner = MueLu::CreateEpetraPreconditioner(A, optionsFile);
```

MUELU parameters are generally Epetra/Tpetra agnostic, thus the XML parameter file could be the same as §3.3.1.

Furthermore, we assume that a right hand side vector and a solution vector with the initial guess are defined.

```
Teuchos::RCP<const Epetra_Map> map = A->DomainMap();
Teuchos::RCP<Epetra_Vector> B = Teuchos::rcp(new Epetra_Vector(map));
Teuchos::RCP<Epetra_Vector> X = Teuchos::rcp(new Epetra_Vector(map));
X->PutScalar(0.0);
```

Then, an Epetra_LinearProblem can be defined.

23

```
Epetra_LinearProblem epetraProblem(A.get(), X.get(), B.get());
```

The following code constructs an AZTECOO CG solver.

```
AztecOO aztecSolver(epetraProblem);
aztecSolver.SetAztecOption(AZ_solver, AZ_cg);
aztecSolver.SetPrecOperator(mueLuPreconditioner.get());
```

Finally, the linear system is solved.

```
int maxIts = 100;
double tol = 1e-10;
aztecSolver.Iterate(maxIts, tol);
```

3.3.3 MUELU's structured algorithms

Some users might use structured meshes to discretize their problems. In such cases it can be advantageous to use the structured grid algorithms provided in MUELU. To use these algorithms the user has to provide extra information to MUELU such as the number of spatial dimensions in the problem and the number of nodes in each direction on the local rank. As demonstrated in the code bellow MUELU expect these additional inputs to be stored in a sublist called "user data".

```
const std::string userName = "user data";
Teuchos::ParameterList& userParamList = paramList.sublist(userName);
userParamList.set<int>("int numDimensions", numDimensions);
userParamList.set<Teuchos::Array<LO> >("Array<LO> lNodesPerDim", lNodesPerDim);
userParamList.set<RCP<RealValuedMultiVector> >("Coordinates", coordinates);
H = MueLu::CreateXpetraPreconditioner(A, paramList, paramList);
```

Full examples demonstrating the structured capabilities of MUELU can be found in the TRILINOS source directories

- `packages/muelu/test/structured,`
- `packages/trilinoscouplings/examples/scaling.`

3.3.4 MUELU's Maxwell solver

MUELU can be used to solve Maxwell's equations in eddy current formulation which can be written as

$$\nabla \times \left(\alpha \nabla \times \vec{E} \right) + \beta \vec{E} = \vec{f}, \tag{3.1}$$

where $\vec{E}$ is the unknown electric field, α and β are material parameters, and $\vec{f}$ is the known right-hand side. In order to deal with the large nullspace of the curl-curl operator a specialized multigrid approach is required. For a detailed description of the solver see [5].

A preconditioner for equation 3.1 can be constructed as follows:

```
RCP<Matrix> SM_Matrix = ... ; \\ Edge-mass + curl-curl
RCP<Matrix> D0_Matrix = ... ; \\ Discrete gradient matrix
RCP<Matrix> M0inv_Matrix = ... ; \\ Approximate inverse of node-mass matrix
    with weight 1/alpha
RCP<Matrix> M1_Matrix = ... ; \\ Edge-mass matrix with constant weight 1
RCP<MultiVector> coords = ...; \\ Nodal coordinates
Teuchos::ParameterList params = ...; \\ Parameters

RCP<MueLu::RefMaxwell> preconditioner
= rcp( new MueLu::RefMaxwell(SM_Matrix, D0_Matrix, M0inv_Matrix,
        M1_Matrix, Teuchos::null, coords, params) );
```

An example set of parameters is given below:

```
<ParameterList name="MueLu">

<Parameter name="refmaxwell: mode" type="string" value="additive"/>

<Parameter name="smoother: type" type="string" value="RELAXATION"/>
<ParameterList name="smoother: params">
  <Parameter name="relaxation: type" type="string" value="Symmetric Gauss-Seidel
      "/>
  <Parameter name="relaxation: sweeps" type="int" value="2"/>
</ParameterList>

<ParameterList name="refmaxwell: 11list">
  <Parameter name="number of equations" type="int" value="3"/>
  <Parameter name="aggregation: type" type="string" value="uncoupled"/>
  <Parameter name="coarse: max size"  type="int"   value="2500"/>
  <Parameter name="smoother: type" type="string" value="RELAXATION"/>
  <ParameterList name="smoother: params">
    <Parameter name="relaxation: type" type="string" value="Symmetric Gauss-
        Seidel"/>
    <Parameter name="relaxation: sweeps" type="int" value="2"/>
  </ParameterList>
</ParameterList>

<ParameterList name="refmaxwell: 22list">
  <Parameter name="aggregation: type" type="string" value="uncoupled"/>
  <Parameter name="coarse: max size"  type="int"   value="2500"/>
  <Parameter name="smoother: type" type="string" value="RELAXATION"/>
  <ParameterList name="smoother: params">
```

```
    <Parameter name="relaxation: type" type="string" value="Symmetric Gauss-
        Seidel"/>
    <Parameter name="relaxation: sweeps" type="int" value="2"/>
  </ParameterList>
 </ParameterList>

</ParameterList>
```

Further examples of how to use MUELU to solve Maxwell's equations can be found in the TRILI-NOS source directories

- `packages/muelu/test/maxwell`,
- `packages/panzer/mini-em/example/BlockPrec` and
- `packages/trilinoscouplings/examples/scaling`.

3.3.5 Further remarks

This section is only meant to give a brief introduction on how to use MUELU as a preconditioner within the TRILINOS packages for iterative solvers. There are other, more complicated, ways to use MUELU as a preconditioner for BELOS and AZTECOO through the XPETRA interface. Of course, MUELU can also work as standalone multigrid solver. For more information on these topics, the reader may refer to the examples and tests in the MUELU source directory (`packages/muelu/example` and `packages/muelu/test`) and in the trilinosCouplings source directory (`packages/trilinosCouplings`), as well as to the MUELU tutorial [19]. For in-depth details of MUELU applied to multiphysics problems, please see [18].

Chapter 4

Performance tips

In practice, it can be very challenging to find an appropriate set of multigrid parameters for a specific problem, especially if few details are known about the underlying linear system. In this Chapter, we provide some advice for improving multigrid performance.

☞ For optimizing multigrid parameters, it is highly recommended to set the verbosity to `high` or `extreme` for MUELU to output more information (e.g., for the effect of the chosen parameters to the aggregation and coarsening process).

Some general advice:

- Choose appropriate iterative linear solver (e.g., GMRES for non-symmetric problems). If available, set options to perform as few all-reduces as possible. (E.g. Use `Single Reduction` in BELOS.)

- Start with the recommended settings for particular problem types. See Table 5.2.

- Choose reasonable basic multigrid parameters (see §5.3), including maximum number of multigrid levels (`max levels`) and maximum allowed coarse size of the problem (`coarse: max size`). Take fine level problem size and sparsity pattern into account for a reasonable choice of these parameters.

- Select an appropriate transfer operator strategy (see §5.7). For symmetric problems, you should start with smoothed aggregation multigrid. For non-symmetric problems, a Petrov-Galerkin smoothed aggregation method is a good starting point, though smoothed aggregation may also perform well.

- Try `unsmoothed` operators instead of smoothed aggregation (`sa`). Scalability in terms of iterations performed will suffer from this, but solution times might go down since the operators are less dense, and less communication is performed.

- Enable implicit restrictor construction (`transpose: use implicit`) for symmetric problems.

- Enable triple matrix products instead of two matrix-matrix products for the construction of coarse operators (`rap: triple product`). This is beneficial as long as the involved operators are not too dense. For `unsmoothed` hierarchies, it is always faster.

- Find good level smoothers (see §5.4). If a problem is symmetric positive definite, choose a smoother with a matrix-vector computational kernel, such as the Chebyshev polynomial smoother. If you are using relaxation smoothers, we recommend starting with optimizing the damping parameter. Once you have found a good damping parameter for your problem, you can increase the number of smoothing iterations.

- Adjust aggregation parameters if you experience bad coarsening ratios (see §5.5). Particularly, try adjusting the minimum (`aggregation: min agg size`) and maximum (`aggregation: max agg size`) aggregation parameters. For a 2D (3D) isotropic problem on a regular mesh, the aggregate size should be about 9 (27) nodes per aggregate.

- Replace a direct solver with an iterative method (`coarse: type`) if your coarse level solution becomes too expensive (see §5.4).

- If on-node parallelism is required, make sure to enable the KOKKOS code path (`use kokkos refactor`). If Gauss-Seidel smoothing is used, switch to multi-threaded Gauss-Seidel (see §5.4).

Some advice for parallel runs include:

1. Enable matrix rebalancing when running in parallel (`repartition: enable`).

2. Use smoothers invariant to the number of processors, such as polynomial smoothing, if possible.

3. Use `uncoupled` aggregation instead of `coupled`, as the latter requires significantly more communication.

4. Adjust rebalancing parameters (see §5.6). Try choosing rebalancing parameters so that you end up with one processor on the coarsest level for the direct solver (this avoids additional communication).

5. If the `multijagged` algorithm from ZOLTAN 2 is used, try setting the premigration option.

6. Enable implicit rebalancing of prolongators and restrictors (`repartition: rebalance P and R`).

Chapter 5

MUELU options

In this section, we report the complete list of MUELU input parameters. It is important to notice, however, that MUELU relies on other TRILINOS packages to provide support for some of its algorithms. For instance, IFPACK/IFPACK2 provide standard smoothers like Jacobi, Gauss-Seidel or Chebyshev, while AMESOS/AMESOS2 provide access to direct solvers. The parameters affecting the behavior of the algorithms in those packages are simply passed by MUELU to a routine from the corresponding library. Please consult corresponding packages' documentation for a full list of supported algorithms and corresponding parameters.

5.1   Using parameters on individual levels

Some of the parameters that affect the preconditioner can in principle be different from level to level. By default, parameters affect all levels in a multigrid hierarchy.

The settings on a particular level can be changed by using level sublists. A level sublist is a `ParameterList` sublist with the name "level XX", where XX is the level number. The parameter names in the sublist do not require any modifications. For example, the following fragment of code

```
<ParameterList name="level 2">
  <Parameter name="smoother: type" type="string" value="CHEBYSHEV"/>
</ParameterList>
```

changes the smoother for level 2 to be a Chebyshev-type polynomial smoother.

5.2   Parameter validation

By default, MUELU validates the input parameter list. A parameter that is misspelled is unknown. A parameter with an incorrect value type is also treated as invalid. Both cases will cause an exception to be thrown and execution to halt.

☞ Spaces are important within a parameter's name. Please separate words by just one space, and make sure there are no leading or trailing spaces.

The option `print initial parameters` prints the initial list given to the interpreter. The option `print unused parameters` prints the list of unused parameters.

5.3   General options

Verbosity level	Description
none	No output
low	Errors, important warnings, and some statistics
medium	Same as low, but with more statistics
high	Errors, all warnings, and all statistics
extreme	Same as high, but also includes output from other packages (*i.e.*, ZOLTAN)

Table 5.1. Verbosity levels.

Problem type	Multigrid algorithm	Block size	Smoother
unknown	–	–	–
Poisson-2D	Smoothed aggregation	1	Chebyshev
Poisson-3D	Smoothed aggregation	1	Chebyshev
Elasticity-2D	Smoothed aggregation	2	Chebyshev
Elasticity-3D	Smoothed aggregation	3	Chebyshev
Poisson-2D-complex	Smoothed aggregation	1	Symmetric Gauss-Seidel
Poisson-3D-complex	Smoothed aggregation	1	Symmetric Gauss-Seidel
Elasticity-2D-complex	Smoothed aggregation	2	Symmetric Gauss-Seidel
Elasticity-3D-complex	Smoothed aggregation	3	Symmetric Gauss-Seidel
ConvectionDiffusion	Petrov-Galerkin AMG	1	Gauss-Seidel
MHD	Unsmoothed aggregation	–	Additive Schwarz method with one level of overlap and ILU(0) as a subdomain solver

Table 5.2. Supported problem types ("–" means not set).

problem: type   [string] Type of problem to be solved. Possible values: see Table 5.2. **Default:** "unknown".

`verbosity`	[string] Control of the amount of printed information. Possible values: see Table 5.1. **Default:** "high".
`number of equations`	[int] Number of PDE equations at each grid node. Only constant block size is considered. **Default:** 1.
`max levels`	[int] Maximum number of levels in a hierarchy. **Default:** 10.
`cycle type`	[string] Multigrid cycle type. Possible values: "V", "W". **Default:** "V".
`problem: symmetric`	[bool] Symmetry of a problem. This setting affects the construction of a restrictor. If set to true, the restrictor is set to be the transpose of a prolongator. If set to false, underlying multigrid algorithm makes the decision. **Default:** true.
`xml parameter file`	[string] An XML file from which to read additional parameters. In case of a conflict, parameters manually set on the list will override parameters in the file. If the string is empty a file will not be read. **Default:** "".
`hierarchy label`	[string] Label for the hierarchy. Is applied to timer labels. **Default:** "".

5.4 Smoothing and coarse solver options

MUELU relies on other TRILINOS packages to provide level smoothers and coarse solvers. IFPACK and IFPACK2 provide standard smoothers (see Table 5.3), and AMESOS and AMESOS2 provide direct solvers (see Table 5.4). Note that it is completely possible to use any level smoother as a direct solver.

MUELU checks parameters `smoother: * type` and `coarse: type` to determine:

- what package to use (i.e., is it a smoother or a direct solver);

- (possibly) transform a smoother type

☛ IFPACK and IFPACK2 use different smoother type names, e.g., "point relaxation stand-alone" vs "RELAXATION". MUELU tries to follow IFPACK2 notation for smoother types. Please consult IFPACK2 documentation [12] for more information.

The parameter lists `smoother: * params` and `coarse: params` are passed directly to the corresponding package without any examination of their content. Please consult the documentation of the corresponding packages for a list of possible values.

By default, MUELU uses one sweep of symmetric Gauss-Seidel for both pre- and post-smoothing, and SuperLU for coarse system solver.

smoother: type	
RELAXATION	Point relaxation smoothers, including Jacobi, Gauss-Seidel, symmetric Gauss-Seidel, multithreaded (coloring-based) Gauss-Seidel, etc. The exact smoother is chosen by specifying `relaxation: type` parameter in the `smoother: params` sublist.
CHEBYSHEV	Chebyshev polynomial smoother.
ILUT, RILUK	Local (processor-based) incomplete factorization methods.

Table 5.3. Commonly used smoothers provided by IF-PACK/IFPACK2. Note that these smoothers can also be used as coarse grid solvers.

coarse: type	AMESOS	AMESOS2	
KLU	x		Default AMESOS solver [8].
KLU2		x	Default AMESOS2 solver [4].
SuperLU	x	x	Third-party serial sparse direct solver [9].
SuperLU_dist	x	x	Third-party parallel sparse direct solver [9].
Umfpack	x		Third-party solver [7].
Mumps	x		Third-party solver [2].

Table 5.4. Commonly used direct solvers provided by AMESOS/AMESOS2

In certain cases, the user may want to do no smoothing on a particular level, or do no solve on the coarsest level.

- To skip smoothing, use the option `smoother: pre or post` with value `none`.

- To skip the coarse grid solve, use the option `coarse: type` with value `none`.

When a problem can be solved using structured aggregation algorithms it is also possible to use the structured line detection factory, this will allow MUELU to pass additional information to IFPACK 2 enabling it to perform line smoothing. An example of line smoothing is provided in `packages/trilinoscouplings/examples/scaling/muelu_Ifpack2_line_detection.xml`.

`smoother: pre or post`	[string] Pre- and post-smoother combination. Possible values: "pre" (only pre-smoother), "post" (only post-smoother), "both" (both pre-and post-smoothers), "none" (no smoothing). **Default:** "both".
`smoother: type`	[string] Smoother type. Possible values: see Table 5.3. **Default:** "RELAXATION".
`smoother: pre type`	[string] Pre-smoother type. Possible values: see Table 5.3. **Default:** "RELAXATION".
`smoother: post type`	[string] Post-smoother type. Possible values: see Table 5.3. **Default:** "RELAXATION".
`smoother: params`	[ParameterList] Smoother parameters. For standard smoothers, MUELU passes them directly to the appropriate package library.
`smoother: pre params`	[ParameterList] Pre-smoother parameters. For standard smoothers, MUELU passes them directly to the appropriate package library.
`smoother: post params`	[ParameterList] Post-smoother parameters. For standard smoothers, MUELU passes them directly to the appropriate package library.
`smoother: overlap`	[int] Smoother subdomain overlap. **Default:** 0.
`smoother: pre overlap`	[int] Pre-smoother subdomain overlap. **Default:** 0.

`smoother: post overlap`	[int] Post-smoother subdomain overlap. **Default:** 0.
`coarse: max size`	[int] Maximum dimension of a coarse grid. MUELU will stop coarsening once it is achieved. **Default:** 2000.
`coarse: type`	[string] Coarse solver. Possible values: see Table 5.4. **Default:** "SuperLU".
`coarse: params`	[ParameterList] Coarse solver parameters. MUELU passes them directly to the appropriate package library.
`coarse: overlap`	[int] Coarse solver subdomain overlap. **Default:** 0.

5.5 Aggregation options

`structured`	Attempts to construct hexahedral aggregates on a structured mesh using a default coarsening rate of 3 in each spatial dimension.
`hybrid`	This option takes in a user parameter that varies on each rank and that specifies whether the local aggregation scheme should be `structured` or `unstructured`.
`uncoupled`	Attempts to construct aggregates of optimal size (3^d nodes in d dimensions). Each process works independently, and aggregates cannot span multiple processes.
`coupled`	Attempts to construct aggregates of optimal size (3^d nodes in d dimensions). Aggregates are allowed to cross processor boundaries. **Use carefully**. If unsure, use `uncoupled` instead.
`brick`	Attempts to construct rectangular aggregates

Table 5.5. Available coarsening schemes.

`aggregation: type`	[string] Aggregation scheme. Possible values: see Table 5.5. **Default:** "uncoupled".

`aggregation: mode`		[string] Controls whether aggregates are allowed to cross processor boundaries. Possible values: "coupled" aggregates can cross processor boundaries, "uncoupled" aggregates cannot cross processor boundaries. **Default:** "uncoupled".
`aggregation: ordering`		[string] Node ordering strategy. Possible values: "natural" (local index order), "graph" (filtered graph breadth-first order), "random" (random local index order). **Default:** "natural".
`aggregation: drop scheme`		[string] Connectivity dropping scheme for a graph used in aggregation. Possible values: "classical", "distance laplacian". **Default:** "classical".
`aggregation: drop tol`		[double] Connectivity dropping threshold for a graph used in aggregation. **Default:** 0.0.
`aggregation: min agg size`		[int] Minimum size of an aggregate. **Default:** 2.
`aggregation: max agg size`		[int] Maximum size of an aggregate (-1 means unlimited). **Default:** -1.
`aggregation: brick x size`		[int] Number of points for x axis in "brick" aggregation (limited to 3). **Default:** 2.
`aggregation: brick y size`		[int] Number of points for y axis in "brick" aggregation (limited to 3). **Default:** 2.
`aggregation: brick z size`		[int] Number of points for z axis in "brick" aggregation (limited to 3). **Default:** 2.
`aggregation: Dirichlet threshold`		[double] Threshold for determining whether entries are zero during Dirichlet row detection. **Default:** 0.0.

`aggregation: phase 1 algorithm`	[string] Choice of algorithm for aggregation phase 1. **Default:** Serial.
`aggregation: export visualization data`	[bool] Export data for visualization post-processing. **Default:** false.
`aggregation: output filename`	[string] Filename to write VTK visualization data to. **Default:** "".
`aggregation: output file: time step`	[int] Time step ID for non-linear problems. **Default:** 0.
`aggregation: output file: iter`	[int] Iteration for non-linear problems. **Default:** 0.
`aggregation: output file: agg style`	[string] Style of aggregate visualization. **Default:** Point Cloud.
`aggregation: output file: fine graph edges`	[bool] Whether to draw all fine node connections along with the aggregates. **Default:** false.
`aggregation: output file: coarse graph edges`	[bool] Whether to draw all coarse node connections along with the aggregates. **Default:** false.
`aggregation: output file: build colormap`	[bool] Whether to output a random colormap in a separate XML file. **Default:** false.
`aggregation: mesh layout`	[string] Type of ordering for structured mesh aggregation. Possible values: "Global Lexicographic" and "Local Lexicographic". **Default:** Global Lexicographic.
`aggregation: output type`	[string] Type of object holding the aggregation data. Possible values: "Aggregates" or "CrsGraph". **Default:** Aggregates.
`aggregation: coarsening rate`	[string] Coarsening rate per spatial dimensions, the string must be interpretable as an array by Teuchos. **Default:** 3.

`aggregation: number of spatial dimensions`	[int] The number of spatial dimensions in the problem. **Default:** 3.
`aggregation: coarsening order`	[int] The interpolation order used while constructing these aggregates, this value will be passed to the prolongator factory. **Default:** 0.

5.6 Rebalancing options

`repartition: enable`	[bool] Rebalancing on/off switch. **Default:** false.
`repartition: partitioner`	[string] Partitioning package to use. Possible values: "zoltan" (ZOLTAN library), "zoltan2" (ZOLTAN2 library). **Default:** "zoltan2".
`repartition: params`	[ParameterList] Partitioner parameters. MUELU passes them directly to the appropriate package library. In particular, this allows to choose a partitioning algorithm from ZOLTAN or ZOLTAN 2 or from external packages such as PARMETIS.
`repartition: start level`	[int] Minimum level to run partitioner. MUELU does not rebalance levels finer than this one. **Default:** 2.
`repartition: min rows per proc`	[int] Minimum number of rows per MPI process. If the actual number if smaller, then rebalancing occurs. The value is not used if "repartition: min rows per thread" is positive. **Default:** 800.
`repartition: target rows per proc`	[int] Target number of rows per MPI process after rebalancing. If the value is set to 0, it will use the value of "repartition: min rows per proc" **Default:** 0.

`repartition: min rows per thread`		[int] Minimum number of rows per thread. If the actual number if smaller, then rebalancing occurs. If the value is set to 0, no repartitioning based on thread count will occur. **Default:** 0.
`repartition: target rows per thread`		[int] Target number of rows per thread after rebalancing. If the value is set to 0, it will use the value of "repartition: min rows per thread". **Default:** 0.
`repartition: max imbalance`		[double] Maximum nonzero imbalance ratio. If the actual number is larger, the rebalancing occurs. **Default:** 1.2.
`repartition: remap parts`		[bool] Postprocessing for partitioning to reduce data migration. **Default:** true.
`repartition: rebalance P and R`		[bool] Explicit rebalancing of R and P during the setup. This speeds up the solve, but slows down the setup phases. **Default:** false.

5.7 Multigrid algorithm options

sa	Classic smoothed aggregation [17]
unsmoothed	Aggregation-based, same as sa but without damped Jacobi prolongator improvement step
pg	Prolongator smoothing using A, restriction smoothing using A^T, local damping factors [13]
emin	Constrained minimization of energy in basis functions of grid transfer operator [20, 11]
interp	Interpolation based grid transfer operator, using piece-wise constant or linear interpolation from coarse nodes to fine nodes. This requires the use of structured aggregation.
semicoarsen	Semicoarsening grid transfer operator used to reduce a n-dimensional problem into a (n-1)-dimensional problem by coarsening fully in one of the spacial dimensions [14].
pcoarsen	

Table 5.6. Available multigrid algorithms for generating grid transfer matrices.

multigrid algorithm	[string] Multigrid method. Possible values: see Table 5.6. **Default:** "sa".
semicoarsen: coarsen rate	[int] Rate at which to coarsen unknowns in the z direction. **Default:** 3.
sa: damping factor	[double] Damping factor for smoothed aggregation. **Default:** 1.33.
sa: use filtered matrix	[bool] Matrix to use for smoothing the tentative prolongator. The two options are: to use the original matrix, and to use the filtered matrix with filtering based on filtered graph used for aggregation. **Default:** true.
interp: interpolation order	[int] Interpolation order used to interpolate values from coarse points to fine points. Possible values are 0 for piece-wise constant interpolation and 1 for piece-wise linear interpolation. This parameter is set to 1 by default. **Default:** 1.

`interp: build coarse coordinates`	[bool] If false, skip the calculation of coarse coordinates. **Default:** true.
`filtered matrix: use lumping`	[bool] Lump (add to diagonal) dropped entries during the construction of a filtered matrix. This allows user to preserve constant nullspace. **Default:** true.
`filtered matrix: reuse eigenvalue`	[bool] Skip eigenvalue calculation during the construction of a filtered matrix by reusing eigenvalue estimate from the original matrix. This allows us to skip heavy computation, but may lead to poorer convergence. **Default:** true.
`emin: iterative method`	[string] Iterative method to use for energy minimization of initial prolongator in energy-minimization. Possible values: "cg" (conjugate gradient), "gmres" (generalized minimum residual), "sd" (steepest descent). **Default:** "cg".
`emin: num iterations`	[int] Number of iterations to minimize initial prolongator energy in energy-minimization. **Default:** 2.
`emin: num reuse iterations`	[int] Number of iterations to minimize the reused prolongator energy in energy-minimization. **Default:** 1.
`emin: pattern`	[string] Sparsity pattern to use for energy minimization. Possible values: "AkPtent". **Default:** "AkPtent".
`emin: pattern order`	[int] Matrix order for the "AkPtent" pattern. **Default:** 1.

5.8 Reuse options

Reuse options are a way for a user to speed up the setup stage of multigrid. The main requirement to use reuse is that the matrix' graph structure does not change. Only matrix values are allowed to change.

The reuse options control the degree to which the multigrid hierarchy is preserved for a subsequent setup call.

In addition, please note that not all combinations of multigrid algorithms and reuse options are valid, or even make sense. For instance, the "emin" reuse option should only be used with the "emin" multigrid algorithm.

Table 5.7 contains the information about different reuse options. The options are ordered in increasing number of reuse components, from the no reuse to the full reuse ("full").

`none`	No reuse
`S`	Reuse only the symbolic information of the level smoothers.
`tP`	Reuse tentative prolongator. The graphs of smoothed prolongator and matrices in Galerkin product are reused only if filtering is not being used (*i.e.*, either `sa: use filtered matrix` or `aggregation: drop tol` is false)
`emin`	Reuse old prolongator as an initial guess to energy minimization, and reuse the prolongator pattern
`RP`	Reuse smoothed prolongator and restrictor. Smoothers are recomputed. ☛ RP should reuse matrix graphs for matrix-matrix product, but currently that is disabled as only EPETRA supports it.
`RAP`	Recompute only the finest level smoothers, reuse all other operators
`full`	Reuse everything

Table 5.7. Available reuse options.

`reuse: type`	[string] Reuse options for consecutive hierarchy construction. This speeds up the setup phase, but may lead to poorer convergence. Possible values: see Table 5.7. **Default:** "none".

5.9   Miscellaneous options

`export data`   [ParameterList] Exporting a subset of the hierarchy data in a file. Currently, the list can contain any of the following parameter names ("A", "P", "R", "Nullspace", "Coordinates") of type string and value "{levels separated by commas}". A matrix/multivector with a name "X" is saved in two or three three Matrix-Market files: a) data is saved in *X_level.mm*; b) its row map is saved in *rowmap_X_level.mm*; c) its column map (for matrices) is saved in *colmap_X_level.mm*.

`print initial parameters`   [bool] Print parameters provided for a hierarchy construction. **Default:** true.

`print unused parameters`   [bool] Print parameters unused during a hierarchy construction. **Default:** true.

`transpose: use implicit`   [bool] Use implicit transpose for the restriction operator. **Default:** false.

`use kokkos refactor`   [bool] Switch on the new KOKKOS based version for on-node parallelism. **Default:** false.

`rap: triple product`   [bool] Use all-at-once triple matrix product kernel **Default:** false.

5.10   Maxwell solver options

`refmaxwell: mode`   [string] Specifying the order of solve of the block system. Allowed values are: "additive" (default), "121", "212", "1", "2" **Default:** "additive".

`refmaxwell: disable addon`   [bool] Specifing whether the addon should be built for stabilization **Default:** true.

`refmaxwell: 11list`   [ParameterList] Specifies the multigrid solver for the 11 block

`refmaxwell: 22list`	[ParameterList] Specifies the multigrid solver for the 22 block
`refmaxwell: use as preconditioner`	[bool] Assume zero initial guess **Default:** false.
`refmaxwell: dump matrices`	[bool] Dump matrices to disk. **Default:** false.
`refmaxwell: subsolves on subcommunicators`	[bool] Redistribute the two subsolves to disjoint sub-communicators (so that the additive solve can occur in parallel). **Default:** false.
`refmaxwell: ratio AH / A22 subcommunicators`	[double] Ratio for the split into sub-communicators. **Default:** 1.0.

Chapter 6

MUEMEX: The MATLAB Interface for MUELU

MUEMEX is MUELU's interface to the MATLAB environment. It allows access to a limited set of routines either MUELU as a preconditioner, Belos as a solver and Epetra or Tpetra for data structures. It is designed to provide access to MUELU's aggregation and solver routines from MATLAB and does little else. MUEMEX allows users to setup and solve arbitrarily many problems, so long as memory suffices. More than one problem can be set up simultaneously.

6.1 CMake Configure and Make

To use MUEMEX, Trilinos must be configured with (at least) the following options:

```
export TRILINOS_HOME=/path/to/your/Trilinos/source/directory
cmake \
    -D Trilinos_ENABLE_EXPLICIT_INSTANTIATION:BOOL=ON \
    -D Trilinos_ENABLE_Amesos:BOOL=ON \
    -D Trilinos_ENABLE_Amesos2:BOOL=ON \
    -D Amesos2_ENABLE_KLU2:BOOL=ON \
    -D Trilinos_ENABLE_AztecOO:BOOL=ON \
    -D Trilinos_ENABLE_Epetra:BOOL=ON \
    -D Trilinos_ENABLE_EpetraExt:BOOL=ON \
    -D Trilinos_ENABLE_Fortran:BOOL=OFF \
    -D Trilinos_ENABLE_Ifpack:BOOL=ON \
    -D Trilinos_ENABLE_Ifpack2:BOOL=ON \
    -D Trilinos_ENABLE_MueLu:BOOL=ON \
    -D Trilinos_ENABLE_Teuchos:BOOL=ON \
    -D Trilinos_ENABLE_Tpetra:BOOL=ON \
    -D TPL_ENABLE_MPI:BOOL=OFF \
    -D TPL_ENABLE_MATLAB:BOOL=ON \
    -D MATLAB_ROOT:STRING="<my matlab root>" \
    -D MATLAB_ARCH:STRING="<my matlab os string>" \
    -D Trilinos_EXTRA_LINK_FLAGS="-lrt -lm -lgfortran" \
${TRILINOS_HOME}
```

Since MUEMEX supports both the Epetra and Tpetra linear algebra libraries, you have to have both enabled in order to build MUEMEX.

☞ If you turn off either Epetra or Tpetra then you will run into an error message: *MueMex requires Epetra, Tpetra and MATLAB.*

Most additional options can be specified as well. It is important to note that MUEMEX does not work properly with MPI, hence MPI must be disabled in order to compile MUEMEX. The MATLAB_ARCH option is new to the cmake build system, and involves the MATLAB-specific architecture code for your system. There is currently no automatic way to extract this, so it must be user-specified. As of MATLAB 7.9 (R2009b), common arch codes are:

Code	OS
glnx86	32-bit Linux (intel/amd)
glnxa64	64-bit Linux (intel/amd)
maci64	64-bit MacOS
maci	32-bit MacOS

On 64-bit Intel/AMD architectures, Trilinos and all relevant TPLs (note: this includes BLAS and LAPACK) must be compiled with the `-fPIC` option. This necessitates adding:

```
-D CMAKE_CXX_FLAGS:STRING="-fPIC" \
-D CMAKE_C_FLAGS:STRING="-fPIC" \
-D CMAKE_Fortran_FLAGS:STRING="-fPIC" \
```

to the cmake configure line.

The additional linker flags specified in `Trilinos_EXTRA_LINK_FLAGS` may slightly vary depending on the system and the exact configuration. But the given parameters may work for most Linux based systems. If you encounter an error message like *Target "muemex.mexa64" links to item "-Wl,-rpath-link,/opt/matlab/bin/glnxa64 " which has leading or trailing whitespace.* you have to add some options to the `Trilinos_EXTRA_LINK_FLAGS` variable. At least adding `-lm` should be safe and fix the error message.

6.1.1 BLAS & LAPACK Option #1: Static Builds

Trilinos does not play nicely with MATLAB's default LAPACK and BLAS on 64-bit machines. If MUEMEX randomly crashes when you run with any Krylov method that has orthogonalization, chances are MUEMEX is finding the wrong BLAS/LAPACK libraries. This leaves you with one of two options. The first is to build them both *statically* and then specify them as follows:

```
-D LAPACK_LIBRARY_DIRS:STRING="<path to my lapack.a>" \
-D BLAS_LIBRARY_DIRS:STRING="<path to my blas.a>" \
```

Using static linking for LAPACK and BLAS prevents MATLAB's default libraries to take precedence.

6.1.2   BLAS & LAPACK Option #2: LD_PRELOAD

The second option is to use `LD_PRELOAD` to tell MATLAB exactly which libraries to use. For this option, you can use the dynamic libraries installed on your system. Before starting MATLAB, set `LD_PRELOAD` to the paths of libstdc++.so corresponding to the version of GCC used to build Trilinos, and the paths of libblas.so and liblapack.so on your local system.

For example, if you use bash, you'd do something like this

```
export LD_PRELOAD=<path>/libstdc++.so:<path>/libblas.so:<path>/liblapack.so
```

For csh / tcsh, do this

```
setenv LD_PRELOAD <path>/libstdc++.so:<path>/libblas.so:<path>/liblapack.so
```

6.1.3   Running MATLAB

Before you run MATLAB you have to make sure that MATLAB is using the same libraries that have been used for compiling MUEMEX. This includes the `libstdc++.so` and depending whether you turned on/off fortran also `libgfortran.so`. Please make sure that the correct libraries and paths are declared in the `LD_PRELOAD` variable. You can refer to section 6.1.2 to see how the `LD_PRELOAD` variable is set.

For a 64 bit Linux system using the bash the command should look like

```
export LD_PRELOAD=/usr/lib64/libstdc++.so.6:/usr/lib64/libgfortran.so.3:
   $LD_PRELOAD
```

to add the `libstdc++.so` and `libgfortran.so` to the existing `LD_PRELOAD` variable. Then run the MATLAB executable in the same shell window.

☞ Note, that this step is necessary even if you statically linked BLAS and LAPACK.

If you are unsure which libraries have to be set in the `LD_PRELOAD` variable you will find out latest if you start MATLAB and try to run MUEMEX. It will throw some error messages with the missing library names. For a 64 bit Linux system the standard libraries usually can be found in `/usr/lib64` or `/usr/lib` (for a 32 bit system).

6.2   Using MUEMEX

MUEMEX is designed to be interfaced with via the MATLAB script `muelu.m`. There are five modes in which MUEMEX can be run:

1. Setup Mode — Performs the problem setup for MUELU. Depending on whether or not the `Linear Algebra` option is used, MUEMEX creates either an unpreconditioned Epetra problem, an Epetra problem with MUELU, or a Tpetra problem with MUELU. The default is `tpetra`. The `epetra` mode only supports real-valued matrices, while `tpetra` supports both real and complex and will infer the scalar type from the matrix passed during setup. This call returns a problem handle used to reference the problem in the future, and (optionally) the operator complexity, if a preconditioner is being used.

2. Solve Mode — Given a problem handle and a right-hand side, MUEMEX solves the problem specified. Setup mode must be called before solve mode.

3. Cleanup Mode — Frees the memory allocated to internal MUELU, Epetra and Tpetra objects. This can be called with a particular problem handle, in which case it frees that problem, or without one, in which case all MUEMEX memory is freed.

4. Status Mode — Prints out status information on problems which have been set up. Like cleanup, it can be called with or without a particular problem handle.

5. Get Mode — Get information from a MueLu hierarchy that has been generated. Given the problem handle, a level number and the name of the field, returns the appropriate array or scalar as a MATLAB object.

All of these modes, with the exception of status and cleanup take option lists which will be directly converted into `Teuchos::ParameterList` objects by MUEMEX, as key-value pairs. Options passed during setup will apply to the MUELU preconditioner, and options passed during a solve will apply to Belos.

6.2.1   Setup Mode

Setup mode is called as follows:

```
>> [h, oc] = muelu('setup', A[, 'parameter', value,...])
```

The parameter `A` represents the sparse matrix to perform aggregation on and the parameter/value pairs represent standard MUELU options.

The routine returns a problem handle, `h`, and the operator complexity `oc` for the operator. In addition to the standard options, setup mode has one unique option of its own:

Linear Algebra	[string] Whether to use 'epetra unprec', 'epetra', or 'tpetra'. Default is 'epetra' for real matrix and 'tpetra' for complex matrix.

6.2.2   Solve Mode

Solve mode is called as follows:

```
>> [x, its] = muelu(h[, A], b[, 'parameter', value,...])
```

The parameter `h` is a problem handle returned by the setup mode call, `A` is the sparse matrix with which to solve and `b` is the right-hand side. Parameter/value pairs to configure the Belos solver are listed as above. If A is not supplied, the matrix provided when setting up the problem will be used. `x` is the solution multivector with the same dimensions as `b`, and `its` is the number of iterations Belos needed to solve the problem.

All of these options are taken directly from Belos, so consult its manual for more information. Belos output style and verbosity settings are implemented as enums, but can be set as strings in MUEMEX. For example:

```
>> x = muelu(0, b, 'Verbosity', 'Warnings + IterationDetails', ...
                   'Output Style', 'Brief');
```

Verbosity settings can be separated by spaces, '+' or ','. Belos::Brief is the default output style.

6.2.3   Cleanup Mode

Cleanup mode is called as follows:

```
>> muelu('cleanup'[, h])
```

The parameter `h` is a problem handle returned by the setup mode call and is optional. If h is provided, that problem is cleaned up. If the option is not provided all currently set up problems are cleaned up.

6.2.4   Status Mode

Status mode is called as follows:

```
>> muelu('status'[, h])
```

The parameter h is a problem handle returned by the setup mode call and is optional. If h is provided, status information for that problem is printed. If the option is not provided all currently set up problems have status information printed.

6.2.5   Get Mode

Get mode is called as follows:

```
>> muelu('get', h, level, fieldName[, typeHint])
```

The parameter h is the problem handle, and level is an integer that identifies the level within the hierarchy containing the desired data. fieldName is a string that identifies the field within the level, e.g. 'Nullspace'. typeHint is an optional parameter that tells MueMex what data type to expect from the level. This is a string, with possible values 'matrix', 'multivector', 'lovector' (ordinal vector), or 'scalar'. MueMex will attempt to guess the type from fieldName but typeHint may be required.

6.2.6   Tips and Tricks

Internally, MATLAB represents all data as doubles unless you go through efforts to do otherwise. MUEMEX detects integer parameters by a relative error test, seeing if the relative difference between the value from MATLAB and the value of the int-typecast value are less than 1e-15. Unfortunately, this means that MUEMEX will choose the incorrect type for parameters which are doubles that happen to have an integer value (a good example of where this might happen would be the parameter 'smoother Chebyshev: alpha', which defaults to 30.0). Since MUEMEX does no internal typechecking of parameters (it uses MUELU's internal checks), it has no way of detecting this conflict. From the user's perspective, avoiding this is as simple as adding a small perturbation (greater than a relative 1e-15) to the parameter that makes it non-integer valued.

References

[1] Trilinos CMake Quickstart. http://trilinos.org/build_instructions.html, 2014.

[2] Patrick R. Amestoy, Iain S. Duff, Jean-Yves L'Excellent, and Jacko Koster. MUMPS: a general purpose distributed memory sparse solver. In *Applied Parallel Computing. New Paradigms for HPC in Industry and Academia*, pages 121–130. Springer, 2001.

[3] Roscoe A. Bartlett. Teuchos::RCP beginner's guide. Technical Report SAND2004-3268, Sandia National Labs, 2010.

[4] Eric Bavier, Mark Hoemmen, Sivasankaran Rajamanickam, and Heidi Thornquist. Amesos2 and Belos: Direct and iterative solvers for large sparse linear systems. *Scientific Programming*, 20(3):241–255, 2012.

[5] P. B. Bochev, J. J. Hu, C. M. Siefert, and R. S. Tuminaro. An algebraic multigrid approach based on a compatible gauge reformulation of Maxwell's equations. *SIAM Journal on Scientific Computing*, 31(1):557–583, 2008.

[6] William L. Briggs, Van Emden Henson, and Steve F. McCormick. *A multigrid tutorial*. SIAM, 2nd edition, 2000.

[7] Timothy A. Davis. Algorithm 832: UMFPACK v4.3 — an unsymmetric-pattern multifrontal method. *ACM Transactions on Mathematical Software (TOMS)*, 30(2):196–199, 2004.

[8] Timothy A. Davis and Ekanathan Palamadai Natarajan. Algorithm 907: KLU, a direct sparse solver for circuit simulation problems. *ACM Transactions on Mathematical Software (TOMS)*, 37(3):36, 2010.

[9] Xiaoye S. Li, James W. Demmel, John R. Gilbert, Laura Grigori, Meiyue Shao, and Ichitaro Yamazaki. SuperLU Users' Guide. 2011.

[10] Paul Lin, Matthew Bettencourt, Stefan Domino, Travis Fisher, Mark Hoemmen, Jonathan J. Hu, Eric Phipps, Andrey Prokopenko, Sivasankaran Rajamanickam, Christopher Siefert, and Stephen Kennon. Towards extreme-scale simulations for low Mach fluids with second-generation Trilinos. *Parallel Processing Letters*, 24(04):1442005, 2014.

[11] L. Olson, J. Schroder, and R. Tuminaro. A general interpolation strategy for algebraic multigrid using energy minimization. *SIAM Journal on Scientific Computing*, 33(2):966–991, 2011.

[12] Andrey Prokopenko, Christopher M. Siefert, Jonathan J. Hu, Mark Hoemmen, and Alicia Klinvex. Ifpack2 User's Guide 1.0. Technical Report SAND2016-5338, Sandia National Labs, 2016.

[13] Marzio Sala and Raymond S. Tuminaro. A new Petrov–Galerkin smoothed aggregation preconditioner for nonsymmetric linear systems. *SIAM Journal on Scientific Computing*, 31(1):143–166, 2008.

[14] I. K. Tezaur, M. Perego, A. G. Salinger, R. S. Tuminaro, and S. F. Price. Albany/FELIX: a parallel, scalable and robust, finite element, first-order Stokes approximation ice sheet solver built for advanced analysis. *Geoscientific Model Development*, 8:1197–1220, 2015.

[15] Heidi Thornquist, Roscoe A. Bartlett, Mark Hoemmen, Christopher Baker, and Michael Heroux. Teuchos: Trilinos tools library. http://trilinos.org/packages/teuchos, 2014.

[16] Ulrich Trottenberg, Cornelis Oosterlee, and Anton Schüller. *Multigrid*. Elsevier Academic Press, 2001.

[17] P. Vaněk, J. Mandel, and M. Brezina. Algebraic multigrid based on smoothed aggregation for second and fourth order problems. *Computing*, 56:179–196, 1996.

[18] Tobias A. Wiesner. *Flexible Aggregation-based Algebraic Multigrid Methods for Contact and Flow Problems*. PhD thesis, 2014.

[19] Tobias A. Wiesner, Michael W. Gee, Andrey Prokopenko, and Jonathan J. Hu. The MueLu tutorial. http://trilinos.org/packages/muelu/muelu-tutorial, 2014. SAND2014-18624R.

[20] Tobias A. Wiesner, Raymond S. Tuminaro, Wolfgang A. Wall, and Michael W. Gee. Multigrid transfers for nonsymmetric systems based on Schur complements and Galerkin projections. *Numerical Linear Algebra with Applications*, 21(3):415–438, 2014.

Appendix A

Copyright and License

MueLu: A package for multigrid based preconditioning

Copyright 2012 Sandia Corporation

Under the terms of Contract DE–AC04–94AL85000 with Sandia Corporation, the U.S. Government retains certain rights in this software.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

3. Neither the name of the Corporation nor the names of the contributors may be used to endorse or promote products derived from this software without specific prior written permission.

Appendix B

ML compatibility

MUELU provides a basic compatibility layer for ML parameter lists. This allows ML users to quickly perform some experiments with MUELU.

First and most important: Long term, we would like to have users use the new MUELU interface, as that is where most of new features will be made accessible. One should make note of the fact that it may not be possible to make ML deck do exactly same things in ML and MUELU, as internally ML implicitly makes some decision that we have no control over and which could be different from MUELU.

There are basically two distinct ways to use ML input parameters with MUELU:

MLParameterListInterpreter: This class is the pendant of the `ParameterListInterpreter` class for the MUELU parameters. It accepts parameter lists or XML files with ML parameters and generates a MUELU multigrid hierarchy. It supports only a well-defined subset of ML parameters which have a equivalent parameter in MUELU.

ML2MueLuParameterTranslator: This class is a simple wrapper class which translates ML parameters to the corresponding MUELU parameters. It has to be used in combination with the MUELU `ParameterListInterpreter` class to generate a MUELU multigrid hierarchy. It is also meant to be used in combination with the `CreateEpetraPreconditioner` and `CreateTpetraPreconditioner` routines (see §3.3). It supports only a small subset of the ML parameters.

B.1 Usage of ML parameter lists with MUELU

B.1.1 MLParameterListInterpreter

The `MLParameterListInterpreter` directly accepts a `ParameterList` containing ML parameters. It also interprets the `null space: vectors` and the `null space: dimension` ML parameters. However, it is recommended to provide the near null space vectors directly in the MUELU way as shown in the following code snippet.

```
Teuchos::RCP<Tpetra::CrsMatrix<> > A;
// create A here ...

// XML file containing ML parameters
std::string xmlFile = "mlParameters.xml"
Teuchos::ParameterList paramList;
Teuchos::updateParametersFromXmlFileAndBroadcast(xmlFile, Teuchos::Ptr<Teuchos
    ::ParameterList>(&paramList), *comm);

// use ParameterListInterpreter with MueLu parameters as input
Teuchos::RCP<HierarchyManager> mueluFactory = Teuchos::rcp(new
    MLParameterListInterpreter(*paramList));

RCP<Hierarchy> H = mueluFactory->CreateHierarchy();
H->GetLevel(0)->Set<RCP<Matrix> >("A", A);
H->GetLevel(0)->Set("Nullspace", nullspace);
H->GetLevel(0)->Set("Coordinates", coordinates);
mueluFactory->SetupHierarchy(*H);
```

Note that the `MLParameterListInterpreter` only supports a basic set of ML parameters allowing to build smoothed aggregation transfer operators (see §B.2 for a list of compatible ML parameters).

B.1.2 ML2MueLuParameterTranslator

The `Ml2MueLuParameterTranslator` class is a simple wrapper translating ML parameters to the corresponding MUELU parameters. This allows the usage of the simple `CreateEpetraPreconditioner` and `CreateTpetraPreconditioner` interface with ML parameters:

```
Teuchos::RCP<Tpetra::CrsMatrix<> > A;
// create A here ...

// XML file containing ML parameters
std::string xmlFile = "mlParameters.xml"
Teuchos::ParameterList paramList;
Teuchos::updateParametersFromXmlFileAndBroadcast(xmlFile, Teuchos::Ptr<Teuchos
    ::ParameterList>(&paramList), *comm);

// translate ML parameters to MueLu parameters
RCP<ParameterList> mueluParamList = Teuchos::getParametersFromXmlString(MueLu
    ::ML2MueLuParameterTranslator::translate(paramList,"SA"));

Teuchos::RCP<MueLu::TpetraOperator> mueLuPreconditioner =
    MueLu::CreateTpetraPreconditioner(A, mueluParamList);
```

In a similar way, ML input parameters can be used with the standard MUELU parameter list interpreter class. Note that the near null space vectors have to be provided in the MUELU way.

```
Teuchos::RCP<Tpetra::CrsMatrix<> > A;
// create A here ...

// XML file containing ML parameters
std::string xmlFile = "mlParameters.xml"
Teuchos::ParameterList paramList;
Teuchos::updateParametersFromXmlFileAndBroadcast(xmlFile, Teuchos::Ptr<Teuchos
    ::ParameterList>(&paramList), *comm);

// translate ML parameters to MueLu parameters
RCP<ParameterList> mueluParamList = Teuchos::getParametersFromXmlString(MueLu
    ::ML2MueLuParameterTranslator::translate(paramList,"SA"));

// use ParameterListInterpreter with MueLu parameters as input
Teuchos::RCP<HierarchyManager> mueluFactory = Teuchos::rcp(new
    ParameterListInterpreter(*mueluParamList));

RCP<Hierarchy> H = mueluFactory->CreateHierarchy();
H->GetLevel(0)->Set<RCP<Matrix> >("A", A);
H->GetLevel(0)->Set("Nullspace", nullspace);
H->GetLevel(0)->Set("Coordinates", coordinates);
mueluFactory->SetupHierarchy(*H);
```

Note that the set of supported ML parameters is very limited. Please refer to §B.2 for a list of all compatible ML parameters.

B.2   Compatible ML parameters

B.2.1   General ML options

ML output   [int] Control of the amount of printed information.
                                Possible values: 0-10 with 0=no output and 10=maxi-
                                mum verbosity. **Default:** 0 **Compatibility:** MLParam-
                                eterListInterpreter, ML2MueLuParameterTranslator.

PDE equations   [int] Number of PDE equations at each grid node.
                                Only constant block size is considered. **De-
                                fault:** 1 **Compatibility:** MLParameterListInterpreter,
                                ML2MueLuParameterTranslator.

`max levels`	[int] Maximum number of levels in a hierarchy. **Default:** 10 **Compatibility:** MLParameterListInterpreter, ML2MueLuParameterTranslator.
`prec type`	[string] Multigrid cycle type. Possible values: "MGV", "MGW". Other values are NOT supported by MueLu. **Default:** "MGV" **Compatibility:** MLParameterListInterpreter, ML2MueLuParameterTranslator.

B.2.2   Smoothing and coarse solver options

`smoother: type`	[string] Smoother type for fine- and intermedium multigrid levels. Possible values: "Jacobi", "Gauss-Seidel", "symmetric Gauss-Seidel", "Chebyshev", "ILU". **Default:** "symmetric Gauss-Seidel" **Compatibility:** MLParameterListInterpreter, ML2MueLuParameterTranslator.
`smoother: sweeps`	[int] Number of smoother sweeps for relaxation based level smoothers. In case of Chebyshev smoother it denotes the polynomial degree. **Default:** 2 **Compatibility:** MLParameterListInterpreter, ML2MueLuParameterTranslator.
`smoother: damping factor`	[double] Damping factor for relaxation based level smoothers. **Default:** 1.0 **Compatibility:** MLParameterListInterpreter, ML2MueLuParameterTranslator.
`smoother: Chebyshev alpha`	[double] Eigenvalue ratio for Chebyshev level smoother. **Default:** 20 **Compatibility:** MLParameterListInterpreter, ML2MueLuParameterTranslator.
`smoother: pre or post`	[string] Pre- and post-smoother combination. Possible values: "pre" (only pre-smoother), "post" (only post-smoother), "both" (both pre-and post-smoothers). **Default:** "both" **Compatibility:** MLParameterListInterpreter, ML2MueLuParameterTranslator.

`max size`	[int] Maximum dimension of a coarse grid. ML will stop coarsening once it is achieved. **Default:** 128 **Compatibility:** MLParameterListInterpreter, ML2MueLuParameterTranslator.
`coarse: type`	[string] Solver for coarsest level. Possible values: "Amesos-KLU", "Amesos-Superlu" (depending on MUELU installation). **Default:** "Amesos-KLU" **Compatibility:** MLParameterListInterpreter, ML2MueLuParameterTranslator.

B.2.3 Transfer operator options

`energy minimization: enable`	[int] Enable energy minimization transfer operators (using Petrov-Galerkin approach). **Default:** 0 **Compatibility:** MLParameterListInterpreter, ML2MueLuParameterTranslator.
`aggregation: damping factor`	[double] Damping factor for smoothed aggregation. **Default:** 1.33 **Compatibility:** MLParameterListInterpreter, ML2MueLuParameterTranslator.

B.2.4 Rebalancing options

`repartition: enable`	[int] Rebalancing on/off switch. Only limited support for repartitioning. Does not use provided node coordinates. **Default:** 0 **Compatibility:** MLParameterListInterpreter.
`repartition: start level`	[int] Minimum level to run partitioner. MUELU does not rebalance levels finer than this one. **Default:** 1 **Compatibility:** MLParameterListInterpreter.
`repartition: min per proc`	[int] Minimum number of rows per MPI process. If the actual number if smaller, then rebalancing occurs. **Default:** 512 **Compatibility:** MLParameterListInterpreter.

`repartition: max min ratio`	[double] Maximum nonzero imbalance ratio. If the actual number is larger, the rebalancing occurs. **Default:** 1.3 **Compatibility:** MLParameterListInterpreter.

DISTRIBUTION

Email–External (encrypt for OUO)

Name	Company Email Address	Company Name
Matthias Mayr	matthias.mayr@unibw.de	University of the Bundeswehr Munich
Andrey Prokopenko	prokopenkoav@ornl.gov	Oak Ridge National Laboratory
Tobias Wiesner	tobias.wiesner@leica-geosystems.com	Leica Geosystems AG

Email–Internal

Name	Org.	Sandia Email Address
Luc Berger-Vergiat	1442	lberge@sandia.gov
Christopher Siefert	1465	csiefer@sandia.gov
Christian Glusa	1465	caglusa@sandia.gov
Mark Hoemmen	1541	mhoemme@sandia.gov
Jonathan Hu	1465	jhu@sandia.gov
Paul Lin	1422	ptlin@sandia.gov
Ray Tuminaro	1442	rstumin@sandia.gov
Technical Library	9536	libref@sandia.gov