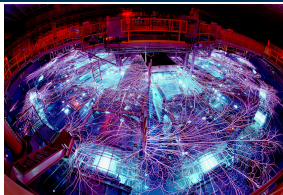


Exceptional service in the national interest



Sandia
National
Laboratories



Lessons Learned from Porting the MiniAero Application to Charm++

David S. Hollman, Janine Bennett (PI), Jeremiah Wilke (Chief Architect),
Ken Franko, Hemanth Kolla, Paul Lin, Greg Sjaardema, Nicole Slattengren,
Keita Teranishi, Nikhil Jain, Eric Mikida

May 7, 2015



Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000. SAND NO. 2015-3671 C

Introduction

Introduction

The Process: Porting an Explicit Aerodynamics Miniapp to the Chare Model

What was easy?

What was harder?

Introduction

The Process: Porting an Explicit Aerodynamics Miniapp to the Chare Model

What was easy?

What was harder?

Preliminary Results and Performance

Introduction

The Process: Porting an Explicit Aerodynamics Miniapp to the Chare Model

What was easy?

What was harder?

Preliminary Results and Performance

Next Steps

Introduction

The Process: Porting an Explicit Aerodynamics Miniapp to the Chare Model

What was easy?

What was harder?

Preliminary Results and Performance

Next Steps

- DHARMA: Distributed asyncHronous Adaptive Resilient Models for Applications



- DHARMA: Distributed asyncHronous Adaptive Resilient Models for Applications
- Project mission: Assess and address fundamental challenges imposed by the need for performant, portable, scalable, fault-tolerant programming models at extreme-scale





- DHARMA: Distributed asyncHronous Adaptive Resilient Models for Applications
- Project mission: Assess and address fundamental challenges imposed by the need for performant, portable, scalable, fault-tolerant programming models at extreme-scale
- Research in programmability, dynamic load-balancing, and fault-tolerance of AMT runtimes



- DHARMA: Distributed asyncHronous Adaptive Resilient Models for Applications
- Project mission: Assess and address fundamental challenges imposed by the need for performant, portable, scalable, fault-tolerant programming models at extreme-scale
- Research in programmability, dynamic load-balancing, and fault-tolerance of AMT runtimes
- Comparative analysis portion of the project:



- DHARMA: Distributed asyncHronous Adaptive Resilient Models for Applications
- Project mission: Assess and address fundamental challenges imposed by the need for performant, portable, scalable, fault-tolerant programming models at extreme-scale
- Research in programmability, dynamic load-balancing, and fault-tolerance of AMT runtimes
- Comparative analysis portion of the project:
 - Assess various asynchronous many-task (AMT) runtimes by implementing mini-apps of interest to Sandia using existing runtimes



- DHARMA: Distributed asyncHronous Addaptive Resilient Models for Applications
- Project mission: Assess and address fundamental challenges imposed by the need for performant, portable, scalable, fault-tolerant programming models at extreme-scale
- Research in programmability, dynamic load-balancing, and fault-tolerance of AMT runtimes
- Comparative analysis portion of the project:
 - Assess various asynchronous many-task (AMT) runtimes by implementing mini-apps of interest to Sandia using existing runtimes
 - First three runtimes to assess:
 - Charm++
 - Legion
 - Uintah



- DHARMA: Distributed asyncHronous Addaptive Resilient Models for Applications
- Project mission: Assess and address fundamental challenges imposed by the need for performant, portable, scalable, fault-tolerant programming models at extreme-scale
- Research in programmability, dynamic load-balancing, and fault-tolerance of AMT runtimes
- Comparative analysis portion of the project:
 - Assess various asynchronous many-task (AMT) runtimes by implementing mini-apps of interest to Sandia using existing runtimes
 - First three runtimes to assess:
 - Charm++
 - Legion
 - Uintah
 - First mini-app for assessment: MiniAero



- DHARMA: Distributed asyncHronous Addaptive Resilient Models for Applications
- Project mission: Assess and address fundamental challenges imposed by the need for performant, portable, scalable, fault-tolerant programming models at extreme-scale
- Research in programmability, dynamic load-balancing, and fault-tolerance of AMT runtimes
- Comparative analysis portion of the project:
 - Assess various asynchronous many-task (AMT) runtimes by implementing mini-apps of interest to Sandia using existing runtimes
 - First three runtimes to assess:
 - Charm++
 - Legion
 - Uintah
 - First mini-app for assessment: MiniAero

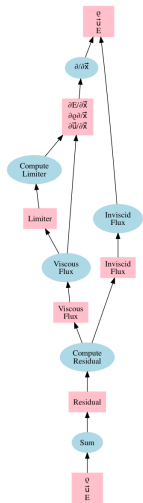


- MiniAero is a proxy app illustrating the common computation and communication patterns in unstructured mesh codes of interest to Sandia

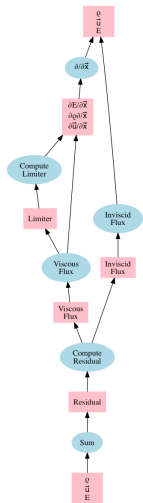
- MiniAero is a proxy app illustrating the common computation and communication patterns in unstructured mesh codes of interest to Sandia
- 3D, unstructured, finite volume computational fluid dynamics code
 - Uses Runge-Kutta fourth-order time marching
 - Has options for first and second order spatial discretization
 - Includes inviscid Roe and viscous Newtonian flux options

- MiniAero is a proxy app illustrating the common computation and communication patterns in unstructured mesh codes of interest to Sandia
- 3D, unstructured, finite volume computational fluid dynamics code
 - Uses Runge-Kutta fourth-order time marching
 - Has options for first and second order spatial discretization
 - Includes inviscid Roe and viscous Newtonian flux options
 - Baseline application is about 3800 lines of C++ code using MPI and Kokkos

- MiniAero is a proxy app illustrating the common computation and communication patterns in unstructured mesh codes of interest to Sandia
- 3D, unstructured, finite volume computational fluid dynamics code
 - Uses Runge-Kutta fourth-order time marching
 - Has options for first and second order spatial discretization
 - Includes inviscid Roe and viscous Newtonian flux options
 - Baseline application is about 3800 lines of C++ code using MPI and Kokkos
- Very little task parallelism; mostly a data parallel problem



- MiniAero is a proxy app illustrating the common computation and communication patterns in unstructured mesh codes of interest to Sandia
- 3D, unstructured, finite volume computational fluid dynamics code
 - Uses Runge-Kutta fourth-order time marching
 - Has options for first and second order spatial discretization
 - Includes inviscid Roe and viscous Newtonian flux options
 - Baseline application is about 3800 lines of C++ code using MPI and Kokkos
- Very little task parallelism; mostly a data parallel problem
- Communication: ghost exchanges, unstructured mesh



- Porting MiniAero to Charm++ began with a “bootcamp”:

- Porting MiniAero to Charm++ began with a “bootcamp”:
 - March 9-12, 2015

- Porting MiniAero to Charm++ began with a “bootcamp”:
 - March 9-12, 2015
 - Led by Nikhil Jain and Eric Mikida

- Porting MiniAero to Charm++ began with a “bootcamp”:
 - March 9-12, 2015
 - Led by Nikhil Jain and Eric Mikida
 - About 10 Sandia scientists in attendance

- Porting MiniAero to Charm++ began with a “bootcamp”:
 - March 9-12, 2015
 - Led by Nikhil Jain and Eric Mikida
 - About 10 Sandia scientists in attendance
- Since the workshop, we've had one scientist working 50% time on the port and a couple others working 10-20% time

- Porting MiniAero to Charm++ began with a “bootcamp”:
 - March 9-12, 2015
 - Led by Nikhil Jain and Eric Mikida
 - About 10 Sandia scientists in attendance
- Since the workshop, we've had one scientist working 50% time on the port and a couple others working 10-20% time
- Current state of the code:

- Porting MiniAero to Charm++ began with a “bootcamp”:
 - March 9-12, 2015
 - Led by Nikhil Jain and Eric Mikida
 - About 10 Sandia scientists in attendance
- Since the workshop, we've had one scientist working 50% time on the port and a couple others working 10-20% time
- Current state of the code:
 - running and passes test suite

- Porting MiniAero to Charm++ began with a “bootcamp”:
 - March 9-12, 2015
 - Led by Nikhil Jain and Eric Mikida
 - About 10 Sandia scientists in attendance
- Since the workshop, we've had one scientist working 50% time on the port and a couple others working 10-20% time
- Current state of the code:
 - running and passes test suite
 - most immediately apparent optimizations done

- Porting MiniAero to Charm++ began with a “bootcamp”:
 - March 9-12, 2015
 - Led by Nikhil Jain and Eric Mikida
 - About 10 Sandia scientists in attendance
- Since the workshop, we've had one scientist working 50% time on the port and a couple others working 10-20% time
- Current state of the code:
 - running and passes test suite
 - most immediately apparent optimizations done
 - SMP version does not work (Kokkos incompatibility)

Introduction

The Process: Porting an Explicit Aerodynamics Miniapp to the Chare Model

What was easy?

What was harder?

Preliminary Results and Performance

Next Steps

What was easy?

- Load balancing

- Load balancing (synchronous)

■ Load balancing (synchronous)

```
1 // at the end of the timestep...
2 if(doLoadBalance && timestepCounter % loadBalanceInterval == 0) {
3   serial {
4     // Do the actual load rebalancing
5     this->AtSync();
6   }
7   // Called when load balancing is completed (required)
8   when ResumeFromSync() { }
9 }
```

■ Load balancing (synchronous)

```
1 // at the end of the timestep...
2 if(doLoadBalance && timestepCounter % loadBalanceInterval == 0) {
3   serial {
4     // Do the actual load rebalancing
5     this->AtSync();
6   }
7   // Called when load balancing is completed (required)
8   when ResumeFromSync() { }
9 }
```

■ Checkpointing

■ Load balancing (synchronous)

```
1 // at the end of the timestep...
2 if(doLoadBalance && timestepCounter % loadBalanceInterval == 0) {
3   serial {
4     // Do the actual load rebalancing
5     this->AtSync();
6   }
7   // Called when load balancing is completed (required)
8   when ResumeFromSync() { }
9 }
```

■ Checkpointing

- To disk

■ Load balancing (synchronous)

```
1 // at the end of the timestep...
2 if(doLoadBalance && timestepCounter % loadBalanceInterval == 0) {
3   serial {
4     // Do the actual load rebalancing
5     this->AtSync();
6   }
7   // Called when load balancing is completed (required)
8   when ResumeFromSync() { }
9 }
```

■ Checkpointing

- To disk: `CkStartCheckpoint(...)`

■ Load balancing (synchronous)

```
1 // at the end of the timestep...
2 if(doLoadBalance && timestepCounter % loadBalanceInterval == 0) {
3   serial {
4     // Do the actual load rebalancing
5     this->AtSync();
6   }
7   // Called when load balancing is completed (required)
8   when ResumeFromSync() { }
9 }
```

■ Checkpointing

- To disk: `CkStartCheckpoint(...)`
- In memory (to partner node)

■ Load balancing (synchronous)

```
1 // at the end of the timestep...
2 if(doLoadBalance && timestepCounter % loadBalanceInterval == 0) {
3   serial {
4     // Do the actual load rebalancing
5     this->AtSync();
6   }
7   // Called when load balancing is completed (required)
8   when ResumeFromSync() { }
9 }
```

■ Checkpointing

■ To disk: `CkStartCheckpoint(...)`

■ In memory (to partner node): `CkStartMemCheckpoint(...)`

■ Load balancing (synchronous)

```
1 // at the end of the timestep...
2 if(doLoadBalance && timestepCounter % loadBalanceInterval == 0) {
3   serial {
4     // Do the actual load rebalancing
5     this->AtSync();
6   }
7   // Called when load balancing is completed (required)
8   when ResumeFromSync() { }
9 }
```

■ Checkpointing

- To disk: `CkStartCheckpoint(...)`
- In memory (to partner node): `CkStartMemCheckpoint(...)`
- Must be done synchronously

■ Load balancing (synchronous)

```
1 // at the end of the timestep...
2 if(doLoadBalance && timestepCounter % loadBalanceInterval == 0) {
3   serial {
4     // Do the actual load rebalancing
5     this->AtSync();
6   }
7   // Called when load balancing is completed (required)
8   when ResumeFromSync() { }
9 }
```

■ Checkpointing

- To disk: `CkStartCheckpoint(...)`
 - In memory (to partner node): `CkStartMemCheckpoint(...)`
 - Must be done synchronously
- Both of these were key features we wanted to test in AMT runtimes, both were done essentially on the first day of coding

■ Load balancing (synchronous)

```
1 // at the end of the timestep...
2 if(doLoadBalance && timestepCounter % loadBalanceInterval == 0) {
3   serial {
4     // Do the actual load rebalancing
5     this->AtSync();
6   }
7   // Called when load balancing is completed (required)
8   when ResumeFromSync() { }
9 }
```

■ Checkpointing

- To disk: `CkStartCheckpoint(...)`
- In memory (to partner node): `CkStartMemCheckpoint(...)`
- Must be done synchronously

- Both of these were key features we wanted to test in AMT runtimes, both were done essentially on the first day of coding

Both of these require only serialization on the user side

MPI \Rightarrow Charm++ is relatively straightforward

- “Quick start” implementation: map one chore to one MPI process

- “Quick start” implementation: map one chare to one MPI process
 - Just like in MPI, data dependencies are expressed in terms of messages:

- “Quick start” implementation: map one chare to one MPI process
 - Just like in MPI, data dependencies are expressed in terms of messages:
 - sends become message-like function calls on proxies of `array` members

- “Quick start” implementation: map one chore to one MPI process
 - Just like in MPI, data dependencies are expressed in terms of messages:
 - sends become message-like function calls on proxies of `array` members
 - receives become `when` clauses

- “Quick start” implementation: map one chore to one MPI process
 - Just like in MPI, data dependencies are expressed in terms of messages:
 - sends become message-like function calls on proxies of `array` members
 - receives become `when` clauses

```
class OldStuffDoer {
  /* ... */
  void do_stuff() {
    generate_data();
    /* ... */
    MPI_Irecv(data, n_send,
              MPI_DOUBLE, partner, /*...*/);
    MPI_Send(other_data, n_recv,
              MPI_DOUBLE, partner, /*...*/);
    use_other_data();
  }
};
```

- “Quick start” implementation: map one chore to one MPI process
 - Just like in MPI, data dependencies are expressed in terms of messages:
 - sends become message-like function calls on proxies of `array` members
 - receives become `when` clauses

```
class OldStuffDoer {
  /* ... */
  void do_stuff() {
    generate data();
    /* ... */
    MPI_Irecv(data, n_send,
              MPI_DOUBLE, partner, /*...*/);
    MPI_Send(other_data, n_recv,
             MPI_DOUBLE, partner, /*...*/);
    use_other_data();
  }
};
```

```
array [1D] NewStuffDoer {
  entry void receive_data(int src,
                          int ndata, double data[ndata]);
  entry void do_stuff_1() {
    generate data();
    /* ... */
    thisProxy[partner].receive_data(
      n_send, data
    );
  };
  entry void do_stuff_2() {
    when receive_data(int partner,
                     int n_send, double* other_data)
    serial {
      memcpy(other_data_, data,
             n_send * sizeof(double));
      use_other_data();
    }
  };
};
```

- “Quick start” implementation: map one chare to one MPI process
 - Just like in MPI, data dependencies are expressed in terms of messages:
 - sends become message-like function calls on proxies of `array` members
 - receives become `when` clauses

“Gotchas”:

```
class OldStuffDoer {
  /* ... */
  void do_stuff() {
    generate_data();
    /* ... */
    MPI_Irecv(&data, 1, MPI_DOUBLE, partner, 0, MPI_STATUS_IGNORE);
    MPI_Send(&other_data, 1, MPI_DOUBLE, partner, /*...*/);
    use_other_data();
  }
};
```

```
array [1D] NewStuffDoer {
  entry void receive_data(int src, int n_recv, double data[n_data]);
  void receive_data(int src, int n_recv, double data[n_data]) {
    /* ... */
    use_other_data();
  }
};

void NewStuffDoer::receive_data(
  int src, int n_recv, double data[n_data]) {
  /* ... */
  use_other_data();
}

int main() {
  /* ... */
  NewStuffDoer new_stuff_doer;
  new_stuff_doer.receive_data(
    /* ... */
  );
};
```

- “Quick start” implementation: map one chare to one MPI process
 - Just like in MPI, data dependencies are expressed in terms of messages:
 - sends become message-like function calls on proxies of `array` members
 - receives become `when` clauses

“Gotchas”:

- static variables

```
class OldStuffDoer {
  /* ... */
  void do_stuff() {
    generate_data();
    /* ... */
    MPI_Irecv(&data, 1, MPI_DOUBLE,
              MPI_ANY_SOURCE, MPI_ANY_TAG,
              MPI_Send(other_data, 1, MPI_DOUBLE, partner, /*...*/);
    use_other_data();
  }
};
```

```
array [1D] NewStuffDoer {
  entry void receive_data(int src,
                          double data[n_data]);
  void receive_data(int src, double data[n_data]) {
    // ...
    .receive_data(
    // ...
  }
};

NewStuffDoer {
  int n_send, double* other_data)
  serial {
    memcpy(other_data_, data,
           n_send * sizeof(double));
    use_other_data();
  }
};
```

- “Quick start” implementation: map one chare to one MPI process
 - Just like in MPI, data dependencies are expressed in terms of messages:
 - sends become message-like function calls on proxies of `array` members
 - receives become `when` clauses

“Gotchas”:

- static variables
- conditional communication

```
class OldStuffDoer {
  /* ... */
  void do_stuff() {
    generate_data();
    /* ... */
    MPI_Irecv(&data, 1, MPI_DOUBLE,
             MPI_ANY_SOURCE, MPI_ANY_TAG,
             MPI_Send(other_data, 1, MPI_DOUBLE, partner, /*...*/);
    use_other_data();
  }
};
```

```
array [1D] NewStuffDoer {
  entry void receive_data(int src,
                          int n_recv, double data[ndata]);
  void receive_data(int src, int n_recv, double data[ndata]) {
    // ...
    .receive_data(
        src, n_recv, data);
  }
};

NewStuffDoer {
  int n_send, double* other_data;
  serial {
    memcpy(other_data, data,
           n_send * sizeof(double));
    use_other_data();
  }
};
```

- “Quick start” implementation: map one chore to one MPI process
 - Just like in MPI, data dependencies are expressed in terms of messages:
 - sends become message-like function calls on proxies of `array` members
 - receives become `when` clauses

```
class OldStuffDoer {
  /* ... */
  void do_stuff() {
    generate_data();
    /* ... */
    MPI_Irecv(&data, 1, MPI_DOUBLE, partner, 0, comm);
    MPI_Send(&other_data, 1, MPI_DOUBLE, partner, /*...*/);
    use_other_data();
  }
};
```

“Gotchas”:

- static variables
- conditional communication
- a lot of size and metadata communication setup can be skipped

```
array [1D] NewStuffDoer {
  entry void receive_data(int src,
    double data[ndata]);
  void receive_data(int src, double data[ndata]) {
    // ...
    .receive_data(
      src, data);
  }
};

NewStuffDoer {
  int n_send, double* other_data;
  serial {
    memcpy(other_data_, data,
      n_send * sizeof(double));
    use_other_data();
  }
};
```

- Is this the best approach for our workloads, or does it lead to unnecessary synchronizations being left over from the MPI version?

- Is this the best approach for our workloads, or does it lead to unnecessary synchronizations being left over from the MPI version?
- Two approaches:

- Is this the best approach for our workloads, or does it lead to unnecessary synchronizations being left over from the MPI version?
- Two approaches:
 - “Bottom up”: Map sends and receives to function calls and `when s`

- Is this the best approach for our workloads, or does it lead to unnecessary synchronizations being left over from the MPI version?
- Two approaches:
 - “Bottom up”: Map sends and receives to function calls and `when s`
 - “Top down”: Think about task structure and dependencies of code, write this into the `.ci` file

- Is this the best approach for our workloads, or does it lead to unnecessary synchronizations being left over from the MPI version?
- Two approaches:
 - “Bottom up”: Map sends and receives to function calls and `when s`
 - “Top down”: Think about task structure and dependencies of code, write this into the `.ci` file
- Clearly, “top down” approach will lead to better, more efficient code in most cases, but...

- Is this the best approach for our workloads, or does it lead to unnecessary synchronizations being left over from the MPI version?
- Two approaches:
 - “Bottom up”: Map sends and receives to function calls and `when s`
 - “Top down”: Think about task structure and dependencies of code, write this into the `.ci` file
- Clearly, “top down” approach will lead to better, more efficient code in most cases, but...
- For production code, a complete “top down” overhaul is completely impractical

- Is this the best approach for our workloads, or does it lead to unnecessary synchronizations being left over from the MPI version?
- Two approaches:
 - “Bottom up”: Map sends and receives to function calls and `when s`
 - “Top down”: Think about task structure and dependencies of code, write this into the `.ci` file
- Clearly, “top down” approach will lead to better, more efficient code in most cases, but...
- For production code, a complete “top down” overhaul is completely impractical
- Is there a good middle ground?

- Is this the best approach for our workloads, or does it lead to unnecessary synchronizations being left over from the MPI version?
- Two approaches:
 - “Bottom up”: Map sends and receives to function calls and `when s`
 - “Top down”: Think about task structure and dependencies of code, write this into the `.ci` file
- Clearly, “top down” approach will lead to better, more efficient code in most cases, but...
- For production code, a complete “top down” overhaul is completely impractical
- Is there a good middle ground?
 - “Bottom up”-ness vs. “top down”-ness of approach should be assessed before writing too much code (in any porting project)

What was harder: Kokkos integration and templated code

- Kokkos is a performance portability layer aimed primarily at on-node parallelism

- Kokkos is a performance portability layer aimed primarily at on-node parallelism
 - handles memory layout and loop structure to produce optimized kernels on multiple devices

- Kokkos is a performance portability layer aimed primarily at on-node parallelism
 - handles memory layout and loop structure to produce optimized kernels on multiple devices
- Application developer implements generic code, Kokkos library implements device-specific specializations

- Kokkos is a performance portability layer aimed primarily at on-node parallelism
 - handles memory layout and loop structure to produce optimized kernels on multiple devices
- Application developer implements generic code, Kokkos library implements device-specific specializations

```
1 template <typename Device>
2 struct ddot {
3     const Kokkos::View<Device>& A, B;
4     double result;
5
6     ddot(
7         const Kokkos::View<Device>& A_in,
8         const Kokkos::View<Device>& B_in
9     ) : A(A_in), B(B_in), result(0)
10    { }
11
12    inline void operator()(int i) {
13        result += A(i) * B(i);
14    }
15 };
16
17 void do_stuff() {
18     /* ... */
19     Kokkos::parallel_for(
20         num_items,
21         ddot<Kokkos::Cuda>(v1, v2)
22     );
23 }
```

- Kokkos is a performance portability layer aimed primarily at on-node parallelism
 - handles memory layout and loop structure to produce optimized kernels on multiple devices
- Application developer implements generic code, Kokkos library implements device-specific specializations
- MiniAero was originally written in “MPI+Kokkos”

```
1 template <typename Device>
2 struct ddot {
3     const Kokkos::View<Device>& A, B;
4     double result;
5
6     ddot(
7         const Kokkos::View<Device>& A_in,
8         const Kokkos::View<Device>& B_in
9     ) : A(A_in), B(B_in), result(0)
10    { }
11
12    inline void operator()(int i) {
13        result += A(i) * B(i);
14    }
15 };
16
17 void do_stuff() {
18     /* ... */
19     Kokkos::parallel_for(
20         num_items,
21         ddot<Kokkos::Cuda>(v1, v2)
22     );
23 }
```

- Kokkos is a performance portability layer aimed primarily at on-node parallelism
 - handles memory layout and loop structure to produce optimized kernels on multiple devices
- Application developer implements generic code, Kokkos library implements device-specific specializations
- MiniAero was originally written in “MPI+Kokkos”
- What happens when you need to write templated code that uses Kokkos?

```
1 template <typename Device>
2 struct ddot {
3     const Kokkos::View<Device>& A, B;
4     double result;
5
6     ddot(
7         const Kokkos::View<Device>& A_in,
8         const Kokkos::View<Device>& B_in
9     ) : A(A_in), B(B_in), result(0)
10    { }
11
12    inline void operator()(int i) {
13        result += A(i) * B(i);
14    }
15};
16
17 void do_stuff() {
18     /* ... */
19     Kokkos::parallel_for(
20         num_items,
21         ddot<Kokkos::Cuda>(v1, v2)
22     );
23 }
```

- Kokkos is a performance portability layer aimed primarily at on-node parallelism
 - handles memory layout and loop structure to produce optimized kernels on multiple devices
- Application developer implements generic code, Kokkos library implements device-specific specializations
- MiniAero was originally written in “MPI+Kokkos”
- What happens when you need to write templated code that uses Kokkos?
 - Explicitly listing all specializations can get out of hand quickly. For instance...

```
1 template <typename Device>
2 struct ddot {
3     const Kokkos::View<Device>& A, B;
4     double result;
5
6     ddot(
7         const Kokkos::View<Device>& A_in,
8         const Kokkos::View<Device>& B_in
9     ) : A(A_in), B(B_in), result(0)
10    { }
11
12    inline void operator()(int i) {
13        result += A(i) * B(i);
14    }
15 };
16
17 void do_stuff() {
18     /* ... */
19     Kokkos::parallel_for(
20         num_items,
21         ddot<Kokkos::Cuda>(v1, v2)
22     );
23 }
```

- The MiniAero solver has five different ghost exchanges.

- The MiniAero solver has five different ghost exchanges.
- Each communicates a different `Kokkos::View` type

- The MiniAero solver has five different ghost exchanges.
- Each communicates a different `Kokkos::View` type, so we want an entry method prototype that looks something like this:

- The MiniAero solver has five different ghost exchanges.
- Each communicates a different `Kokkos::View` type, so we want an entry method prototype that looks something like this:

```
1 template <typename ViewType>  
2 entry [local] void receive_ghost_data(ViewType& v);
```

- The MiniAero solver has five different ghost exchanges.
- Each communicates a different `Kokkos::View` type, so we want an entry method prototype that looks something like this:

```
1 template <typename ViewType>  
2 entry [local] void receive_ghost_data(ViewType& v);
```

- The solver chore is already parameterized on the Kokkos device type:

- The MiniAero solver has five different ghost exchanges.
- Each communicates a different `Kokkos::View` type, so we want an entry method prototype that looks something like this:

```
1 template <typename ViewType>  
2 entry [local] void receive_ghost_data(ViewType& v);
```

- The solver chore is already parameterized on the Kokkos device type:

```
1 /* solver.ci */  
2 template <typename Device>  
3 array [1D] RK4Solver {  
4     /* ... */  
5 };
```

- The MiniAero solver has five different ghost exchanges.
- Each communicates a different `Kokkos::View` type, so we want an entry method prototype that looks something like this:

```
1 template <typename ViewType>
2 entry [local] void receive_ghost_data(ViewType& v);
```

- The solver chare is already parameterized on the Kokkos device type:

```
1 /* solver.ci */
2 template <typename Device>
3 array [1D] RK4Solver {
4     /* ... */
5 };
```

```
1 /* solver.h */
2 template <typename Device>
3 class RK4Solver
4     : public CBase_RK4Solver<Device>
5 {
6     Kokkos::View<Device, double*, 5> m_data1;
7     Kokkos::View<Device, double*, 5, 3> m_data2;
8     Kokkos::View<Device, int*> m_data3;
9     /* etc... */
10 };
```

- The MiniAero solver has five different ghost exchanges.
- Each communicates a different `Kokkos::View` type, so we want an entry method prototype that looks something like this:

```
1 template <typename ViewType>
2 entry [local] void receive_ghost_data(ViewType& v);
```

- The solver chare is already parameterized on the Kokkos device type:

```
1 /* solver.h */
2 template <typename Device>
3 class RK4Solver
4   : public CBase_RK4Solver<Device>
5   {
6     Kokkos::View<Device, double*, 5> m_data1;
7     Kokkos::View<Device, double*, 5, 3> m_data2;
8     Kokkos::View<Device, int*> m_data3;
9     /* etc... */
10  };

1 /* solver.ci */
2 template <typename Device>
3 array [1D] RK4Solver {
4   /* ... */
5 };
```

- The devices we'd like to test include `Kokkos::Serial`, `Kokkos::Threads`, `Kokkos::Cuda`, and `Kokkos::OpenMP`

- The MiniAero solver has five different ghost exchanges.
- Each communicates a different `Kokkos::View` type, so we want an entry method prototype that looks something like this:

```
1 template <typename ViewType>
2 entry [local] void receive_ghost_data(ViewType& v);
```

- The solver chare is already parameterized on the Kokkos device type:

```
1 /* solver.ci */
2 template <typename Device>
3 array [1D] RK4Solver {
4     /* ... */
5 };
6
7 /* solver.h */
8 template <typename Device>
9 class RK4Solver
10 : public CBase_RK4Solver<Device>
11 {
12     Kokkos::View<Device, double*, 5> m_data1;
13     Kokkos::View<Device, double*, 5, 3> m_data2;
14     Kokkos::View<Device, int*> m_data3;
15     /* etc... */
16 };
```

- The devices we'd like to test include `Kokkos::Serial`, `Kokkos::Threads`, `Kokkos::Cuda`, and `Kokkos::OpenMP`
- That already leads to *20 different explicit signatures* for `receive_ghost_data()`.

- Pattern: templated setup, non-templated entry method, templated cleanup

- Pattern: templated setup, non-templated entry method, templated cleanup

```
/* comm_stuff.h */
template <typename Device>
class CommStuffDoer :
public CBase_CommStuffDoer<Device>
{
    Kokkos::View<Device, double*, 3> my_data_1_;
    Kokkos::View<Device, int*, 3, 5> my_data_2_;
    /* ... */
};
```


- Pattern: templated setup, non-templated entry method, templated cleanup

```
/* comm_stuff.h */
template <typename Device>
class CommStuffDoer :
public CBase_CommStuffDoer<Device>
{
    Kokkos::View<Device, double*, 3> my_data_1_;
    Kokkos::View<Device, int*, 3, 5> my_data_2_;
    /* ... */
};
```

```
/* comm_stuff.ci */
template <typename Device>
array [1D] CommStuffDoer {

    entry void do_stuff() {

};
};
```

■ Pattern: templated setup, non-templated entry method, templated cleanup

```
/* comm_stuff.h */
template <typename Device>
class CommStuffDoer :
public CBase_CommStuffDoer<Device>
{
    Kokkos::View<Device, double*, 3> my_data_1_;
    Kokkos::View<Device, int*, 3, 5> my_data_2_;
    /* ... */
};
```

```
/* comm_stuff.ci */
template <typename Device>
array [1D] CommStuffDoer {

    entry void do_stuff() {
        /* ... */
        serial {
            int src = /*...*/, dest = /*...*/;
            send_it(dest, my_data_1_);

        }

    };
};
```

■ Pattern: templated setup, non-templated entry method, templated cleanup

```
/* comm_stuff.h */
template <typename Device>
class CommStuffDoer :
public CBase_CommStuffDoer<Device>
{
    Kokkos::View<Device, double*, 3> my_data_1_;
    Kokkos::View<Device, int*, 3, 5> my_data_2_;
    /* ... */
};
```

```
/* comm_stuff.ci */
template <typename Device>
array [1D] CommStuffDoer {

    entry void do_stuff() {
        /* ... */
        serial {
            int src = /*...*/, dest = /*...*/;
            send_it(dest, my_data_1_);
            setup_recv(src, my_data_1_);
        }
    };
};
```

■ Pattern: templated setup, non-templated entry method, templated cleanup

```
/* comm_stuff.h */
template <typename Device>
class CommStuffDoer :
public CBase_CommStuffDoer<Device>
{
    Kokkos::View<Device, double*, 3> my_data_1_;
    Kokkos::View<Device, int*, 3, 5> my_data_2_;
    /* ... */
};
```

```
/* comm_stuff.ci */
template <typename Device>
array [1D] CommStuffDoer {

    entry void do_stuff() {
        /* ... */
        serial {
            int src = /*...*/, dest = /*...*/;
            send_it(dest, my_data_1_);
            setup_recv(src, my_data_1_);
            do_recv(src);
        }
    };
};
```

■ Pattern: templated setup, non-templated entry method, templated cleanup

```
/* comm_stuff.h */
template <typename Device>
class CommStuffDoer :
public CBase_CommStuffDoer<Device>
{
    Kokkos::View<Device,double*,3> my_data_1_;
    Kokkos::View<Device,int*,3,5> my_data_2_;
    /* ... */
};
```

```
/* comm_stuff.ci */
template <typename Device>
array [1D] CommStuffDoer {

    entry void do_recv_done();

    entry void do_stuff() {
        /* ... */
        serial {
            int src = /*...*/, dest = /*...*/;
            send_it(dest, my_data_1_);
            setup_recv(src, my_data_1_);
            do_recv(src);
        }
        when do_recv_done() serial {

        }
    }
};
```

■ Pattern: templated setup, non-templated entry method, templated cleanup

```
/* comm_stuff.h */
template <typename Device>
class CommStuffDoer :
{
    public CBase_CommStuffDoer<Device>
{
    Kokkos::View<Device,double*,3> my_data_1_;
    Kokkos::View<Device,int*,3,5> my_data_2_;
    /* ... */
};
```

```
/* comm_stuff.ci */
template <typename Device>
array [1D] CommStuffDoer {

    entry void do_rcv_done();

    entry void do_stuff() {
        /* ... */
        serial {
            int src = /*...*/, dest = /*...*/;
            send_it(dest, my_data_1_);
            setup_rcv(src, my_data_1_);
            do_rcv(src);
        }
        when do_rcv_done() serial {
            finish_rcv(src, my_data_1_);
        }
    };
};
```


■ Pattern: templated setup, non-templated entry method, templated cleanup

```
/* comm_stuff.h */
template <typename Device>
class CommStuffDoer :
public CBase_CommStuffDoer<Device>
{
    Kokkos::View<Device,double*,3> my_data_1_;
    Kokkos::View<Device,int*,3,5> my_data_2_;
    /* ... */

    template <typename ViewT>
    void send_it(int dst, const ViewT& data) {
        size_t size = get_size(data, dst);
        double* data = extract_data(data, dst);
        this->thisProxy[dst].recv_it(
            this->thisIndex, size, data);
    }
};
```

```
/* comm_stuff.ci */
template <typename Device>
array [1D] CommStuffDoer {
    entry void recv_it(int src,
        int size, double data[size]);
    entry void do_recv_done();

    entry void do_stuff() {
        /* ... */
        serial {
            int src = /*...*/, dest = /*...*/;
            send_it(dest, my_data_1_);
            setup_recv(src, my_data_1_);
            do_recv(src);
        }
        when do_recv_done() serial {
            finish_recv(src, my_data_1_);
        }
    };
};
```

■ Pattern: templated setup, non-templated entry method, templated cleanup

```
/* comm_stuff.h */
template <typename Device>
class CommStuffDoer :
public CBase_CommStuffDoer<Device>
{
    Kokkos::View<Device,double*,3> my_data_1_;
    Kokkos::View<Device,int*,3,5> my_data_2_;
    /* ... */
    std::vector<double*> recv_buffers_;
    template <typename ViewT>
    void send_it(int dst, const ViewT& data) {
        size_t size = get_size(data, dst);
        double* data = extract_data(data, dst);
        this->thisProxy[dst].recv_it(
            this->thisIndex, size, data);
    }
    template <typename ViewT>
    void setup_recv(int src, ViewT& data) {
        recv_buffers_[src] =
            get_buffer(data, src);
    }
};
```

```
/* comm_stuff.ci */
template <typename Device>
array [1D] CommStuffDoer {
    entry void recv_it(int src,
        int size, double data[size]);
    entry void do_recv_done();

    entry void do_stuff() {
        /* ... */
        serial {
            int src = /*...*/, dest = /*...*/;
            send_it(dest, my_data_1_);
            setup_recv(src, my_data_1_);
            do_recv(src);
        }
        when do_recv_done() serial {
            finish_recv(src, my_data_1_);
        }
    };
};
```

■ Pattern: templated setup, non-templated entry method, templated cleanup

```
/* comm_stuff.h */
template <typename Device>
class CommStuffDoer :
public CBase_CommStuffDoer<Device>
{
    Kokkos::View<Device,double*,3> my_data_1_;
    Kokkos::View<Device,int*,3,5> my_data_2_;
    /* ... */
    std::vector<double*> recv_buffers_;
    template <typename ViewT>
    void send_it(int dst, const ViewT& data) {
        size_t size = get_size(data, dst);
        double* data = extract_data(data, dst);
        this->thisProxy[dst].recv_it(
            this->thisIndex, size, data);
    }
    template <typename ViewT>
    void setup_recv(int src, ViewT& data) {
        recv_buffers_[src] =
            get_buffer(data, src);
    }
    template <typename ViewT>
    void finish_recv(int src, ViewT& data) {
        insert_data(data, recv_buffers_[src], src);
        delete recv_buffers_[src];
    }
};
```

```
/* comm_stuff.ci */
template <typename Device>
array [1D] CommStuffDoer {
    entry void recv_it(int src,
        int size, double data[size]);
    entry void do_recv_done();

    entry void do_stuff() {
        /* ... */
        serial {
            int src = /*...*/, dest = /*...*/;
            send_it(dest, my_data_1_);
            setup_recv(src, my_data_1_);
            do_recv(src);
        }
        when do_recv_done() serial {
            finish_recv(src, my_data_1_);
        }
    };
};
```

■ Pattern: templated setup, non-templated entry method, templated cleanup

```
/* comm_stuff.h */
template <typename Device>
class CommStuffDoer :
public CBase_CommStuffDoer<Device>
{
    Kokkos::View<Device,double*,3> my_data_1_;
    Kokkos::View<Device,int*,3,5> my_data_2_;
    /* ... */
    std::vector<double*> recv_buffers_;
    template <typename ViewT>
    void send_it(int dst, const ViewT& data) {
        size_t size = get_size(data, dst);
        double* data = extract_data(data, dst);
        this->thisProxy[dst].recv_it(
            this->thisIndex, size, data);
    }
    template <typename ViewT>
    void setup_recv(int src, ViewT& data) {
        recv_buffers_[src] =
            get_buffer(data, src);
    }
    template <typename ViewT>
    void finish_recv(int src, ViewT& data) {
        insert_data(data, recv_buffers_[src], src);
        delete recv_buffers_[src];
    }
};
```

```
/* comm_stuff.ci */
template <typename Device>
array [1D] CommStuffDoer {
    entry void recv_it(int src,
        int size, double data[size]);
    entry void do_recv_done();
    entry [local] void do_recv(int src) {
        when recv_it[src](int s, int size,
            double data[size]) serial {
            memcpy(recv_buffers[src], data,
                size*sizeof(double));
            do_recv_done();
        }
    };
    entry void do_stuff() {
        /* ... */
        serial {
            int src = /*...*/, dest = /*...*/;
            send_it(dest, my_data_1_);
            setup_recv(src, my_data_1_);
            do_recv(src);
        }
        when do_recv_done() serial {
            finish_recv(src, my_data_1_);
        }
    };
};
```

■ Pattern: templated setup, non-templated entry method, templated cleanup

```
/* comm_stuff.h */
template <typename Device>
class CommStuffDoer :
public CBase_CommStuffDoer<Device>
{
    Kokkos::View<Device,double*,3> my_data_1_;
    Kokkos::View<Device,int*,3,5> my_data_2_;
    /* ... */
    std::vector<double*> recv_buffers_;
    template <typename ViewT>
    void send_it(int dst, const ViewT& data) {
        size_t size = get_size(data,
        double* data = extract_data(
        this->thisProxy[dst].recv_it
        this->thisIndex, size, data
    }
    template <typename ViewT>
    void setup_recv(int src, ViewT& data) {
        recv_buffers_[src] =
        get_buffer(data, src);
    }
    template <typename ViewT>
    void finish_recv(int src, ViewT& data) {
        insert_data(data, recv_buffers_[src], src);
        delete recv_buffers_[src];
    }
};
```

Is this ideal?
Obviously not

```
/* comm_stuff.ci */
template <typename Device>
array [1D] CommStuffDoer {
    entry void recv_it(int src,
        int size, double data[size]);
    entry void do_recv_done();
    entry [local] void do_recv(int src) {
        when recv_it[src](int s, int size,
            double data[size]) serial {
            memcpy(recv_buffers[src], data,
                size*sizeof(double));
            recv_done();
        }
    }
    void do_stuff() {
        /* ... */
        serial {
            int src = /*...*/, dest = /*...*/;
            send_it(dest, my_data_1_);
            setup_recv(src, my_data_1_);
            do_recv(src);
        }
        when do_recv_done() serial {
            finish_recv(src, my_data_1_);
        }
    };
};
```

■ Pattern: templated setup, non-templated entry method, templated cleanup

```
/* comm_stuff.h */
template <typename Device>
class CommStuffDoer :
public CBase_CommStuffDoer<Device>
{
    Kokkos::View<Device, double*, 3> my_data_1_;
    Kokkos::View<Device, int*, 3, 5> my_data_2_;
    /* ... */
    std::vector<double*> r
    template <typename ViewT>
    void send_it(int dst,
                size_t size = get_size(),
                double* data = extra_data,
                this->thisProxy[dst],
                this->thisIndex, s
    }
    template <typename ViewT>
    void setup_recv(int src, ViewT& data,
                   recv_buffers_[src] =
                   get_buffer(data, src);
    }
    template <typename ViewT>
    void finish_recv(int src, ViewT& data) {
        insert_data(data, recv_buffers_[src], src);
        delete recv_buffers_[src];
    }
};
```

```
/* comm_stuff.ci */
template <typename Device>
array [1D] CommStuffDoer {
    entry void recv_it(int src,
                       int size, double data[size]);
    entry void do_recv_done();
    entry [local] void do_recv(int src) {
        when recv_it[src](int s, int size,
                          data[size]) serial {
            buffers[src], data,
            of(double));
            e();
        }
        buff() {
            int src = /*...*/, dest = /*...*/;
            send_it(dest, my_data_1_);
            setup_recv(src, my_data_1_);
            do_recv(src);
        }
        when do_recv_done() serial {
            finish_recv(src, my_data_1_);
        }
    };
};
```

Is this typical of the effort required to make templated code work with an asynchronous many-task runtime system (AMT RTS)?

Maybe

■ Pattern: templated setup, non-templated entry method, templated cleanup

```
/* comm_stuff.h */
template <typename Device>
class CommStuffDoer :
public CBase_CommStuffDoer<Device>
{
Kokkos::View<Device,double*,3> my_data_1_;
Kokkos::View<Device,int*,3,5> my_data_2_;
/* ... */
std::vector<double*> recv_buffers :
template <typename ViewT>
void send_it(int dst,
size_t size = get_size(dst),
double* data = extra_data(dst),
this->thisProxy[dst],
this->thisIndex, size);
}
template <typename ViewT>
void setup_recv(int src, ViewT& data) {
recv_buffers_[src] =
get_buffer(data, src);
}
template <typename ViewT>
void finish_recv(int src, ViewT& data) {
insert_data(data, recv_buffers_[src], src);
delete recv_buffers_[src];
}
};
```

Does Charm++ even support
templated entry methods inside
templated chares?
(We couldn't figure out how to do it)

```
/* comm_stuff.ci */
template <typename Device>
array [1D] CommStuffDoer {
entry void recv_it(int src,
int size, double data[size]);
entry void do_recv_done();
entry [local] void do_recv(int src) {
when recv_it[src](int s, int size,
double data[size]) serial {
recv_buffers[src], data,
do_recv(s, data);
}
}
}
int src = /*...*/, dest = /*...*/;
send_it(dest, my_data_1_);
setup_recv(src, my_data_1_);
do_recv(src);
}
when do_recv_done() serial {
finish_recv(src, my_data_1_);
}
};
```



```
1 entry void do_stuff() {  
2     serial {  
3         do_stuff_1();  
4         do_stuff_2();  
5     }  
6 };
```

- Suppose all of the `do_stuff_*`() methods are ordinary, non-entry methods.

```
1 entry void do_stuff() {  
2     serial {  
3         do_stuff_1();  
4         do_stuff_2();  
5     }  
6 };
```

- Suppose all of the `do_stuff_*`() methods are ordinary, non-entry methods.
 - What happens first?

```
1 entry void do_stuff() {  
2     serial {  
3         do_stuff_1();  
4         do_stuff_2();  
5     }  
6 };
```

- Suppose all of the `do_stuff_*`() methods are ordinary, non-entry methods.
 - What happens first?
- Now suppose `do_stuff_1()` is an entry method and `do_stuff_2()` is a normal method.

```
1 entry void do_stuff() {  
2     serial {  
3         do_stuff_1();  
4         do_stuff_2();  
5     }  
6 };
```

- Suppose all of the `do_stuff_*`() methods are ordinary, non-entry methods.
 - What happens first?
- Now suppose `do_stuff_1()` is an entry method and `do_stuff_2()` is a normal method.
 - Now what happens first?

```
1 entry void do_stuff() {  
2     serial {  
3         do_stuff_1();  
4         do_stuff_2();  
5     }  
6 };
```

- Suppose all of the `do_stuff_*`() methods are ordinary, non-entry methods.
 - What happens first?
- Now suppose `do_stuff_1()` is an entry method and `do_stuff_2()` is a normal method.
 - Now what happens first?
- How does the programmer who *didn't* write `do_stuff_1()` know this?

```
1 entry void do_stuff() {  
2     serial {  
3         do_stuff_1();  
4         do_stuff_2();  
5     }  
6 };
```

- Suppose all of the `do_stuff_*`() methods are ordinary, non-entry methods.
 - What happens first?
- Now suppose `do_stuff_1()` is an entry method and `do_stuff_2()` is a normal method.
 - Now what happens first?
- How does the programmer who *didn't* write `do_stuff_1()` know this?
 - Perhaps using naming conventions? (e.g., `EM_*`)

```
1 entry void do_stuff() {
2     serial {
3         do_stuff_1();
4         do_stuff_2();
5     }
6 };
```

```
1 entry void EM_do_stuff() {
2     serial {
3         EM_do_stuff_1();
4         do_stuff_2();
5     }
6 };
```

Distinguishing Entry Methods from Regular Method Calls

- Suppose all of the `do_stuff_*`() methods are ordinary, non-entry methods.
 - What happens first?
- Now suppose `do_stuff_1()` is an entry method and `do_stuff_2()` is a normal method.
 - Now what happens first?
- How does the programmer who *didn't* write `do_stuff_1()` know this?
 - Perhaps using naming conventions? (e.g., `EM_*`())

```

1 entry void do_stuff() {
2     serial {
3         do_stuff_1();
4         do_stuff_2();
5     }
6 };

```

```

1 entry void EM_do_stuff() {
2     serial {
3         EM_do_stuff_1();
4         do_stuff_2();
5     }
6 };

```

In short, mixing entry method calls and regular method calls without using naming conventions makes it *difficult to write self-documenting code*

- Further complication:
non-blocking calls from a
blocking context

- Further complication:
non-blocking calls from a
blocking context

```
1 /* stuff_doer.ci */
2 chare StuffDoer {
3     entry void EM_do_stuff() {
4         serial {
5             EM_do_stuff_1();
6             do_stuff_2();
7             do_stuff_3();
8         }
9     };
10};
```

```
1 /* stuff_doer.h */
2 class StuffDoer
3     : public CBase_StuffDoer {
4     /*...*/
5     void do_stuff_2() {
6         // uh-oh
7         thisProxy.EM_do_stuff_4();
8     }
9};
```

- Further complication:
non-blocking calls from a
blocking context
- In fact, `do_stuff_2()` may
only be blocking *most* of the
time, but occasionally contain
non-blocking calls

```
1 /* stuff_doer.ci */
2 chare StuffDoer {
3     entry void EM_do_stuff() {
4         serial {
5             EM_do_stuff_1();
6             do_stuff_2();
7             do_stuff_3();
8         }
9     };
10};
```

```
1 /* stuff_doer.h */
2 class StuffDoer
3     : public CBase_StuffDoer {
4     /*...*/
5     void do_stuff_2() {
6         // uh-oh
7         thisProxy.EM_do_stuff_4();
8     }
9};
```

- Further complication: non-blocking calls from a blocking context
- In fact, `do_stuff_2()` may only be blocking *most* of the time, but occasionally contain non-blocking calls

```
1 /* stuff_doer.ci */
2 chare StuffDoer {
3     entry void EM_do_stuff() {
4         serial {
5             EM_do_stuff_1();
6             do_stuff_2();
7             do_stuff_3();
8         }
9     };
10 };
```

```
1 /* stuff_doer.h */
2 class StuffDoer
3 : public CBase_StuffDoer {
4     /*...*/
5     void do_stuff_2() {
6         if(some_rare_condition) {
7             thisProxy.EM_do_stuff_4();
8         }
9         /* ... */
10    }
11 };
```

- Further complication:
non-blocking calls from a
blocking context
- In fact, `do_stuff_2()` may
only be blocking *most* or the
time, but occasionally contain
non-blocking calls
- In this case, how does the
programmer make the control
flow of the program apparent to
future programmers?

```
1 /* stuff_doer.ci */
2 chare StuffDoer {
3     entry void EM_do_stuff() {
4         serial {
5             EM_do_stuff_1();
6             do_stuff_2();
7             do_stuff_3();
8         }
9     };
10 };
```

```
1 /* stuff_doer.h */
2 class StuffDoer
3 : public CBase_StuffDoer {
4     /*...*/
5     void do_stuff_2() {
6         if(some_rare_condition) {
7             thisProxy.EM_do_stuff_4();
8         }
9         /* ... */
10    }
11 };
```

- Further complication:
non-blocking calls from a
blocking context
- In fact, `do_stuff_2()` may
only be blocking *most* or the
time, but occasionally contain
non-blocking calls
- In this case, how does the
programmer make the control
flow of the program apparent to
future programmers?
 - Avoid writing code like this?

```
1 /* stuff_doer.ci */
2 chare StuffDoer {
3     entry void EM_do_stuff() {
4         serial {
5             EM_do_stuff_1();
6             do_stuff_2();
7             do_stuff_3();
8         }
9     };
10};
```

```
1 /* stuff_doer.h */
2 class StuffDoer
3 : public CBase_StuffDoer {
4     /*...*/
5     void do_stuff_2() {
6         if(some_rare_condition) {
7             thisProxy.EM_do_stuff_4();
8         }
9         /* ... */
10    }
11};
```

- Further complication:
non-blocking calls from a
blocking context
- In fact, `do_stuff_2()` may
only be blocking *most* of the
time, but occasionally contain
non-blocking calls
- In this case, how does the
programmer make the control
flow of the program apparent to
future programmers?
 - Avoid writing code like this?
 - Avoid naming conventions?
Makes the programmer “get
used to” the idea that *any*
method invocation in a `.ci` file
could be non-blocking

```
1 /* stuff_doer.ci */
2 chare StuffDoer {
3     entry void EM_do_stuff() {
4         serial {
5             EM_do_stuff_1();
6             do_stuff_2();
7             do_stuff_3();
8         }
9     };
10 };
```

```
1 /* stuff_doer.h */
2 class StuffDoer
3 : public CBase_StuffDoer {
4     /*...*/
5     void do_stuff_2() {
6         if(some_rare_condition) {
7             thisProxy.EM_do_stuff_4();
8         }
9     }
10 }
11 };
```

- Further complication:
non-blocking calls from a
blocking context
- In fact, `do_stuff_2()` may
only be blocking *most* of the
time, but occasionally contain
non-blocking calls
- In this case, how does the
programmer make the control
flow of the program apparent to
future programmers?
 - Avoid writing code like this?
 - Avoid naming conventions?
Makes the programmer “get
used to” the idea that *any*
method invocation in a `.ci` file
could be non-blocking
 - Just use comments?

```
1 /* stuff_doer.ci */
2 chare StuffDoer {
3     entry void EM_do_stuff() {
4         serial {
5             EM_do_stuff_1();
6             do_stuff_2();
7             do_stuff_3();
8         }
9     };
10};
```

```
1 /* stuff_doer.h */
2 class StuffDoer
3 : public CBase_StuffDoer {
4     /*...*/
5     void do_stuff_2() {
6         if(some_rare_condition) {
7             thisProxy.EM_do_stuff_4();
8         }
9     }
10 }
11};
```


Introduction

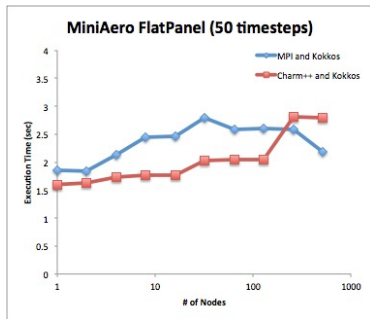
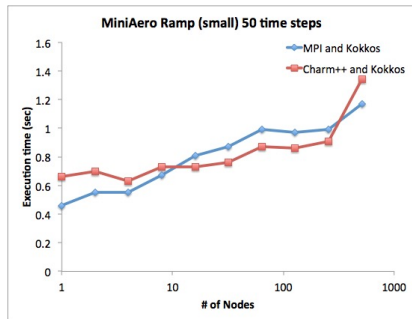
The Process: Porting an Explicit Aerodynamics Miniapp to the Chare Model

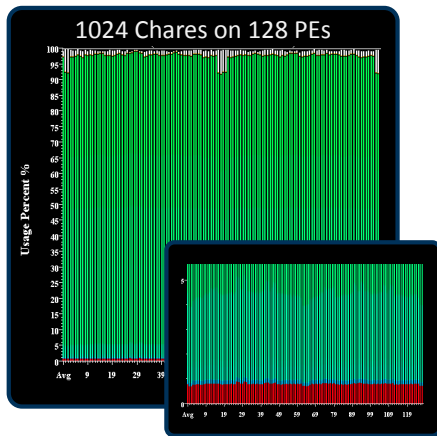
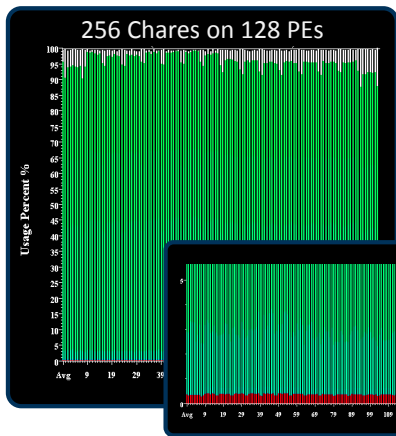
What was easy?

What was harder?

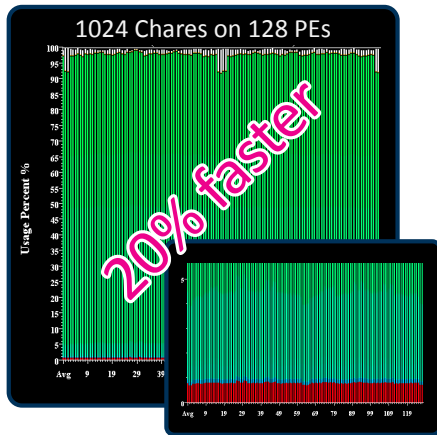
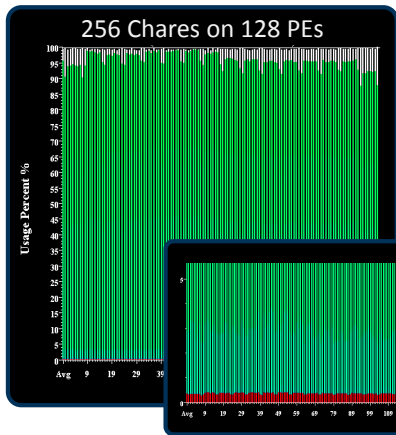
Preliminary Results and Performance

Next Steps





Application code in green, runtime overhead in red, idle time in white
(Insets are enlargements of y-axes)



Application code in green, runtime overhead in red, idle time in white
(Insets are enlargements of y-axes)

Introduction

The Process: Porting an Explicit Aerodynamics Miniapp to the Chare Model

What was easy?

What was harder?

Preliminary Results and Performance

Next Steps

- More performance analysis

- More performance analysis
 - PAPI?

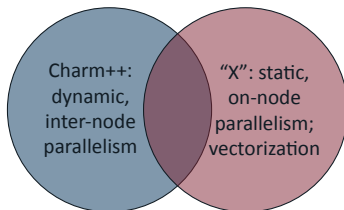
- More performance analysis
 - PAPI?
- More Kokkos devices (Cuda?)

- More performance analysis
 - PAPI?
- More Kokkos devices (Cuda?)
- More miniapps

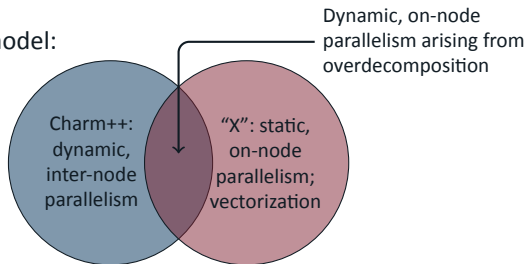
Questions?

Extra Slides

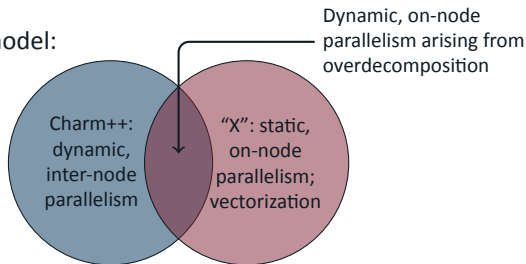
- The “Charm + X” model:



■ The “Charm + X” model:

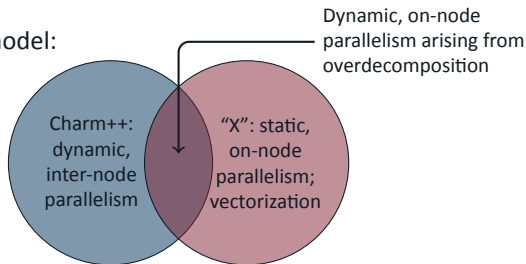


- The “Charm + X” model:



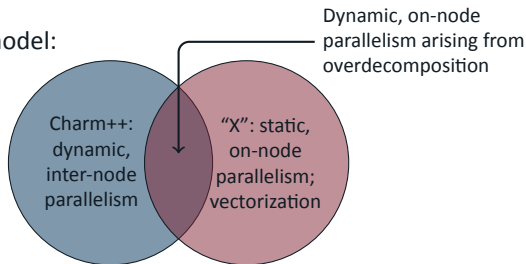
- Zero-copy semantics and some shared data model or data warehouse are critical to mitigating the AMT runtime overhead from overdecomposition

- The “Charm + X” model:



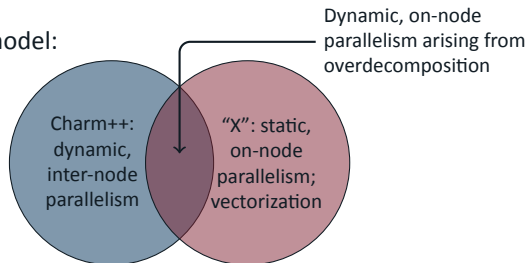
- Zero-copy semantics and some shared data model or data warehouse are critical to mitigating the AMT runtime overhead from overdecomposition
- Charm++ allows zero copy transfer of data between chares on-node using `PackedMessage s`

- The “Charm + X” model:



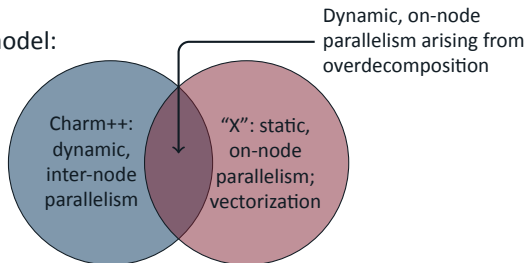
- Zero-copy semantics and some shared data model or data warehouse are critical to mitigating the AMT runtime overhead from overdecomposition
- Charm++ allows zero copy transfer of data between chares on-node using `PackedMessage s`
 - But these are an “advanced feature,” much more difficult than `PUP` ing

- The “Charm + X” model:



- Zero-copy semantics and some shared data model or data warehouse are critical to mitigating the AMT runtime overhead from overdecomposition
- Charm++ allows zero copy transfer of data between chares on-node using `PackedMessage s`
 - But these are an “advanced feature,” much more difficult than `PUP` ing
- The concept of a *shared data block* is missing

- The “Charm + X” model:



- Zero-copy semantics and some shared data model or data warehouse are critical to mitigating the AMT runtime overhead from overdecomposition
- Charm++ allows zero copy transfer of data between chares on-node using `PackedMessage s`
 - But these are an “advanced feature,” much more difficult than `PUP` ing
- The concept of a *shared data block* is missing
 - For instance, `PackedMessage s` have no access privileges (e.g., read only, shared read/write, exclusive read/write)

- The Charm++ compiler

- The Charm++ compiler
 - .ci file compiler issues

- The Charm++ compiler
 - .ci file compiler issues
 - (e.g., `}` inside C++ style comments not ignored?)

- The Charm++ compiler
 - .ci file compiler issues
 - (e.g., } inside C++ style comments not ignored?)
 - Would be nice if it could run like a preprocessor to generate code, then regular compiler could be used after that.

- The Charm++ compiler
 - .ci file compiler issues
 - (e.g., } inside C++ style comments not ignored?)
 - Would be nice if it could run like a preprocessor to generate code, then regular compiler could be used after that.
- SMP version: no way to initialize libraries on main thread?

- How does the programming experience in Charm++ compare to other runtimes?

- How does the programming experience in Charm++ compare to other runtimes?
 - *Inconclusive* so far. Charm++ MiniAero was a port, others were complete rewrites

- How does the programming experience in Charm++ compare to other runtimes?
 - *Inconclusive* so far. Charm++ MiniAero was a port, others were complete rewrites
- How does the performance of Charm++ compare to other runtimes?

- How does the programming experience in Charm++ compare to other runtimes?
 - *Inconclusive* so far. Charm++ MiniAero was a port, others were complete rewrites
- How does the performance of Charm++ compare to other runtimes?
 - *Inconclusive* so far. Other MiniAero versions are in various states of completeness

- How does the programming experience in Charm++ compare to other runtimes?
 - *Inconclusive* so far. Charm++ MiniAero was a port, others were complete rewrites
- How does the performance of Charm++ compare to other runtimes?
 - *Inconclusive* so far. Other MiniAero versions are in various states of completeness
 - *But...* our current implementation is already comparable to MPI