Charm++ Tutorial

Presented by:

Lukasz Wesolowski

Pritish Jetley

PARALLEL PROGRAMMING LABORATORY

Overview

- Introduction
  - Characteristics of a Parallel Language
  - Virtualization
  - Message Driven Execution
- Charm++ features
  - Chares and Chare Arrays
  - Parameter Marshalling
  - Structured Dagger Construct
  - Adaptive MPI
  - Load Balancing
- Tools
  - Parallel Debugger
  - Projections
  - LiveViz
- Conclusion

Outline

- Introduction
- Charm++ features
  - Chares and Chare Arrays
  - Parameter Marshalling
- Structured Dagger Construct
- Adaptive MPI
- Tools
  - Parallel Debugger
  - Projections
- Load Balancing
- LiveViz
- Conclusion

Characteristics of a Parallel Language

- Developing a parallel application involves:
  - decomposition
  - mapping
  - scheduling
  - machine-dependent expression
- Each task is either automated by the system or assigned to the programmer

Charm++ vs. MPI

	Charm++	MPI
Portabiliity	X	X
Scheduling	X	
Mapping	X	
Decomposition		

Virtualization: Object-based Decomposition

- Divide the computation into a large number of pieces
  - Independent of number of processors
  - Typically larger than number of processors
- Let the system map objects to processors

Object-based Parallelization

User is only concerned with interaction between objects

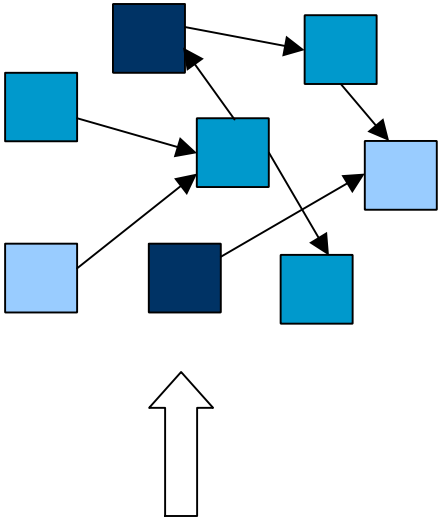

User View

System implementation

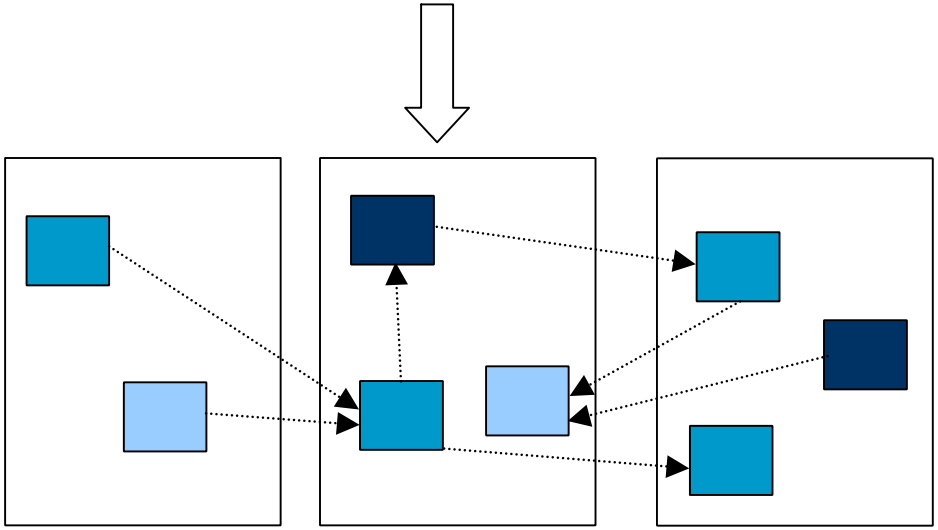

Message-Driven Execution

- Objects communicate asynchronously through remote method invocation
- Encourages non-deterministic execution
- Benefits:
  - Communication latency tolerance
  - Logical structure for scheduling

Message-Driven Execution in Charm++

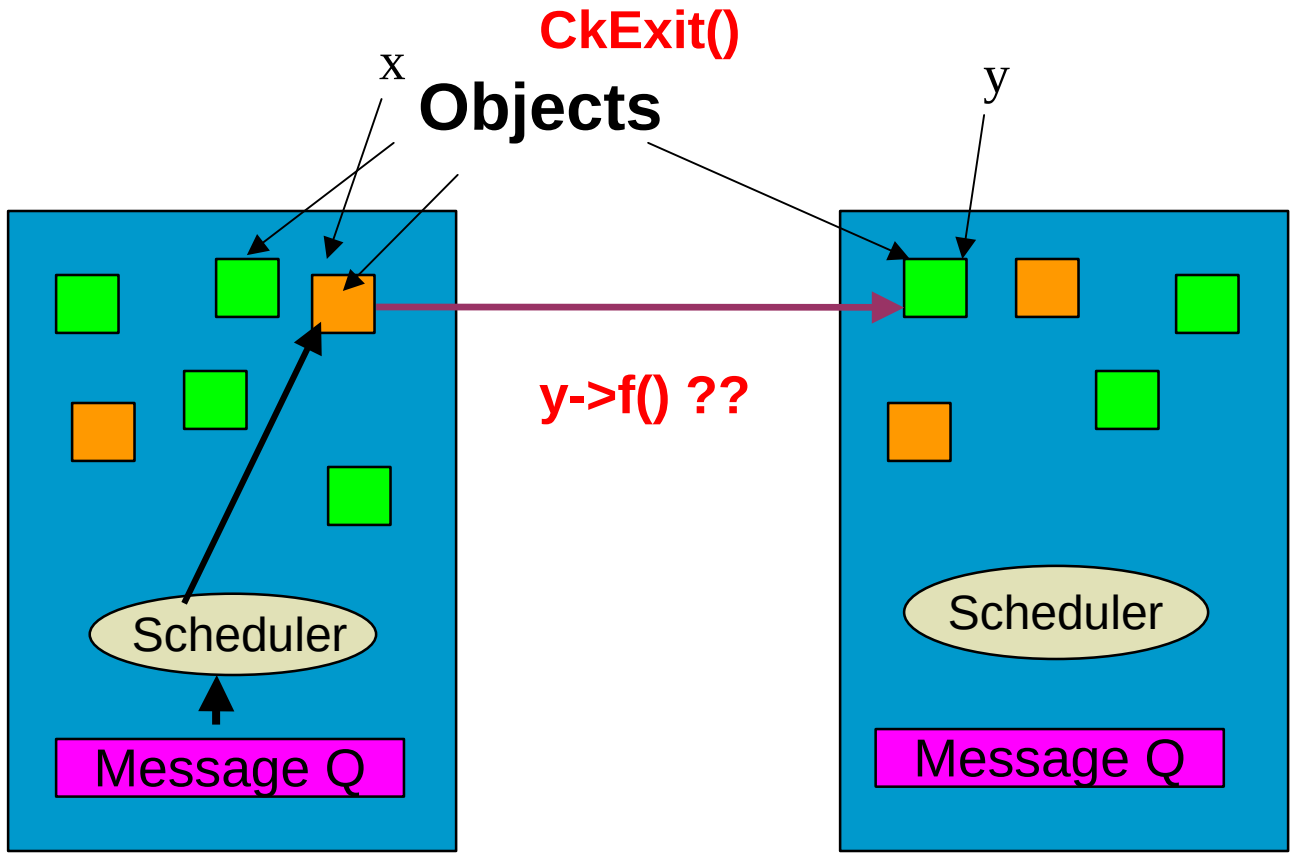

Other Charm++ Characteristics

- Methods execute one at a time
- No need for locks
- Expressing flow of control may be difficult

Outline

- Introduction
- Charm++ features
  - Chares and Chare Arrays
  - Parameter Marshalling
- Structured Dagger Construct
- Adaptive MPI
- Tools
  - Parallel Debugger
  - Projections
- Load Balancing
- LiveViz
- Conclusion

Chares – Concurrent Objects

- Can be dynamically created on any available processor
- Can be accessed from remote processors
- Send messages to each other asynchronously
- Contain “entry methods”

“Hello World”

```
// hello.ci

mainmodule hello
  mainchare myma
    entry mymain
  };
};
```

```
// hello.C file

#include “hello.decl.h”
class mymain : public Chare {
public:
     mymain(CkArgMsg *m)
     {
   ckout <<“Hello World”
              <<endl;
   CkExit();
     }

nclude “hello.def.h”
```

Generates:

hello.decl.h

hello.def.h

Compile and run the program

Compiling

- charmc <options> <source file>
- -o, -g, -language, -module, -tracemode

pgm: pgm.ci pgm.h pgm.C

Example Nodelist File:

group main ++shell ssh
host Host1
host Host2

rm++

To run a CH pgm" on four processors, type:

charmrun pgm +p4 <params>

Nodelist file (for network architecture)

- list of machines to run the program
- host <hostname> <qualifiers>

Charm++ solution: Proxy classes

- Proxy class generated for each chare class
  - For instance, CProxy_Y is the proxy class generated for chare class Y.
  - Proxy objects know where the real object is
  - Methods invoked on this object simply put the data in an “envelope” and send it out to the destination
- Given a proxy p, you can invoke methods
  - p.method(msg);

Chare Arrays

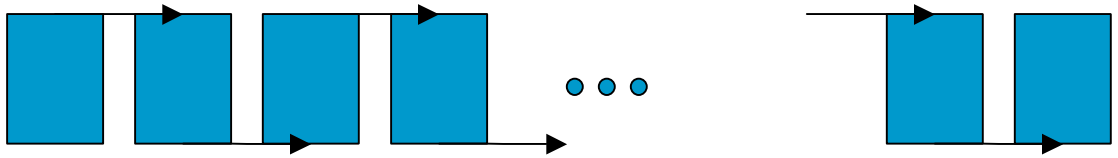

- Array of Objects of the same kind
- Each one communicates with the next one
- Individual chares – cumbersome and not practical

Chare Array:

- with a single global name for the collection
- each member addressed by an index
- mapping of element objects to processors handled by the system

Chare Arrays

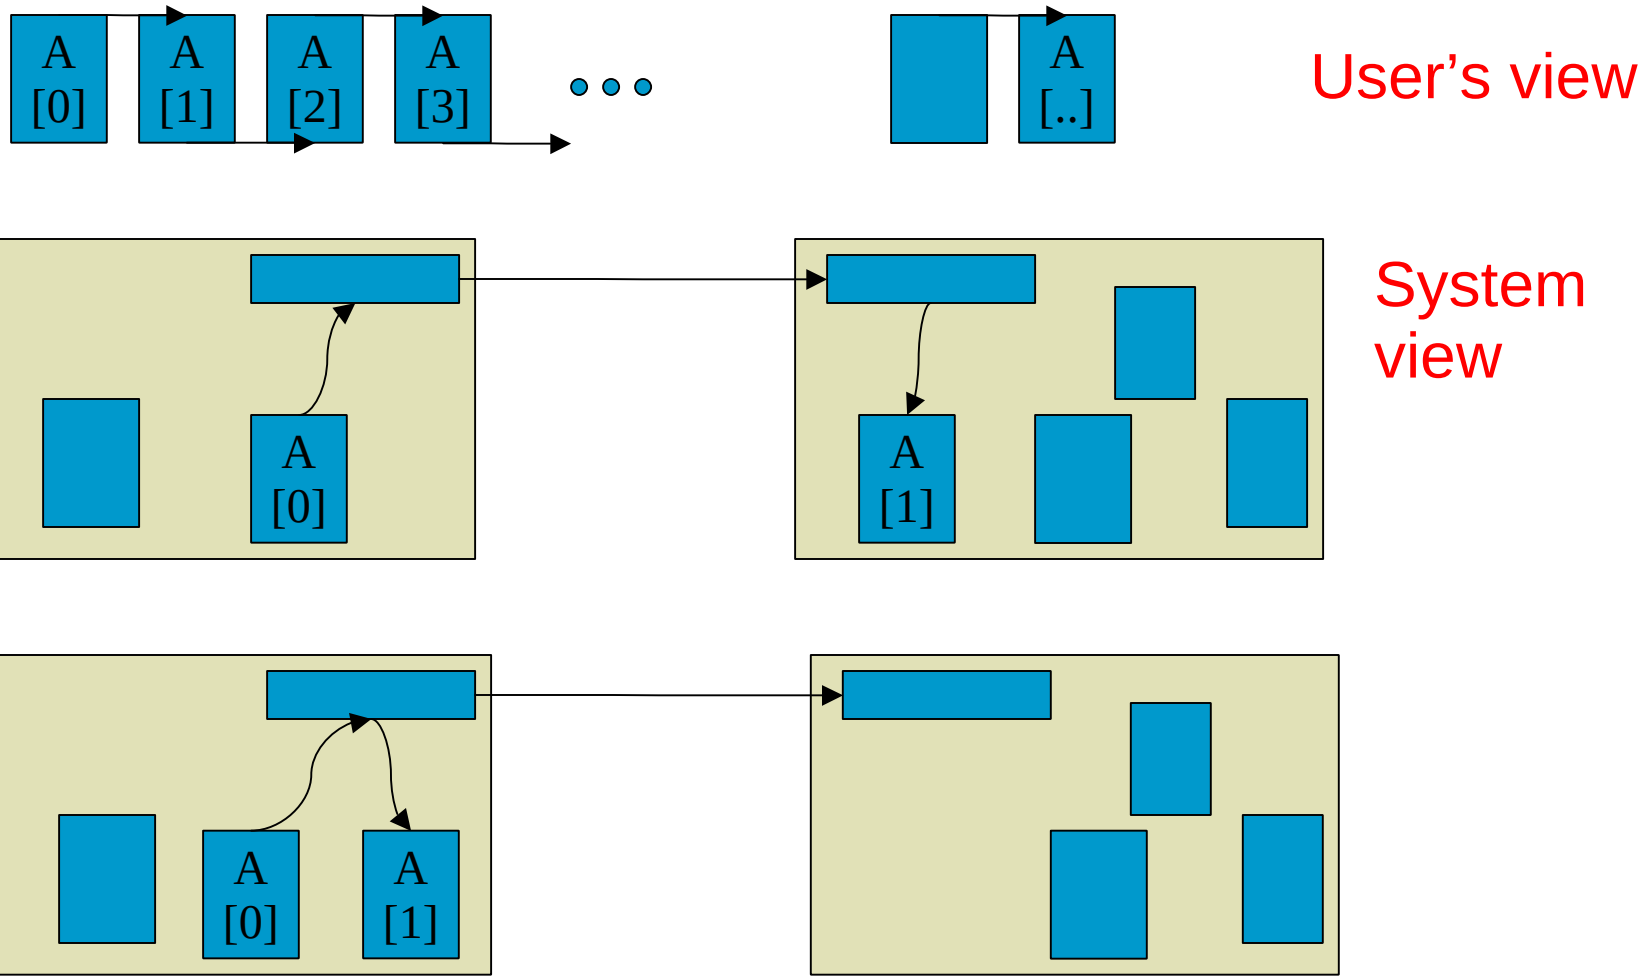

Array Hello

```
mainmodule m {
  readonly CProxy_mymain
mainProxy;
  readonly int nElements;
  mainchare mymain { …. }
  array [1D] Hello {
    entry Hello(void);
    entry void sayHi(int HiNo);
  };
};
```

```
class Hello : public CBase_Hello
{
  public:
    Hello(CkMigrateMessage *m){}
    Hello();
```

```
Class mymain : public Chare
{
    mymain() {
        nElements=4;
        mainProxy = thisProxy;
        CProxy_Hello p =
              CProxy_Hello::ckNew(nElements);
        //Have element 0 say "hi"
        p[0].sayHi(12345);
    }
}
```

In mymain:: mymain()

Array Hello

```
void Hello::sayHi(int hiNo)
{
    ckout <<  hiNo <<"from element" << thisIndex
                                    << endl;
    if (thisIndex < nElements-1)
     //Pass the hello on:
      thisProxy[thisIndex+1].sayHi(hiNo+1);
    else
      //We've been around once-- we're done.
      mainProxy.done();
}
```

Element index

Array Proxy

Read-only

```
void mymain::done(void){
     CkExit();
}
```

Sorting numbers

- Sort n integers in increasing order.
- Create **n** chares, each keeping one number.
- In every **odd iteration** chares numbered 2i swaps with chare 2i+1 if required.
- In every **even iteration** chares 2i swaps with chare 2i-1 if required.
- After each iteration all chares report to the mainchare. After everybody reports mainchares signals next iteration. Sorting completes in **n** iterations.

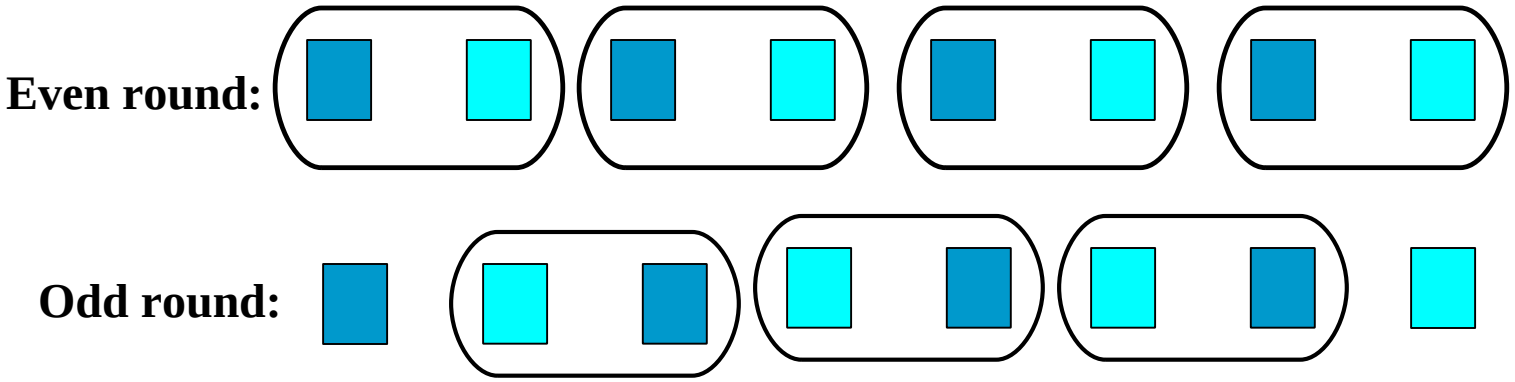

Array Sort

sort.ci

```
mainmodule sort{
    readonly CProxy_myM
    readonly int nEleme

    mainchare myMain {
        entry myMain(Ck
        entry void swap
  };
  array [1D] sort{
     entry sort(void);
int
};
```

sort.h

```
class sort : public CBase_sort{
        private:
                int myValue;
        public:
                sort() ;
                sort(CkMigrateMessage *m);
                void setValue(int number);
                void swap(int round_no);
                e(int from_index,
                nt value);
```

myMain::myMain()

```
        swapcount=0;
        roundsDone=0;
        mainProxy = thishandle;
        CProxy_sort arr =
                CProxy_sort::ckNew(nElements);
        for(int i=0;i<nElements;i++)
                arr[i].setValue(rand());
        arr.swap(0);
```

Array Sort (continued ...)

```
void sort::swap(int roundno)
{
   bool sendright=false;
   if (roundno%2==0 && thisIndex%2==0|| roundno%2==1 &&
thisIndex%2==1)
      sendright=true;   //sendright is true if I have to send to right

   if((sendright && thisIndex=
thisIndex==0))
      mainProxy.swapdone();
   else{
       if(s
        thi
      else
        thi
   }
}
```

```
void sort::swapReceive(int from_index,
int value)
{
                                  =thisIndex-1 &&

                                  value;
                                  =thisIndex+1 &&

                                  value;
```

```
void myMain::swapdone(void) {
    if (++swapcount==nElements) {
        swapcount=0;
        roundsDone++;
        if (roundsDone==nElements)
            CkExit();
        else
            arr.swap(roundsDone);
    }
}
```

Error!

Remember :

- ✓ Message passing is **asynchronous**.
- ✓ Messages can be delivered **out of order**.

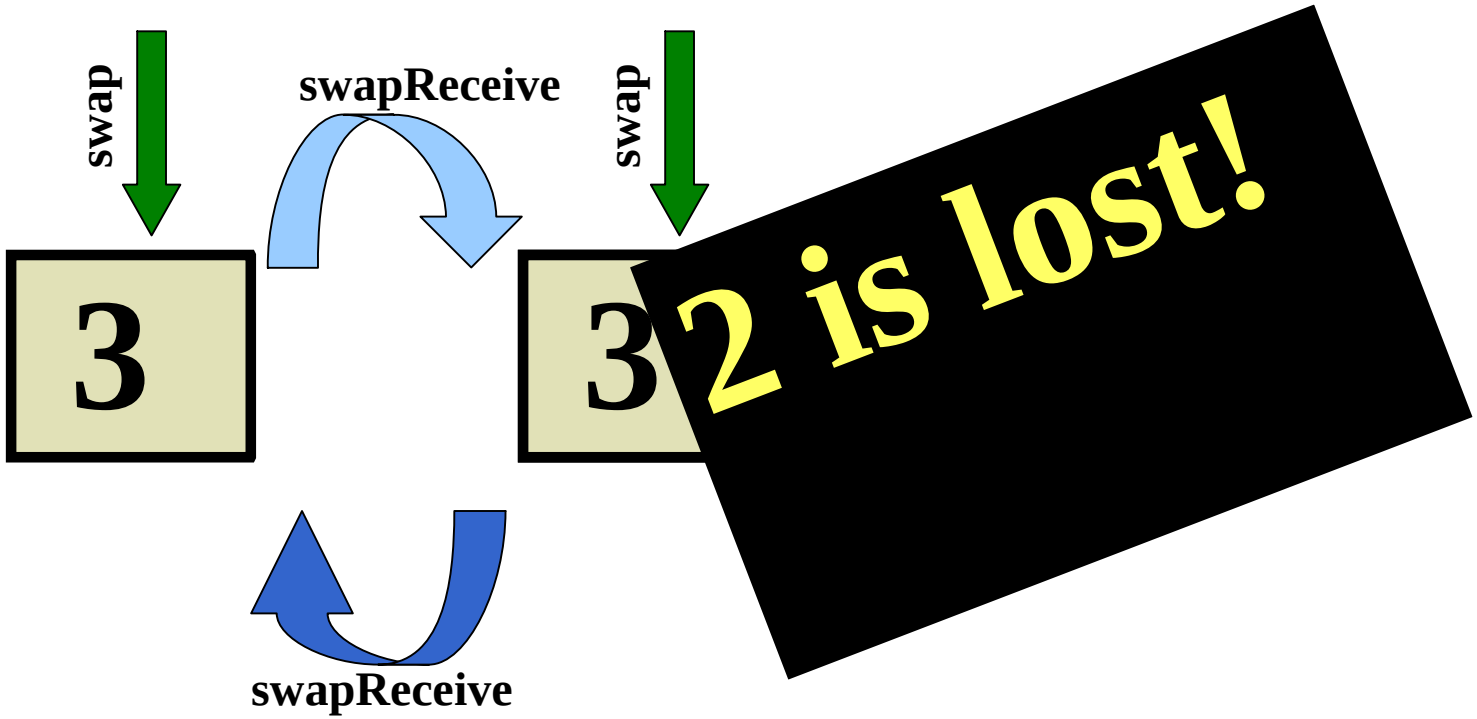

Array Sort (correct)

```
void sort::swap(int roundno)
{
   bool sendright=false;
   if (roundno%2==0 && thisIndex%2==0|| roundno%2==1 &&
thisIndex%2==
      sendrigh
right

   if
thisIn

   } (

   }
}
```

```
void sort::swapReceive(int from_index, int value) {

   if (from_index==thisIndex-1) {
      if (value>myValue) {
                                          swapReceive(thisIndex,

                                          swapReceive(thisIndex,
}
```

```
void myMain::swapdone(void) {
   if (++swapcount==nElements) {
      swapcount=0;
      roundsDone++;
      if (roundsDone==nElements)
         CkExit();
      else
         arr.swap(roundsDone);
   }
}
```

Array Sort II: A Different Approach

- Do not have the chares do work unless it is needed
- All processing is message driven (the result of receiving a message, No for loops)
- Do not continue the sort unless there is work to be done…

Quiescence Detection

"…the state in which no processor is executing an entry point, and no messages are awaiting processing…" --- Charm Manual

- Uses a Callback Function (more on Callback Functions later)
- For now: When Quiescence is detected, the Callback Function will be called and perform the desired task (in this case, print the sorted array and call CkExit() to end the program)

Array Sort II (cont.)

- Membe
  - Rec
  - denyS
    - Rec
  - accep
    - Rec
      acc
  - check
    - Use
      while the chare was already busy taking care of another swap

- When the array is sorted, all ‘requestSwap()’s will be answered with ‘denySwap()’s
- The remaining messages drain from the queues
- Quiescence occurs

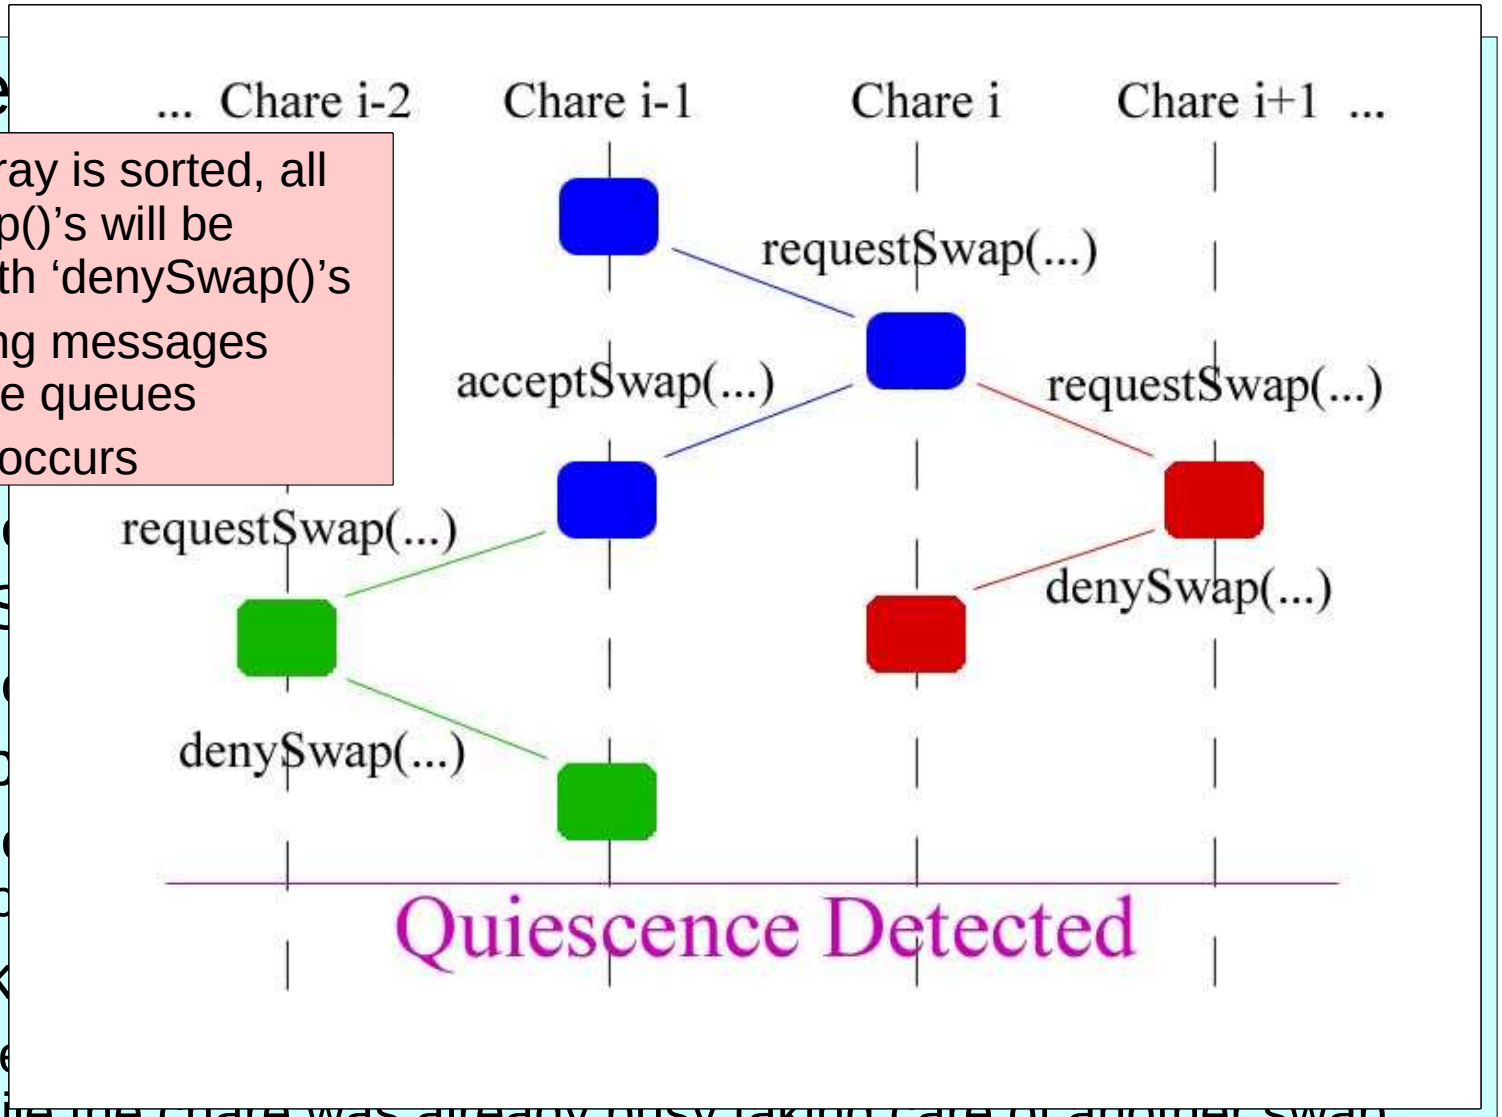

```
void Bubble::checkForPending() {

    // Check to see if there is a pending initiate swap
    if (pendingInitIndex > 0) {
        // There is a pending initiate swap... resend request to self
        // (Note: initSwapSequenceWith() does
        // not call this function so it is safe to do a standard call to it from here.)
        initSwapSequenceWith(pendingInitIndex);
        pendingInitIndex = -1;
    }

    // Check to see if there is a pending request for a swap
    if (pendingRequestIndex >= 0) {

        // There is a pending request for a swap... resend request to self
        // (Note: This function clears pendingRequestIndex and pendingRequestValue
        // before making the call to requestSwap() so when requestSwap() calls this
        // function again at the end, execution will not enter this if statement
        // a second time which means there will not be an infinite loop of calls back
        // and forth between the two functions as one might think at first glance.

        // Also note that isSwappingWith will be -1 if this function is called.)
        int tempIndex = pendingRequestIndex;
        int tempValue = pendingRequestValue;
        pendingRequestIndex = -1;
        pendingRequestValue = -1;
        thisProxy[thisIndex].requestSwap(tempIndex, tempValue);
    }
}
```

```
void Bubble::denySwap(int index) {

    // Finished with the swap so exit the swap sequence
    isSwappingWith = -1;

    // Check to see if there are any pending items
    if (isSwappingWith < 0)
        checkForPending();
}
```

Example: 5-Point 2-D Stencil

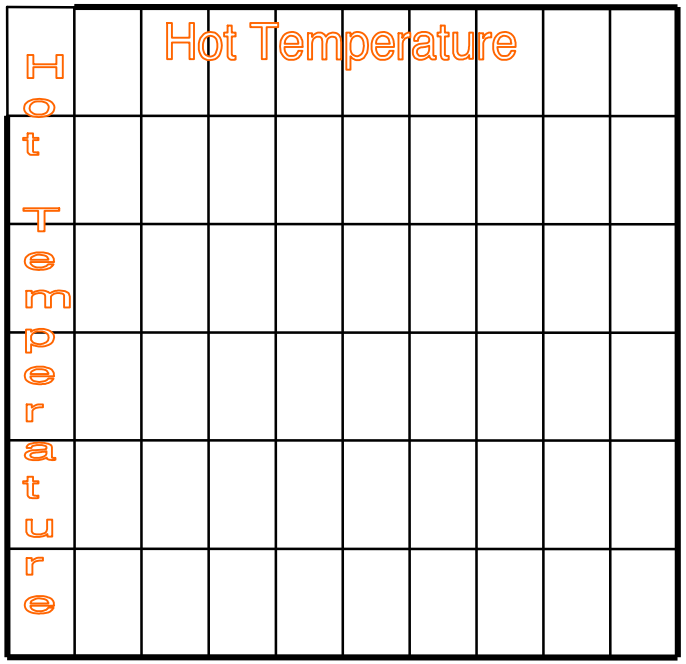

Hot temperature on two sides will slowly spread across the entire grid.

Example: 5-Point 2-D Stencil

- Input: 2D array of values with boundary conditions
- In each iteration, each array element is computed as the average of itself and its neighbors(average on 5 points)
- Iterations are repeated till some threshold difference value is reached

Parallel Solution!

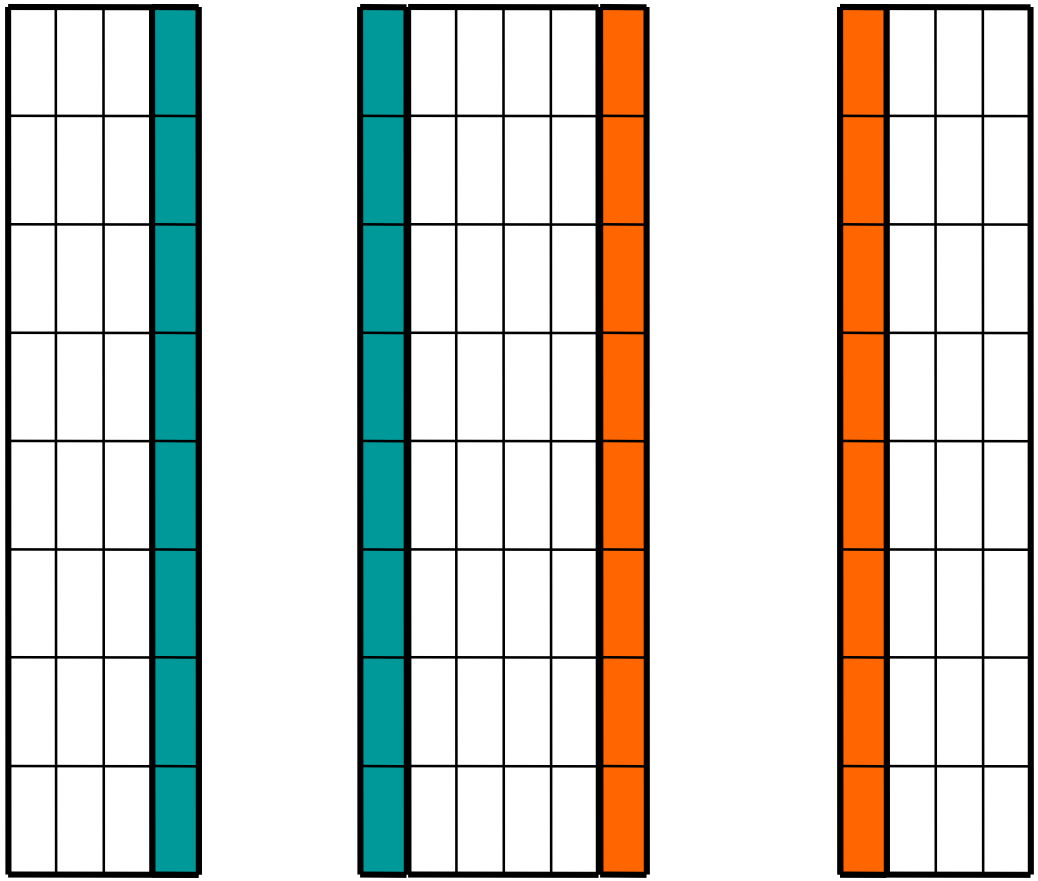

Parallel Solution!

- Slice up the 2D array into sets of columns
- Chare = computations in one set
- At the end of each iteration
  - Chares exchange boundaries
  - Determine maximum change in computation
- Output result at each step or when threshold is reached

Arrays as Parameters

- Array cannot be passed as pointer
- specify the length of the array in the interface file
  - entry void bar(int n,double arr[n])
  - n is size of arr[]

Stencil Code

```
void Ar1::doWork(int sendersID, int n, double arr[])
{
    maxChange = 0.0;
    if (sendersID == thisIndex-1)
    {  leftmsg = 1; }
    //set boolean to indicate we received the left message
    else if (sendersID == thisIndex+1)
    {  rightmsg = 1; }
    //set boolean to indicate we received the right message
    // Rest of the code on a following slide
    …
}
```

Reduction

- Apply a single operation (add, max, min, ...) to data items scattered across many processors
- Collect the result in one place
- Reduce *x* across all elements
  - **contribute(sizeof(***x***), &***x***, CkReduction::***sum_int***);**
- Must create and register a callback function that will receive the final value, in main chare

Types of Reductions

- Predefined Reductions – A number of reductions are predefined, including ones that
  - Sum values or arrays
  - Calculate the product of values or arrays
  - Calculate the maximum contributed value
  - Calculate the minimum contributed value
  - Calculate the logical and of integer values
  - Calculate the logical or of contributed integer values
  - Form a set of all contributed values
  - Concatenate bytes of all contributed values
- Plus, you can create your own

Code (continued …)

```
void Ar1::doWork(int sendersID, int n, double arr[n])
{
   //Code on previous slide
   …
   if (((rightmsg == 1) && (leftmsg == 1)) || ((thisIndex == 0)
   &&
   (rightmsg == 1)) || ((thisIndex ==K-1) && (leftmsg == 1)))
   {
     // Both messages have been received and we can now
         compute the new values of the matrix
        …
     // Use a reduction to find determine if all of the maximum
        errors on each processor had a maximum change that
        is below our threshold value.

   contribute(sizeof(double), &maxChange,
                                        CkReduction::max_double);
   }
}
```

Callbacks

- A generic way to transfer control to a chare after a library(such as reduction) has finished.
- After finishing a reduction, the results have to be passed to some chare's entry method.
- To do this, create an object of type *CkCallback* with chare's ID & entry method index
- Different types of callbacks
- One commonly used type:

CkCallback cb(<chare's entry method>,<chare's proxy>);

Outline

- Introduction
- Charm++ features
  - Chares and Chare Arrays
  - Parameter Marshalling
- Structured Dagger Construct
- Adaptive MPI
- Tools
  - Parallel Debugger
  - Projections
- Load Balancing
- LiveViz
- Conclusion

Structured Dagger

- Motivation:
  - Keeping flags & buffering manually can complicate code in charm++ model.
  - Considerable overhead in the form of thread creation and synchronization

Advantages

- Reduce the complexity of program development
  - Facilitate a clear expression of flow of control
- Take advantage of adaptive message-driven execution
  - Without adding significant overhead

What is it?

- A coordination language built on top of Charm++
  - Structured notation for specifying intra-process control dependences in message-driven programs
- Allows easy expression of dependences among messages, computations and also among computations within the same object using various structured constructs

Structured Dagger Constructs

To Be Covered in Advanced Charm++ Session

- *atomic* {code}
- *overlap* {code}
- *when* <entrylist> {code}
- if/else/for/while
- foreach

Stencil Example Using Structured Dagger

```
stencil.ci
 array[1D] Ar1 {
…
entry void GetMessages () {
    when rightmsgEntry(), leftmsgEntry() {
        atomic { CkPrintf("Got both left and right messages \n");
                 doWork(right, left); }
     }
};

entry void rightmsgEntry();
entry void leftmsgEntry();
…
};
```

AMPI = Adaptive MPI

- Motivation:
  - Typical MPI implementations are not suitable for the new generation parallel applications
    - Dynamically varying: load shifting, adaptive refinement
  - Some legacy codes in MPI can be easily ported and run fast in current new machines
  - Facilitate those who are familiar with MPI

What is it?

- An MPI implementation built on Charm++ (MPI with virtualization)
- To provide benefits of Charm++ Runtime System to standard MPI programs
  - Load Balancing, Checkpointing, Adaptability to dynamic number of physical processors

Sample AMPI Program
Also a valid MPI Program

```
#include <stdio.h>
#include "mpi.h"

int main(int argc, char** argv){
  int ierr, rank, np, myval=0;
  MPI_Status status;

  MPI_Init(&argc, &argv);
  ierr = MPI_Comm_rank(MPI_COMM_WORLD, &rank);
  ierr = MPI_Comm_size(MPI_COMM_WORLD, &np);

  if(rank < np-1) MPI_Send(&myval, 1, MPI_INT, rank+1,1,MPI_COMM_WORLD);
  if(rank > 0) MPI_Recv(&myval,1, MPI_INT, rank-1,1,MPI_COMM_WORLD, &status);

  printf("rank %d completed\n", rank);
  ierr = MPI_Finalize();
}
```

AMPI Compilation

Compile:

charmc sample.c -language ampi -o sample

Run:

charmrun ./sample +p16 +vp 128 [args]

Instead of Traditional MPI equivalent:

mpirun ./sample -np 128 [args]

Comparison to Native MPI

- AMPI Performance
  - Similar to Native MPI
  - Not utilizing any other features of AMPI(load balancing, etc.)
- AMPI Flexibility
  - **AMPI runs on any # of Physical Processors (eg 19, 33, 105). Native MPI needs cube #.**

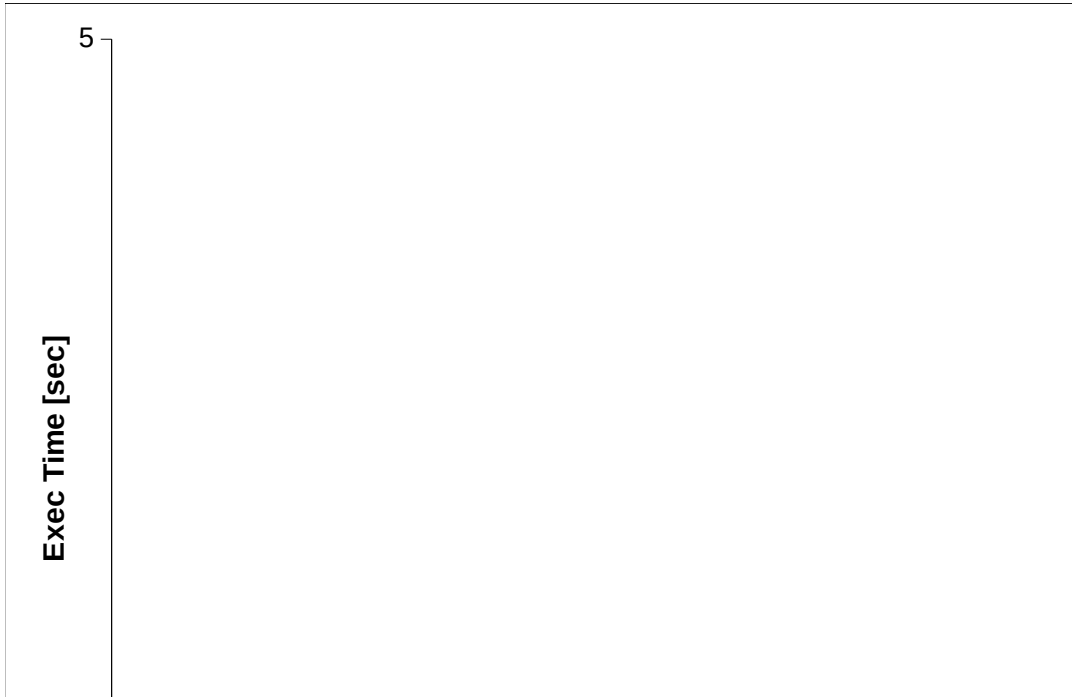

Problem setup: 3D stencil calculation of size 240^3 run on Lemieux.

Current AMPI Capabilities

- Automatic checkpoint/restart mechanism
  - Robust implementation available
- Load Balancing and “process” Migration
- MPI 1.1 compliant, Most of MPI 2 implemented
- Interoperability
  - With Frameworks
  - With *Charm++*
- Performance visualization

More on the next session!

Outline

- Introduction
- Charm++ features
  - Chares and Chare Arrays
  - Parameter Marshalling
- Structured Dagger Construct
- Adaptive MPI
- Tools
  - Parallel Debugger
  - Projections
- Load Balancing
- LiveViz
- Conclusion

Parallel debugging support

- Parallel debugger (*charmdebug*)
- Allows programmer to view the changing state of the parallel program
- Java GUI client

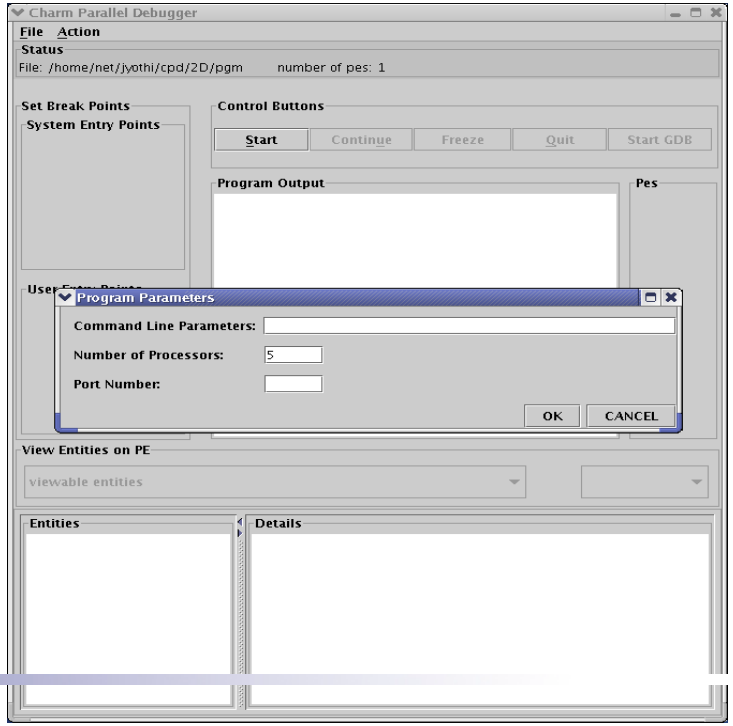

Debugger features

- Provides a means to easily access and view the **major programmer visible entities,** including objects and messages in queues, during program execution
- Provides an interface to **set and remove breakpoints** on remote entry points, which capture the major programmer-visible control flows

Debugger features (contd.)

- Provides the ability to freeze and unfreeze the execution of selected processors of the parallel program, which allows a **consistent snapshot**
- Provides a way to **attach** a sequential debugger (like **GDB**) to a specific subset of processes of the parallel program during execution, which keeps a manageable number of sequential debugger windows open

Alternative debugging support

- Uses gdb for debugging
  - Runs each node under gdb in an xterm window, prompting the user to begin execution
- Charm program has to be compiled using ‘-g’ and run with ‘++debug’ as a command-line option.

Projections: Quick Introduction

(More detailed in later session!)

- Projections is a tool used to **analyze** the **performance** of your application
- The tracemode option is used when you build your application to enable tracing
- You get one log file per processor, plus a separate file with global information
- These files are read by Projections so you can use the Projections views to analyze performance

Screen shots – Load imbalance

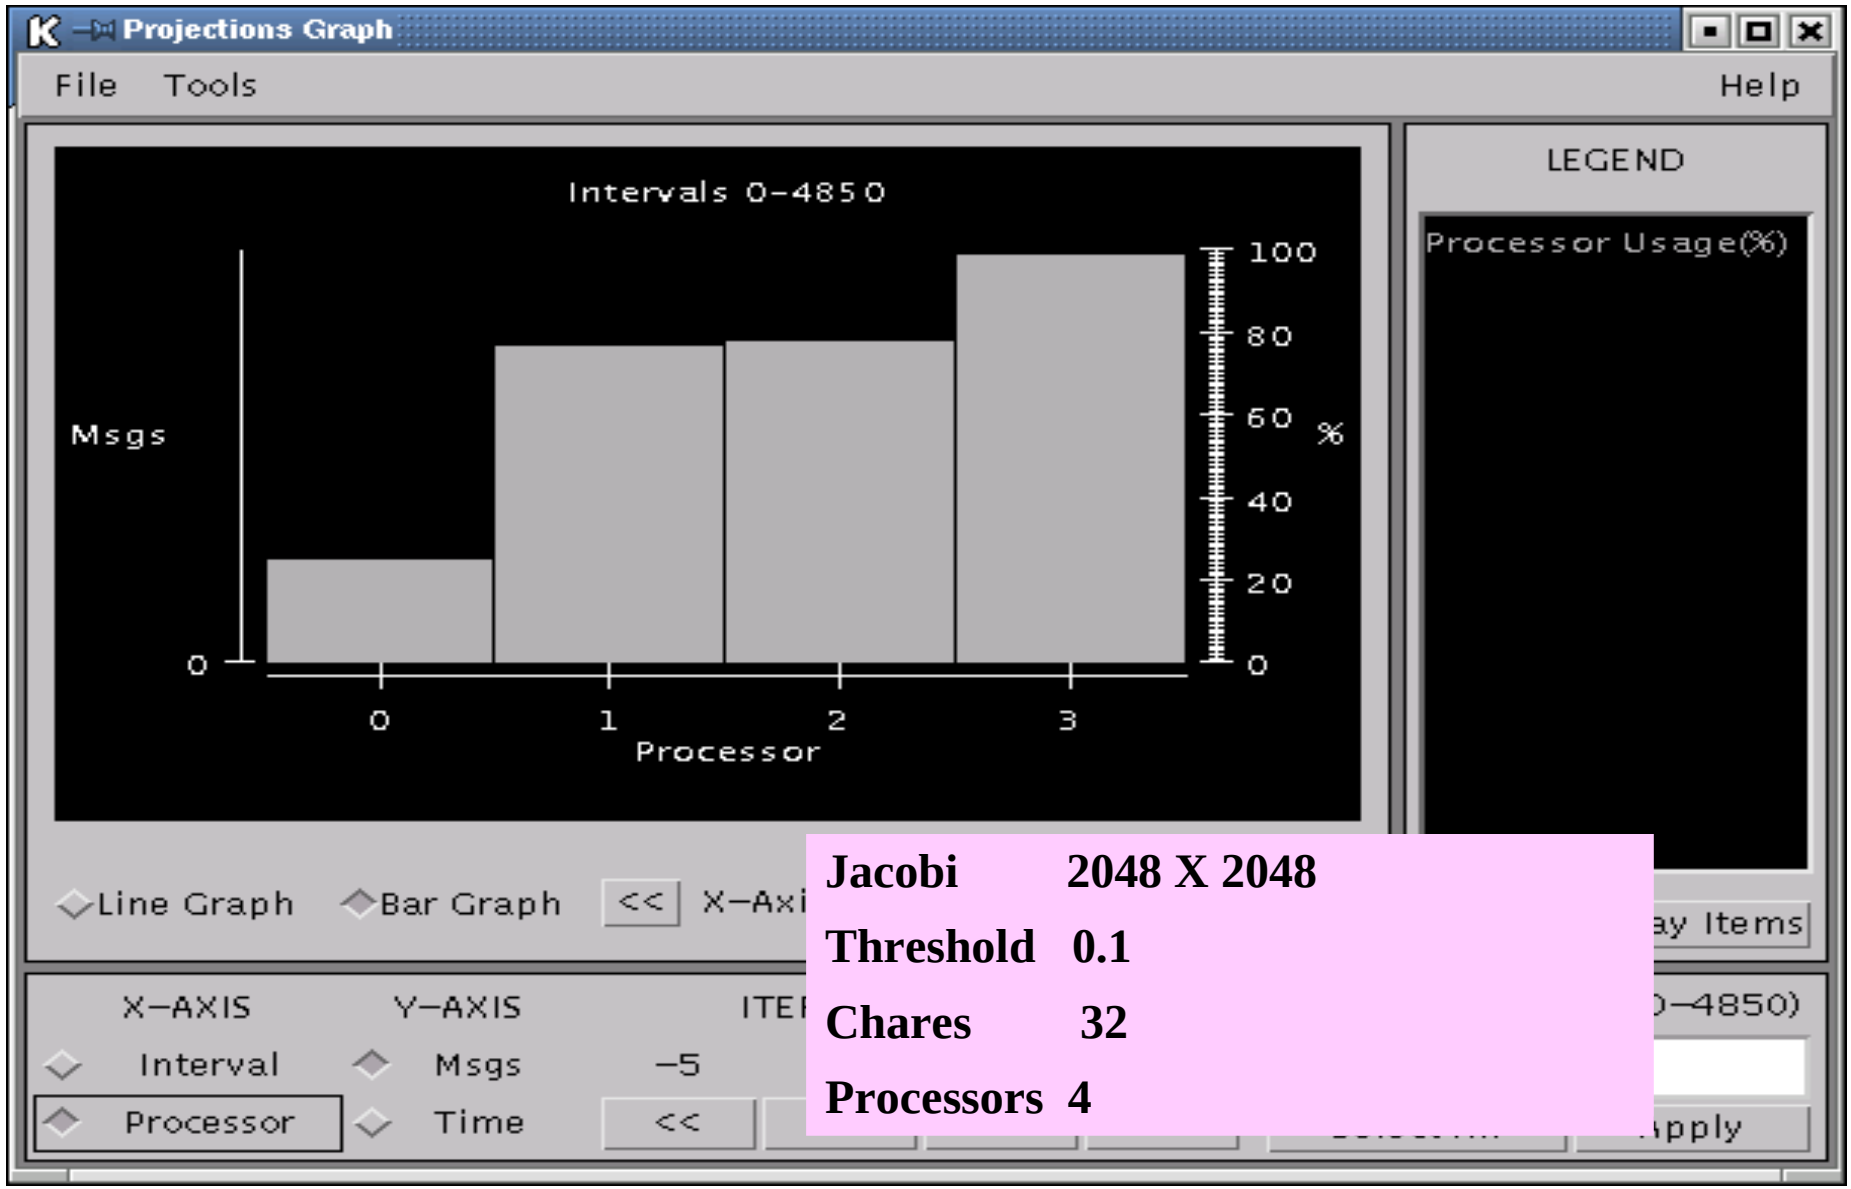

Timelines – load imbalance

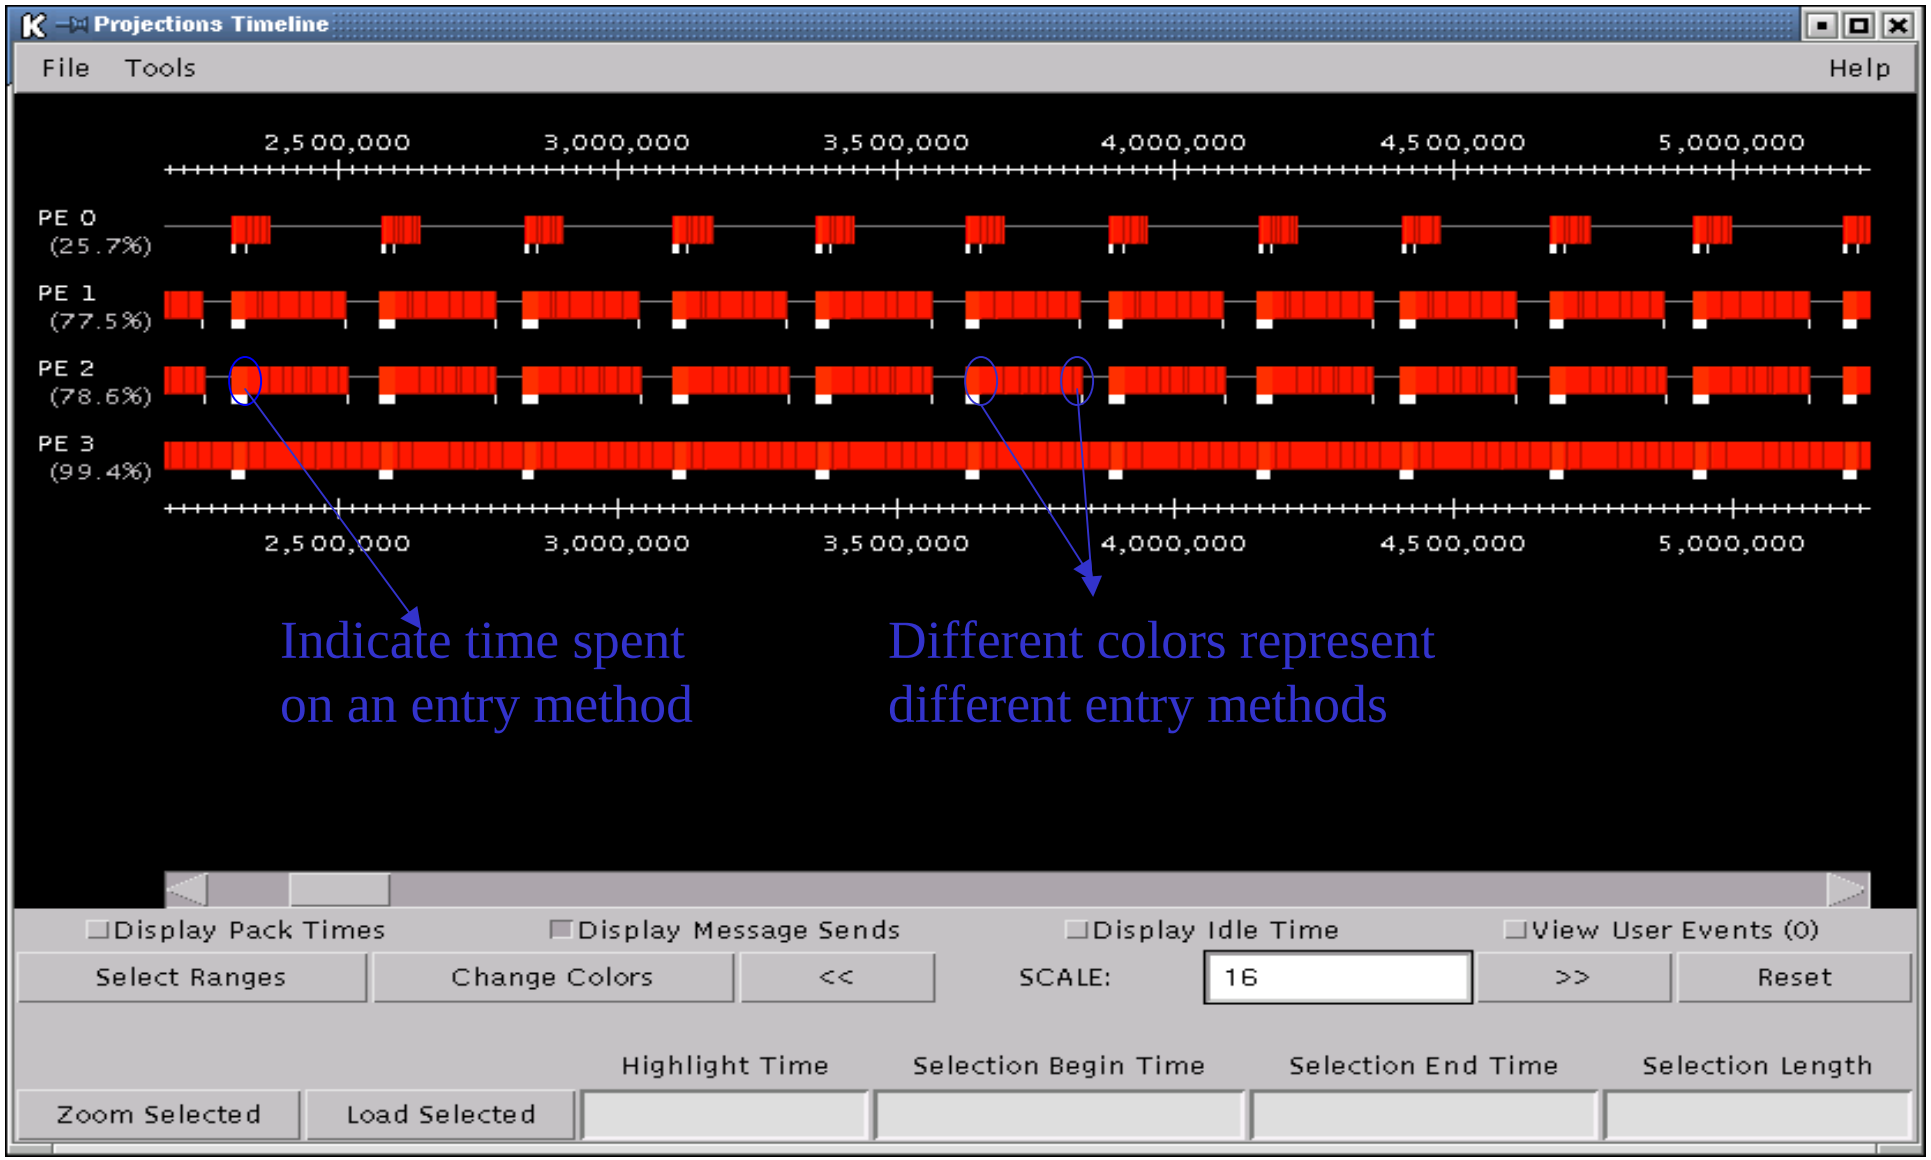

Outline

- Introduction
- Charm++ features
  - Chares and Chare Arrays
  - Parameter Marshalling
- Structured Dagger Construct
- Adaptive MPI
- Tools
  - Parallel Debugger
  - Projections
- Load Balancing
- LiveViz
- Conclusion

Load Balancing

- Goal: **higher processor utilization**
- Object migration allows us to move the work load among processors easily
- Measurement-based Load Balancing
- Two approaches to distributing work:
  - Centralized
  - Distributed
- Principle of Persistence

Migration

- Array objects can **migrate** from one processor to another
- Migration creates a new object on the destination processor while destroying the original
- Need a way of **packing** an object into a message, then **unpacking** it on the receiving processor

PUP

- **PUP is a framework for packing and unpacking migratable objects into messages**
- **To migrate, must implement pack/unpack or *pup* method**
- **Pup method combines 3 functions**
  - **Data structure traversal : compute message size, in bytes**
  - **Pack : write object into message**
  - **Unpack : read object out of message**

Writing a PUP Method

```
Class ShowPup {
   double a;      int x;
   char y;    unsigned long z;
   float q[3];    int *r; // heap allocated memory
public:
    void pup(PUP::er &p) {
    if (p.isUnpacking())
        r = new int[ARRAY_SIZE];
    p | a; p |x; p|y   // you can use | operator
    p(z); p(q, 3)      // or ()
    p(r,ARRAY_SIZE);
    }
};
```

The Principle of Persistence

- Big Idea: the past predicts the future
- Patterns of communication and computation remain nearly constant

- By measuring these patterns we can improve our load balancing techniques

Centralized Load Balancing

- Uses information about activity on all processors to make load balancing decisions
- Advantage: **Global information** gives higher quality balancing
- Disadvantage: Higher **communication costs and latency**
- Algorithms: Greedy, Refine, Recursive Bisection, Metis

Neighborhood Load Balancing

- Load balances among a small set of processors (the neighborhood)
- Advantage: Lower communication costs
- Disadvantage: Could leave a system which is poorly balanced globally

- Algorithms: NeighborLB, WorkstationLB

When to Re-balance Load?

Default: Load balancer will migrate when needed

- **Programmer Control: AtSync load balancing**

  AtSync method: enable load balancing at specific point
  - **Object ready to migrate**
  - **Re-balance if needed**
  - ***AtSync()* called when your chare is ready to be load balanced**
    - **load balancing may not start right away**
  - ***ResumeFromSync()* called when load balancing for this chare has finished**

Using a Load Balancer

- link a LB module
  - ***-module <strategy>***
  - RefineLB, NeighborLB, GreedyCommLB, others…
  - EveryLB will include all load balancing strategies
- compile time option (specify default balancer)
  - **-balancer RefineLB**
- runtime option
  - ***+balancer RefineLB***

Load Balancing in Jacobi2D

```
Main:
   Setup worker array, pass data to them

Workers:
   Start looping

   Send messages to all neighbors with ghost rows

   Wait for all neighbors to send ghost rows to me

   Once they arrive, do the regular Jacobi relaxation

   Calculate maximum error, do a reduction to compute
     global maximum error

   If timestep is a multiple of 64, load balance the
     computation. Then restart the loop.
```

Load Balancing in Jacobi2D (cont.)

```
worker::worker(void) {
     //Initialize other parameters
```

```
Void worker::doCompute(void){
  // do all the jacobi computation
  syncCount++;
  if(syncCount%64==0)
    AtSync();
  else
```

```
c
}
```

```
void worker::ResumeFromSync(void){
 contribute(1*sizeof(float),&errorMax,CkReduction::max_float);
}
```

Processor Utilization: After Load Balance

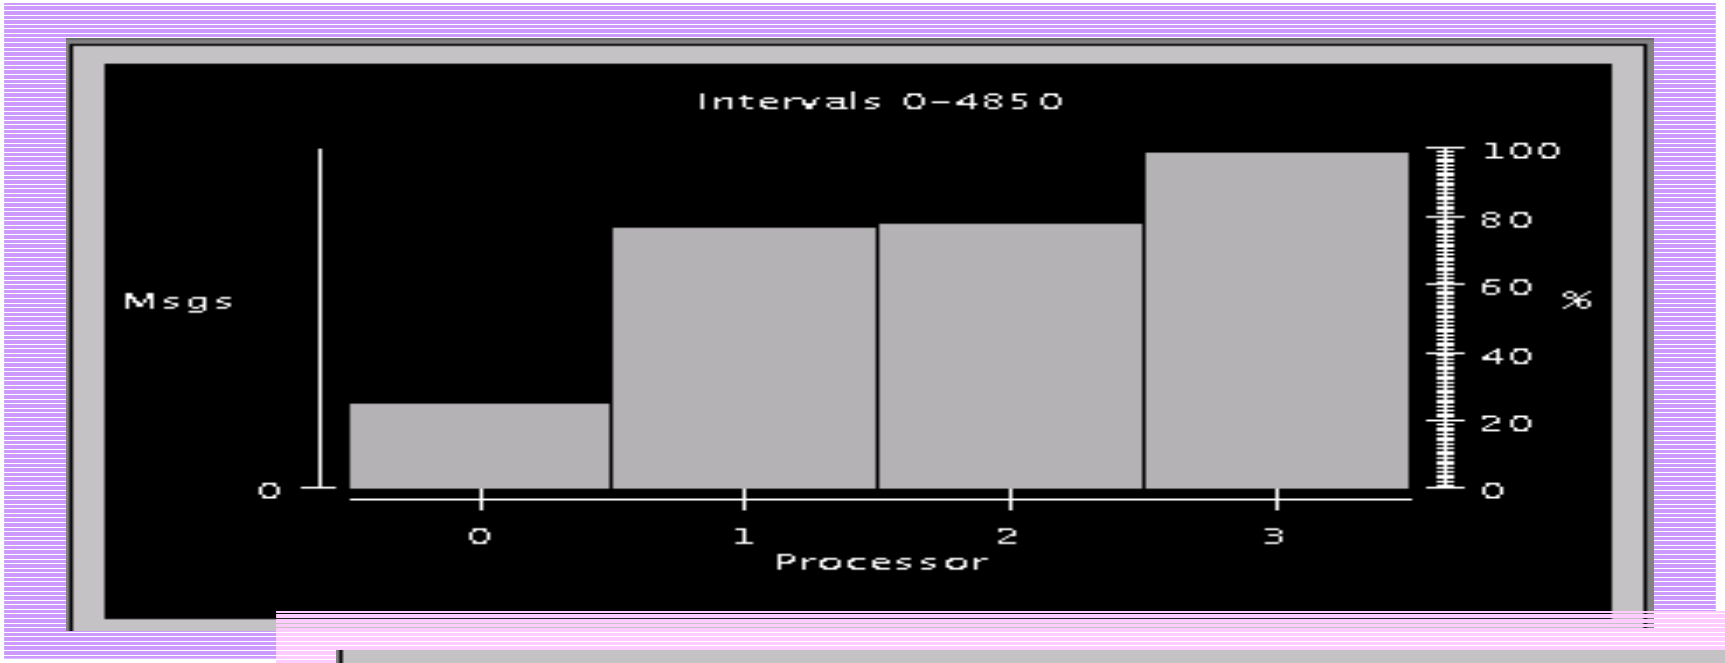

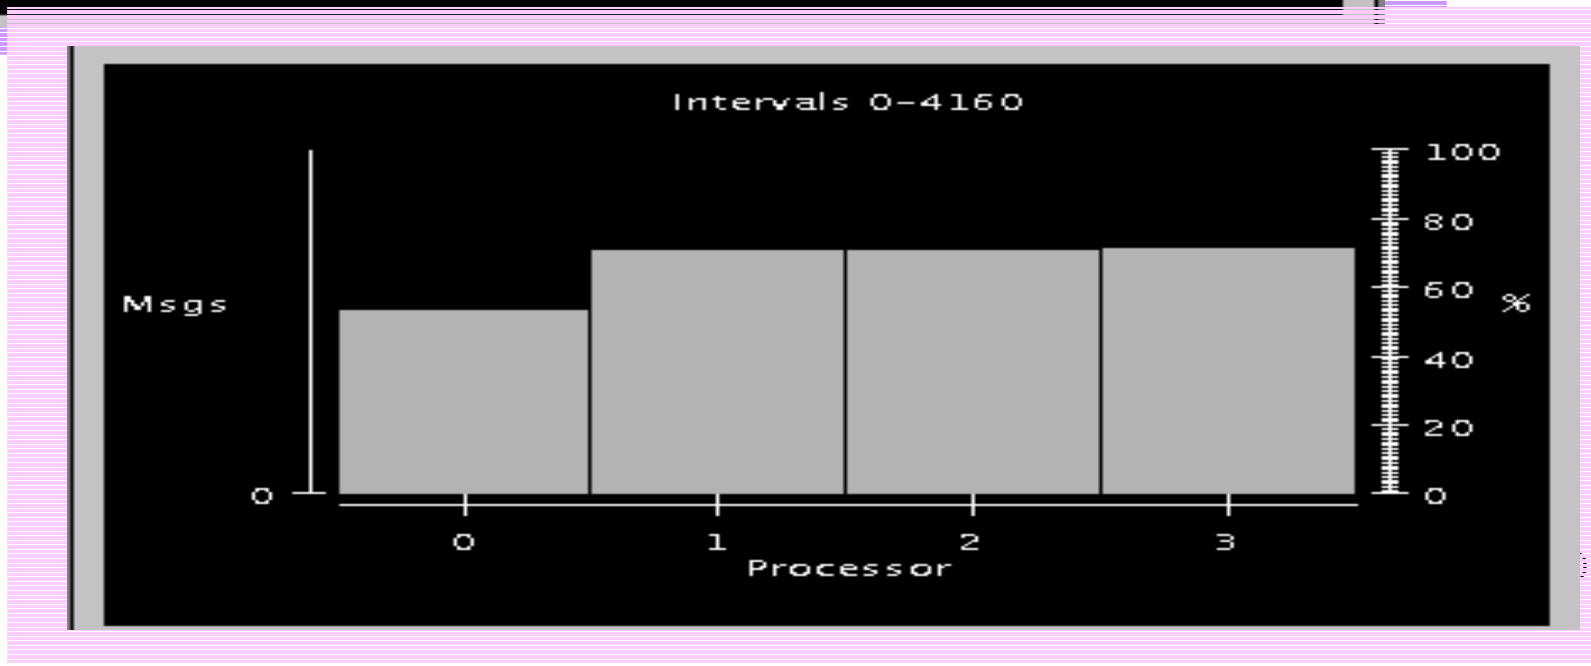

Timelines: Before and After Load Balancing

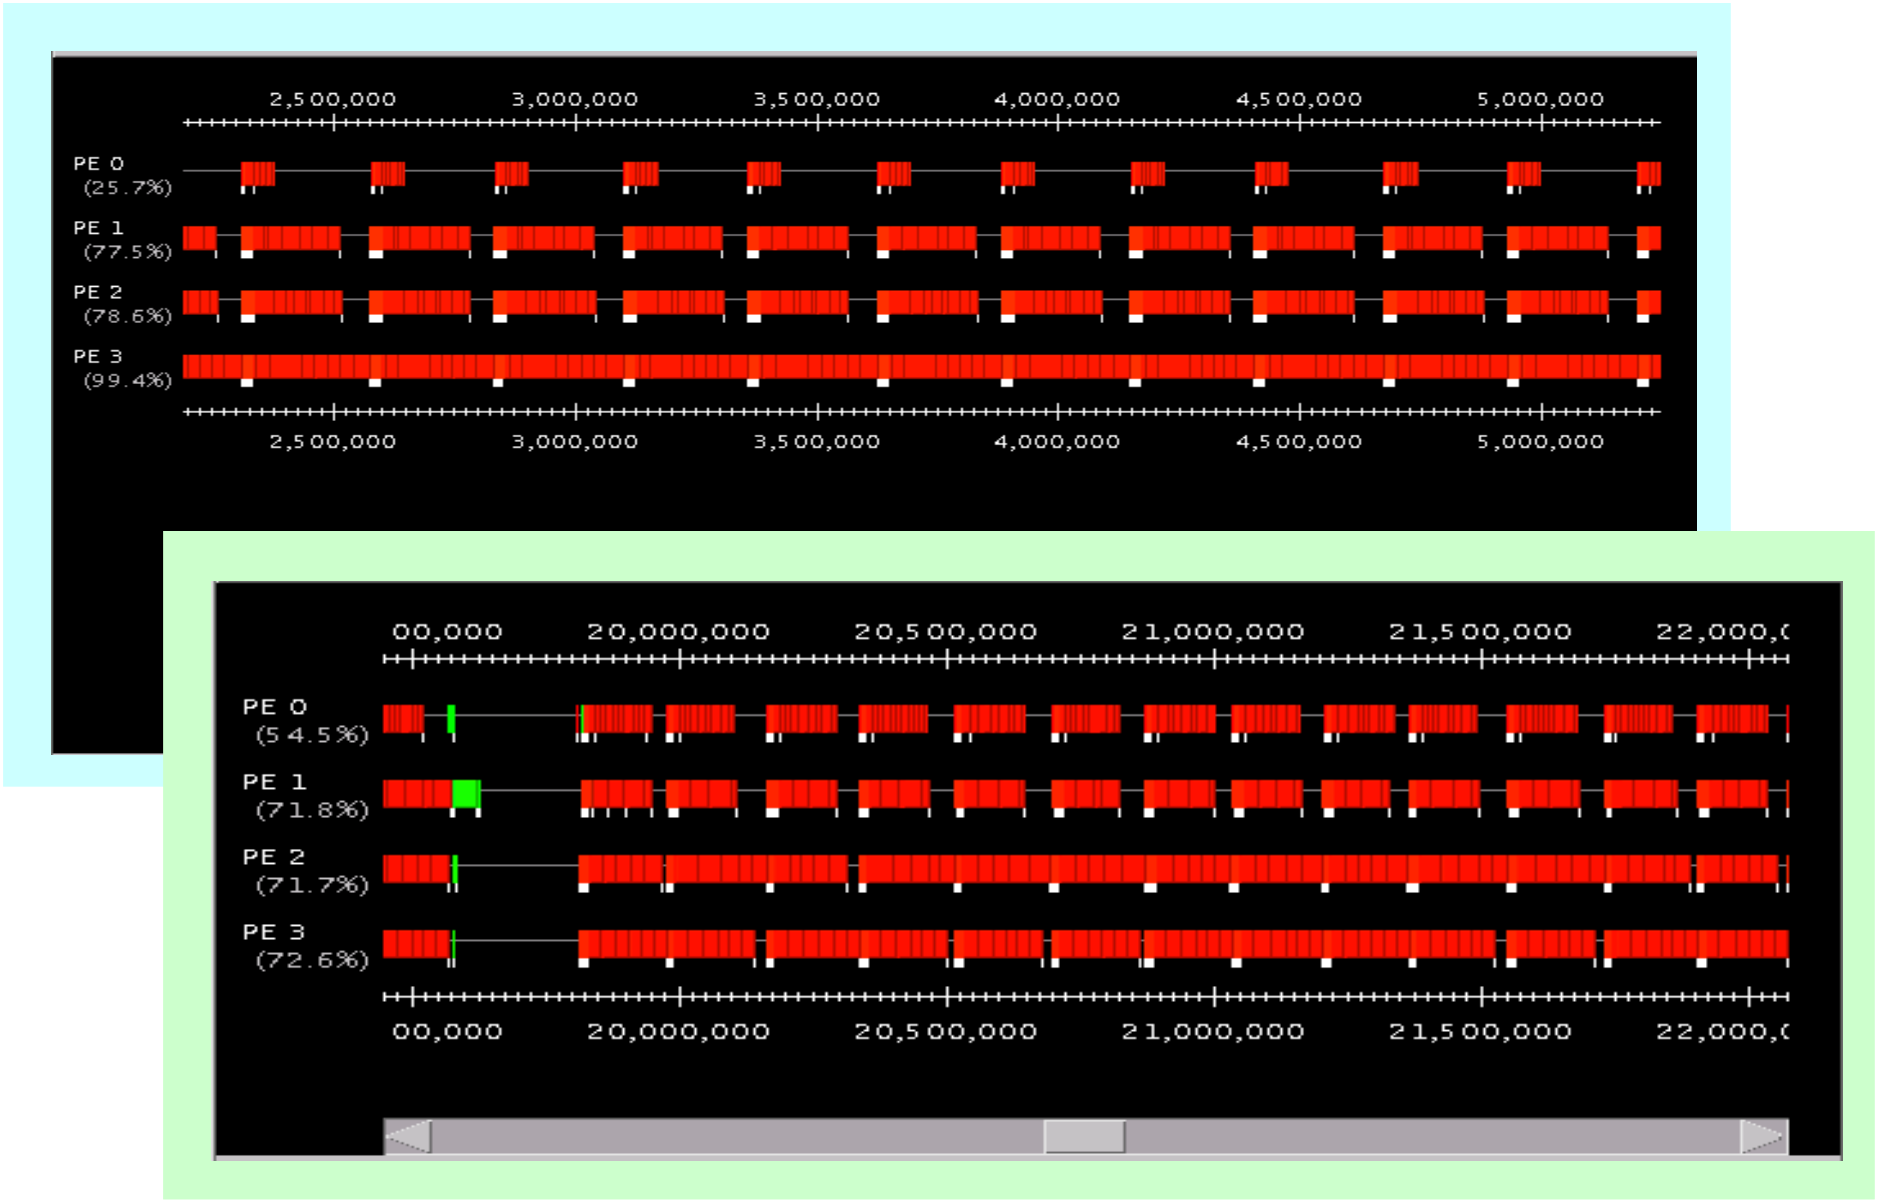

LiveViz – What is it?

- Charm++ library
- Visualization tool
- Inspect your program's current state
- Java client runs on any machine
- You code the image generation
- 2D and 3D modes

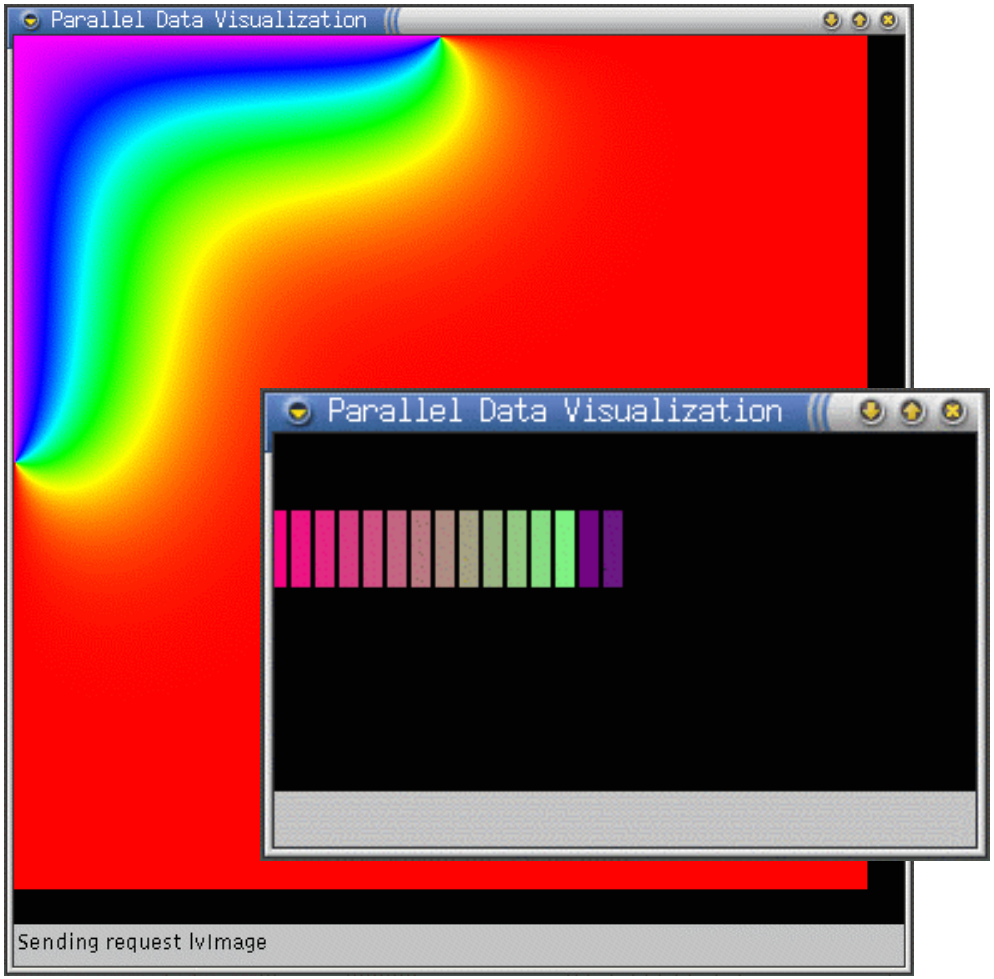

LiveViz – Monitoring Your Application

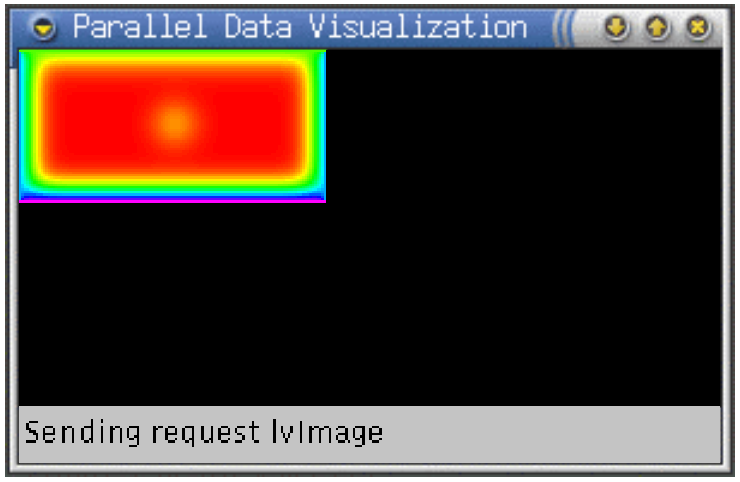

- LiveViz allows you to watch your application's progress
- Doesn't slow down computation when there is no client

LiveViz - Compilation

- LiveViz is part of the standard Charm++ distribution – when you build Charm++, you also get LiveViz

```
nilesh@skill:~/work/charm/src/libs/ck-libs/liveViz> make
../../../../bin/charmc  liveViz.ci
Warning: deprecated use of initcall. Use initnode or initproc instead.
../../../../bin/charmc  liveVizPoll.ci
../../../../bin/charmc  -c -DEXTERIOR_BLACK_PIXEL_ELIMINATION liveViz0.C
../../../../bin/charmc  -c -DEXTERIOR_BLACK_PIXEL_ELIMINATION liveViz.C
../../../../bin/charmc  -c -DEXTERIOR_BLACK_PIXEL_ELIMINATION liveVizPoll.C
../../../../bin/charmc  -c -DEXTERIOR_BLACK_PIXEL_ELIMINATION colorScale.C
../../../../bin/charmc  -c -DEXTERIOR_BLACK_PIXEL_ELIMINATION ImageData.C
../../../../bin/charmc  -c -DEXTERIOR_BLACK_PIXEL_ELIMINATION compat_float2
rgb.C
cp liveViz.h liveViz.decl.h liveVizPoll.decl.h liveViz0.h colorScale.h Imag
eData.h ../../../../include/
touch headers
../../../../bin/charmc  liveViz0.o liveViz.o liveVizPoll.o colorScale.o Ima
geData.o compat_float2rgb.o  -o ../../../../lib/libmoduleliveViz.a
Warning: creating ../../../../lib_so/libmoduleliveViz.so
ar: creating ../../../../lib/libmoduleliveViz.a
nilesh@skill:~/work/charm/src/libs/ck-libs/liveViz>
```

Running LiveViz

- Build and run the server
  - cd examples/charm++/lbServer
  - make
  - ./run_server.sh
- Or in detail…

```
nilesh@skill:~/work/charm/pgms/charm++/ccs/liveViz/pollserver> make
../../../../../bin/charmc  -c app.ci
../../../../../bin/charmc  -c app.C
nilesh@skill:~/work/charm/pgms/charm++/ccs/liveViz/pollserver> ./run_server.sh
Starting up server:
ccs: 2
ccs: Server IP = 128.174.241.98, Server port = 1234 $
CCS getImage handlers registered.  Waiting for clients...
```

Running LiveViz

- Run the client
  - liveViz [<host> [<port>]]
- Brings up a result window:

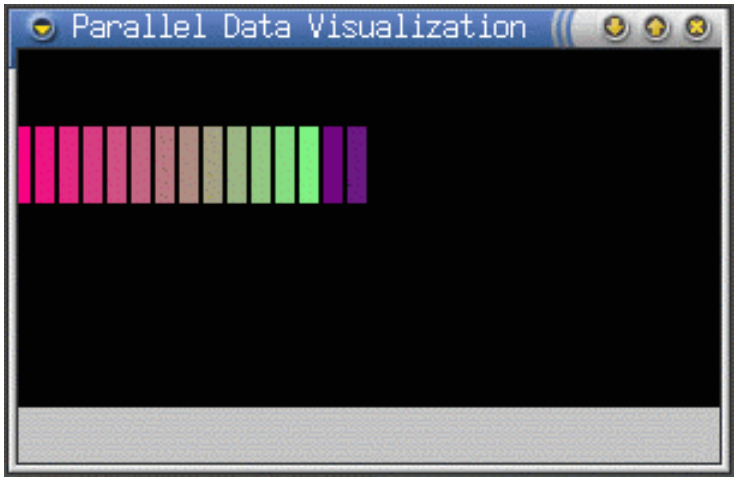

LiveViz Request Model

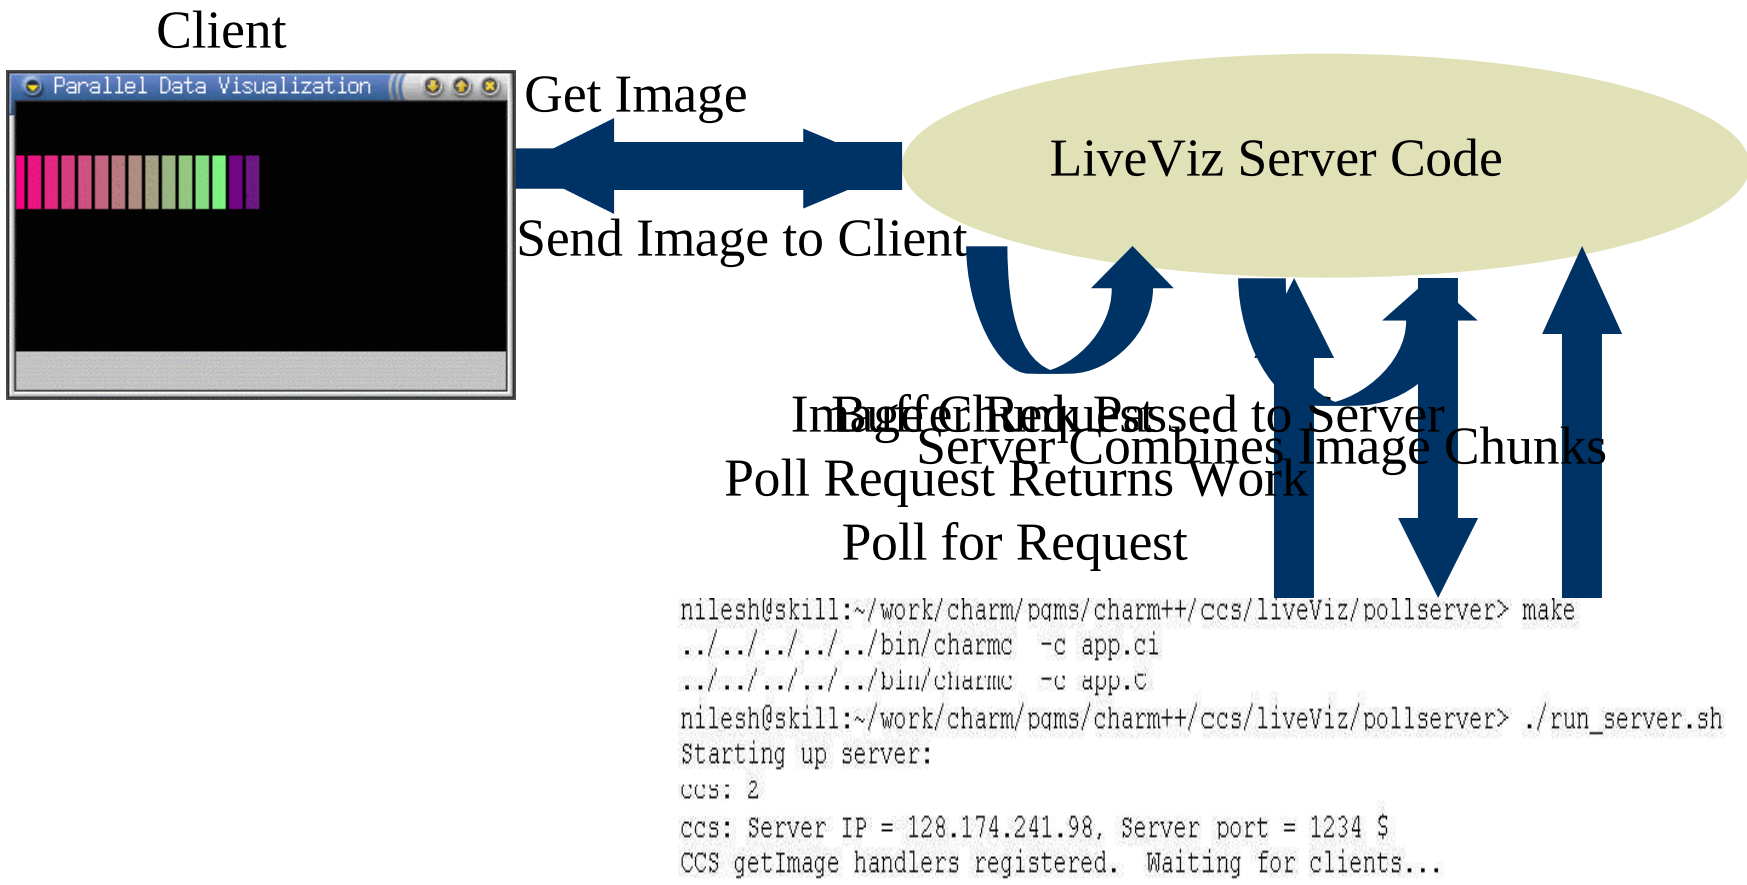

Jacobi 2D Example Structure

```
Main:
    Setup worker array, pass data to them

Workers:
    Start looping

    Send messages to all neighbors with ghost rows

    Wait for all neighbors to send ghost rows to me

    Once they arrive, do the regular Jacobi relaxation

    Calculate maximum error, do a reduction to compute
      global maximum error

    If timestep is a multiple of 64, load balance the
      computation. Then restart the loop.
```

LiveViz Setup

```
#include <liveVizPoll.h>

void main::main(. . .) {
  // Do misc initilization stuff

  // Create the workers and register with liveviz
  CkArrayOptions opts(0);     // By default allocate 0
                  // array elements.
  liveVizConfig cfg(true, true);  // color image = true and
                  // animate image = true
  liveVizPollInit(cfg, opts); // Initialize the library

  // Now create the jacobi 2D array
  work = CProxy_matrix::ckNew(opts);

  // Distribute work to the array, filling it as you do
}
```

Adding LiveViz To Your Code

```
void matrix::startTimeSlice() {
  // Send ghost row north, south, east, west, . . .
  sendMsg(dims.x-2, NORTH, dims.x+1, 1, +0, -1);

  // Now having sent all our ghosts, service liveViz
  // while waiting for neighbor's ghosts to arrive.
  serviceLiveViz();
}
```

```
void matrix::serviceLiveViz() {
  liveVizPollRequestMsg *m;
  while ( (m = liveVizPoll((ArrayElement *)this, timestep))
       != NULL ) {
    sendNextFrame(m);
  }
}
```

Generate an Image For a Request

```
void matrix::sendNextFrame(liveVizPollRequestMsg *m) {
  // Compute the dimensions of the image piece we'll send

  // Compute the image data of the chunk we'll send –
  // image data is just a linear array of bytes in row-major
  // order. For greyscale it's 1 byte, for color it's 3
  // bytes (rgb).

  // The liveViz library routine colorScale(value, min, max,
  // *array) will rainbow-color your data automatically.

  // Finally, return the image data to the library
  liveVizPollDeposit((ArrayElement *)this, timestep, m,
                     loc_x, loc_y, width, height, imageBits);
}
```

Link With The LiveViz Library

```
OPTS=-g
CHARMC=charmc $(OPTS)
LB=-module RefineLB

OBJS = jacobi2d.o

all: jacobi2d

jacobi2d: $(OBJS)
	$(CHARMC) -language charm++ \
	          -o jacobi2d $(OBJS) $(LB) -lm \
	          -module liveViz

jacobi2d.o: jacobi2d.C jacobi2d.decl.h
	$(CHARMC) -c jacobi2d.C
```

LiveViz Summary

- Easy to use visualization library
- Simple code handles any number of clients
- Doesn’t slow computation when there are no clients connected
- Works in parallel, with load balancing, etc.

Advanced Features

- Groups
- Node Groups
- Priorities
- Entry Method Attributes
- Communications Optimization
- Checkpoint/Restart

Conclusions

- Better Software Engineering
  - Logical Units decoupled from number of processors
  - Adaptive overlap between computation and communication
  - Automatic load balancing and profiling
- Powerful Parallel Tools
  - Projections
  - Parallel Debugger
  - LiveViz

More Information

- http://charm.cs.uiuc.edu
  - Manuals
  - Papers
  - Download files
  - FAQs
- ppl@cs.uiuc.edu