

# Chapel: an HPC language in a mainstream multicore world

Brad Chamberlain

UW CSE/MSR Summer Research Institute  
August 6, 2008



# Chapel

**Chapel:** a new parallel language being developed by Cray Inc.

## Themes:

- **general parallel programming**
  - data-, task-, and nested parallelism
  - express general levels of software parallelism
  - target general levels of hardware parallelism
- ***global-view* abstractions**
- ***multiresolution* design**
- **control of locality**
- **reduce gap between mainstream & parallel languages**

# Chapel's Setting: HPCS

## **HPCS:** High *Productivity* Computing Systems (DARPA *et al.*)

- **Goal:** Raise HEC user productivity by 10× for the year 2010
- **Productivity** = Performance
  - + Programmability
  - + Portability
  - + Robustness
- **Phase II:** Cray, IBM, Sun (July 2003 – June 2006)
  - Evaluated the entire system architecture's impact on productivity...
    - processors, memory, network, I/O, OS, runtime, compilers, tools, ...
    - ...and new languages:  
**Cray:** Chapel                      **IBM:** X10                      **Sun:** Fortress
- **Phase III:** Cray, IBM (July 2006 – 2010)
  - Implement the systems and technologies resulting from phase II
  - (Sun also continues work on Fortress, without HPCS funding)

# HPC vs. Mainstream Multicore

- Differences:
  - machine scale
  - performance requirements (?)
  - memory requirements (?)
  - robustness requirements (?)
  - workloads
  - programming community sizes and expertise areas
  
- Some interesting HPC(S) trends:
  - growing desire for software productivity, programmability
  - desire to better support nontraditional users
    - students just out of school with no C/Fortran/vi/emacs experience
    - scientists without strong parallel CS background
  - desire to leverage multicore technologies in larger systems
    - ideally without requiring hybrid programming models
  - and others that match opinions expressed in this meeting...

# Outline

✓ Chapel Context


➤ Defining my terms

☐ Some Chapel features } *(with an emphasis on themes and topics from this meeting)*

☐ Wrap-up



# Global-view vs. Fragmented

**Problem:** “Apply 3-pt stencil to vector”



# Global-view vs. Fragmented

**Problem:** “Apply 3-pt stencil to vector”



# Global-view vs. SPMD Code

**Problem:** “Apply 3-pt stencil to vector”

**global-view**

```
def main() {
  var n: int = 1000;
  var a, b: [1..n] real;

  forall i in 2..n-1 {
    b(i) = (a(i-1) + a(i+1))/2;
  }
}
```

**SPMD**

```
def main() {
  var n: int = 1000;
  var locN: int = n/numProcs;
  var a, b: [0..locN+1] real;

  if (iHaveRightNeighbor) {
    send(right, a(locN));
    rcv(right, a(locN+1));
  }

  if (iHaveLeftNeighbor) {
    send(left, a(1));
    rcv(left, a(0));
  }

  forall i in 1..locN {
    b(i) = (a(i-1) + a(i+1))/2;
  }
}
```



# Global-view vs. SPMD Code


**Problem:** “Apply 3-pt stencil to vector”

Assumes *numProcs* divides *n*;  
a more general version would  
require additional effort

global-view

```
def main() {
  var n: int = 1000;
  var a, b: [1..n] real;

  forall i in 2..n-1 {
    b(i) = (a(i-1) + a(i+1))/2;
  }
}
```




SPMD

```
def main() {
  var n: int = 1000;
  var locN: int = n/numProcs;
  var a, b: [0..locN+1] real;
  var innerLo: int = 1;
  var innerHi: int = locN;

  if (iHaveRightNeighbor) {
    send(right, a(locN));
    recv(right, a(locN+1));
  } else {
    innerHi = locN-1;
  }

  if (iHaveLeftNeighbor) {
    send(left, a(1));
    recv(left, a(0));
  } else {
    innerLo = 2;
  }

  forall i in innerLo..innerHi {
    b(i) = (a(i-1) + a(i+1))/2;
  }
}
```



# MPI SPMD pseudo-code

**Problem:** “Apply 3-pt stencil to vector”

## SPMD (pseudocode + MPI)

```


var n: int = 1000, locN: int = n/numProcs;
var a, b: [0..locN+1] real;
var innerLo: int = 1, innerHi: int = locN;
var numProcs, myPE: int;
var retval: int;
var status: MPI_Status;

MPI_Comm_size(MPI_COMM_WORLD, &numProcs);
MPI_Comm_rank(MPI_COMM_WORLD, &myPE);
if (myPE < numProcs-1) {
    retval = MPI_Send(&a(locN), 1, MPI_FLOAT, myPE+1, 0, MPI_COMM_WORLD);
    if (retval != MPI_SUCCESS) { handleError(retval); }
    retval = MPI_Recv(&a(locN+1), 1, MPI_FLOAT, myPE+1, 1, MPI_COMM_WORLD, &status);
    if (retval != MPI_SUCCESS) { handleErrorWithStatus(retval, status); }
} else
    innerHi = locN-1;
if (myPE > 0) {
    retval = MPI_Send(&a(1), 1, MPI_FLOAT, myPE-1, 1, MPI_COMM_WORLD);
    if (retval != MPI_SUCCESS) { handleError(retval); }
    retval = MPI_Recv(&a(0), 1, MPI_FLOAT, myPE-1, 0, MPI_COMM_WORLD, &status);
    if (retval != MPI_SUCCESS) { handleErrorWithStatus(retval, status); }
} else
    innerLo = 2;
forall i in (innerLo..innerHi) {
    b(i) = (a(i-1) + a(i+1))/2;
}


```





Communication becomes geometrically more complex for higher-dimensional arrays

# NAS MG *rprj3* stencil




=




	= $W_0$
	= $W_1$
	= $W_2$
	= $W_3$


=



+



+



# NAS MG *rprj3* stencil in Fortran + MPI

```

subroutine comm3(u,n1,n2,n3,kk)
use caf_intrinsics

implicit none
include 'cafnpb.h'
include 'globals.h'

integer n1, n2, n3, kk
double precision u(n1,n2,n3)
integer axis

if( .not. dead(kk) ) then
do axis = 1, 3
if( nprocs .ne. 1 ) then
call sync_all()
call give3( axis, +1, u, n1, n2, n3, kk )
call give3( axis, -1, u, n1, n2, n3, kk )
call sync_all()
call take3( axis, -1, u, n1, n2, n3 )
call take3( axis, +1, u, n1, n2, n3 )
else
call commlp( axis, u, n1, n2, n3, kk )
endif
enddo
else
do axis = 1, 3
call sync_all()
call sync_all()
enddo
call zero3(u,n1,n2,n3)
return
end

subroutine give3( axis, dir, u, n1, n2, n3, k )
use caf_intrinsics

implicit none
include 'cafnpb.h'
include 'globals.h'

integer axis, dir, n1, n2, n3, k, ierr
double precision u( n1, n2, n3 )

integer i3, i2, i1, buff_len, buff_id

buff_id = 2 + dir
buff_len = 0

if( axis .eq. 1 ) then
if( dir .eq. -1 ) then

do i3=2,n3-1
do i2=2,n2-1
buff_len = buff_len + 1
buff(buff_len, buff_id) = u( 2, i2, i3 )
enddo
enddo

buff(1:buff_len, buff_id+1) [nbr(axis, dir, k)] =
buff(1:buff_len, buff_id)

>
else if( dir .eq. +1 ) then

do i3=2,n3-1
do i2=2,n2-1
buff_len = buff_len + 1
buff(buff_len, buff_id) = u( n1-1, i2, i3 )
enddo
enddo

buff(1:buff_len, buff_id+1) [nbr(axis, dir, k)] =
buff(1:buff_len, buff_id)

>
endif
endif

if( axis .eq. 2 ) then
if( dir .eq. -1 ) then
do i3=2,n3-1
do i1=1,n1
buff_len = buff_len + 1
buff(buff_len, buff_id) = u( i1, 2, i3 )
enddo
enddo

buff(1:buff_len, buff_id+1) [nbr(axis, dir, k)] =
buff(1:buff_len, buff_id)

>
else if( dir .eq. +1 ) then

do i3=2,n3-1
do i1=1,n1
buff_len = buff_len + 1
buff(buff_len, buff_id) = u( i1, i2, 2 )
enddo
enddo

buff(1:buff_len, buff_id+1) [nbr(axis, dir, k)] =
buff(1:buff_len, buff_id)

>
else if( dir .eq. +1 ) then

do i2=1,n2
do i1=1,n1
buff_len = buff_len + 1
buff(buff_len, buff_id) = u( i1, i2, 2 )
enddo
enddo

buff(1:buff_len, buff_id+1) [nbr(axis, dir, k)] =
buff(1:buff_len, buff_id)

>
else if( dir .eq. -1 ) then

do i2=1,n2
do i1=1,n1
buff_len = buff_len + 1
buff(buff_len, buff_id) = u( i1, i2, n3-1 )
enddo
enddo

buff(1:buff_len, buff_id+1) [nbr(axis, dir, k)] =
buff(1:buff_len, buff_id)

>
endif
endif

subroutine take3( axis, dir, u, n1, n2, n3 )
use caf_intrinsics

implicit none
include 'cafnpb.h'
include 'globals.h'

integer axis, dir, n1, n2, n3
double precision u( n1, n2, n3 )

integer i3, i2, i1, buff_len, buff_id

buff_id = 3 + dir
buff_len = 0

if( axis .eq. 1 ) then
if( dir .eq. -1 ) then

do i3=2,n3-1
do i2=2,n2-1
buff_len = buff_len + 1
buff(buff_len, buff_id) = u( n1-1, i2, i3 )
enddo
enddo

buff(1:buff_len, buff_id) = 0.0D0

dir = +1
buff_id = 3 + dir
buff_len = nm2

do i=1, nm2
buff( i, buff_id ) = 0.0D0
enddo

dir = +1
buff_id = 3 + dir
buff_len = nm2

do i=1, nm2
buff( i, buff_id ) = 0.0D0
enddo

dir = +1
buff_id = 2 + dir
buff_len = 0

if( axis .eq. 1 ) then
do i3=2,n3-1
do i2=2,n2-1
buff_len = buff_len + 1
buff(buff_len, buff_id) = u( n1-1, i2, i3 )
enddo
enddo

else if( dir .eq. +1 ) then
do i3=2,n3-1
do i2=2,n2-1
buff_len = buff_len + 1
buff(buff_len, buff_id) = u( n1-1, i2, i3 )
enddo
enddo

endif
endif

if( axis .eq. 2 ) then
if( dir .eq. -1 ) then
do i3=2,n3-1
do i1=1,n1
buff_len = buff_len + 1
buff(buff_len, buff_id) = u( i1, n2-1, i3 )
enddo
enddo

buff(1:buff_len, buff_id+1) [nbr(axis, dir, k)] =
buff(1:buff_len, buff_id)

>
else if( dir .eq. +1 ) then

do i3=2,n3-1
do i1=1,n1
buff_len = buff_len + 1
buff(buff_len, buff_id) = u( i1, n2, i3 )
enddo
enddo

buff(1:buff_len, buff_id+1) [nbr(axis, dir, k)] =
buff(1:buff_len, buff_id)

>
endif
endif

if( axis .eq. 3 ) then
if( dir .eq. -1 ) then

do i2=1,n2
do i1=1,n1
buff_len = buff_len + 1
buff(buff_len, buff_id) = u( i1, i2, n3-1 )
enddo
enddo

buff_id = 2 + dir
buff_len = 0

if( axis .eq. 1 ) then
do i3=2,n3-1
do i2=2,n2-1
buff_len = buff_len + 1
buff(buff_len, buff_id) = u( 2, i2, i3 )
enddo
enddo

if( axis .eq. 2 ) then
do i3=2,n3-1
do i1=1,n1
buff_len = buff_len + 1
buff(buff_len, buff_id) = u( i1, i2, 2 )
enddo
enddo

if( axis .eq. 3 ) then
do i2=1,n2
do i1=1,n1
buff_len = buff_len + 1
buff(buff_len, buff_id) = u( i1, i2, 2 )
enddo
enddo

endif
endif

subroutine commlp( axis, u, n1, n2, n3, kk )
use caf_intrinsics

implicit none
include 'cafnpb.h'
include 'globals.h'

integer axis, dir, n1, n2, n3
double precision u( n1, n2, n3 )

integer i3, i2, i1, buff_len, buff_id
integer i, kk, indx

dir = -1
buff_id = 3 + dir
buff_len = nm2

do i=1, nm2
buff( i, buff_id ) = 0.0D0
enddo

dir = +1
buff_id = 3 + dir
buff_len = nm2

do i=1, nm2
buff( i, buff_id ) = 0.0D0
enddo

dir = +1
buff_id = 2 + dir
buff_len = 0

if( axis .eq. 1 ) then
do i3=2,n3-1
do i2=2,n2-1
buff_len = buff_len + 1
buff(buff_len, buff_id) = u( n1-1, i2, i3 )
enddo
enddo

12, i3 )
enddo
enddo

endif
endif

if( axis .eq. 2 ) then
do i3=2,n3-1
do i1=1,n1
buff_len = buff_len + 1
buff(buff_len, buff_id) = u( i1, n2-1, i3 )
enddo
enddo

1, i3 )
enddo
enddo

endif
endif

if( axis .eq. 3 ) then
do i2=1,n2
do i1=1,n1
buff_len = buff_len + 1
buff(buff_len, buff_id) = u( i1, i2, i3 )
enddo
enddo

endif
endif

subroutine comm3(s,m1j,m2j,m3j,k)
implicit none
include 'cafnpb.h'
include 'globals.h'

integer mk, m2k, m3k, m1j, m2j, m3j, k

double precision r(mk,m2k,m3k), s(m1j,m2j,m3j)
integer j3, j2, j1, i3, i2, i1, d1, d2, d3, j
double precision x1(m), y1(m), x2,y2

if(m1k.eq.3) then
d1 = 2
else
d1 = 1
endif

if(m2k.eq.3) then
d2 = 2
else
d2 = 1
endif

if(m3k.eq.3) then
d3 = 2
else
d3 = 1
endif

do j3=2,m3j-1
i3 = 2*j3-d3
do j2=2,m2j-1
i2 = 2*j2-d2
do j1=2,m1j-1
i1 = 2*j1-d1

x1(i1-1) = r(i1-1,i2-1,i3) ) + r(i1-1,i2+1,i3 )
>
+ r(i1-1,i2, i3-1) + r(i1-1,i2, i3+1)
>
y1(i1-1) = r(i1-1,i2-1,i3-1) + r(i1-1,i2-1,i3+1)
>
+ r(i1-1,i2+1,i3-1) + r(i1-1,i2+1,i3+1)

enddo
enddo
do j1=2,m1j-1
i1 = 2*j1-d1
y2 = r(i1, i2-1,i3-1) + r(i1, i2-1,i3+1)
>
+ r(i1, i2+1,i3-1) + r(i1, i2+1,i3+1)
>
x2 = r(i1, i2-1, i3 ) + r(i1, i2+1, i3 )
>
+ r(i1, i2, i3-1) + r(i1, i2, i3+1)
>
s(j1,j2,j3) =
>
0.5D0 * r(i1,i2,i3)
>
+ 0.25D0 * ( r(i1-1,i2,i3) + r(i1+1,i2,i3) + x2 )
>
+ 0.125D0 * ( x1(i1-1) + x1(i1+1) + y2 )
>
+ 0.0625D0 * ( y1(i1-1) + y1(i1+1) )

enddo
enddo
enddo
j = k-1
call comm3(s,m1j,m2j,m3j,j)
return
end

```


# NAS MG *rprj3* stencil in Chapel

```

def rprj3(S, R) {
  param Stencil = [-1..1, -1..1, -1..1],
    w: [0..3] real = (0.5, 0.25, 0.125, 0.0625),
    w3d = [(i,j,k) in Stencil] w((i!=0) + (j!=0) + (k!=0));

  forall ijk in S.domain do
    S(ijk) = + reduce [offset in Stencil]
      (w3d(offset) * R(ijk + offset*R.stride));
}

```



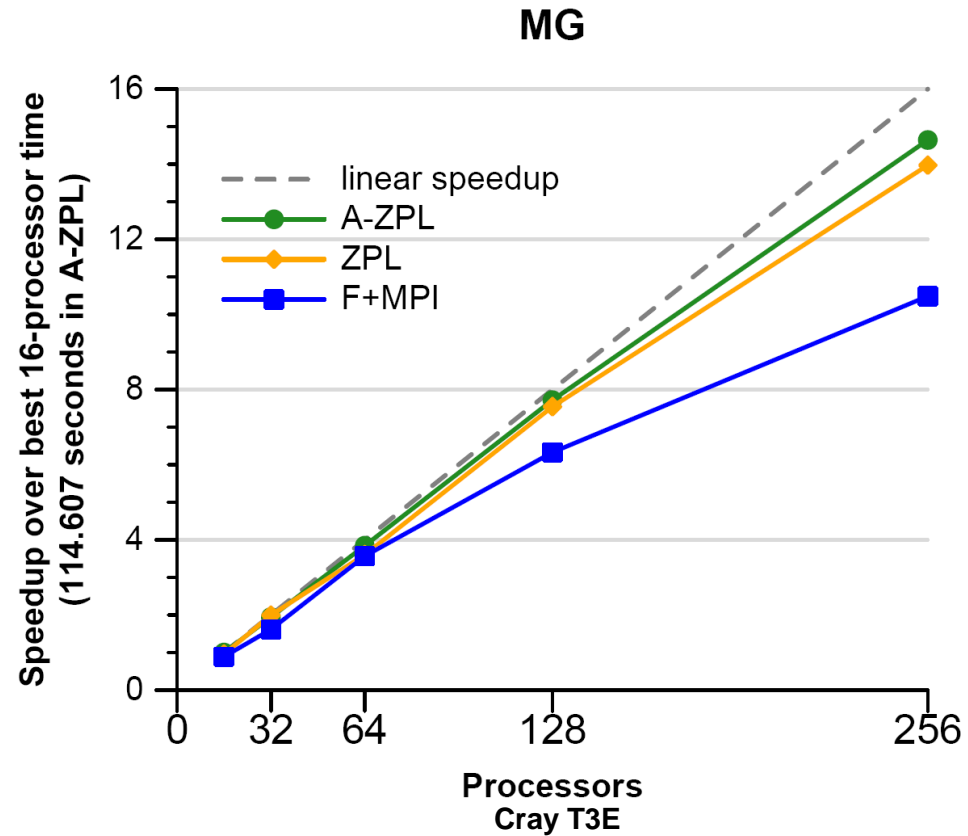
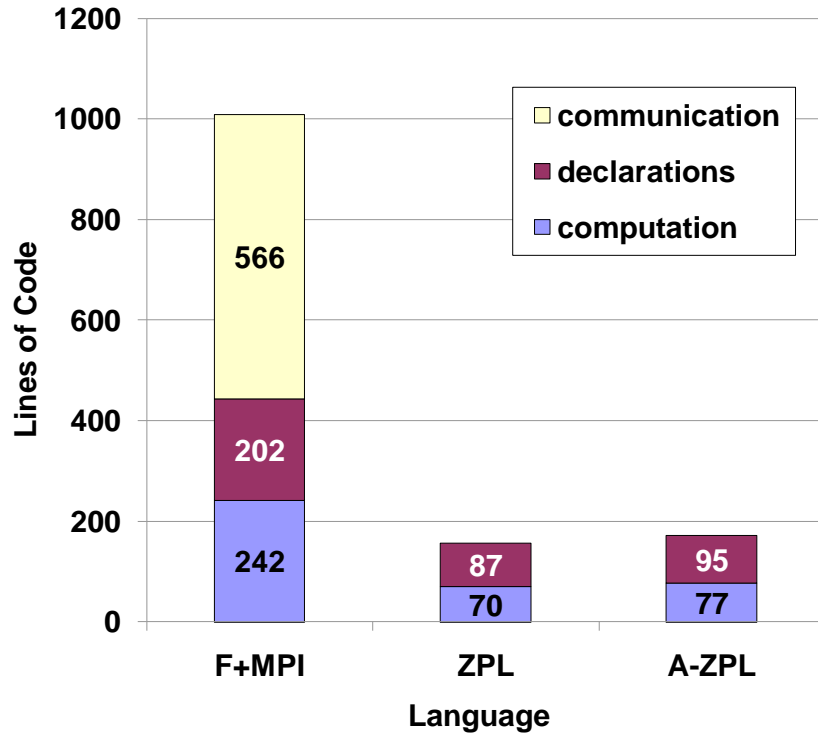
# NAS MG *rprj3* stencil in ZPL

```


procedure rprj3(var S,R: [, , ] double;
                 d: array [] of direction);
begin
  S := 0.5 * R
    + 0.25 * (R@^d[ 1, 0, 0] + R@^d[ 0, 1, 0] + R@^d[ 0, 0, 1] +
              R@^d[-1, 0, 0] + R@^d[ 0, -1, 0] + R@^d[ 0, 0, -1])
    + 0.125 * (R@^d[ 1, 1, 0] + R@^d[ 1, 0, 1] + R@^d[ 0, 1, 1] +
              R@^d[ 1, -1, 0] + R@^d[ 1, 0, -1] + R@^d[ 0, 1, -1] +
              R@^d[-1, 1, 0] + R@^d[-1, 0, 1] + R@^d[ 0, -1, 1] +
              R@^d[-1, -1, 0] + R@^d[-1, 0, -1] + R@^d[ 0, -1, -1])
    + 0.0625 * (R@^d[ 1, 1, 1] + R@^d[ 1, 1, -1] +
                R@^d[ 1, -1, 1] + R@^d[ 1, -1, -1] +
                R@^d[-1, 1, 1] + R@^d[-1, 1, -1] +
                R@^d[-1, -1, 1] + R@^d[-1, -1, -1]);
end;

```


# NAS MG: Fortran + MPI vs. ZPL



# Parallel Programming Models: Two Camps



“Why is everything so painful?”



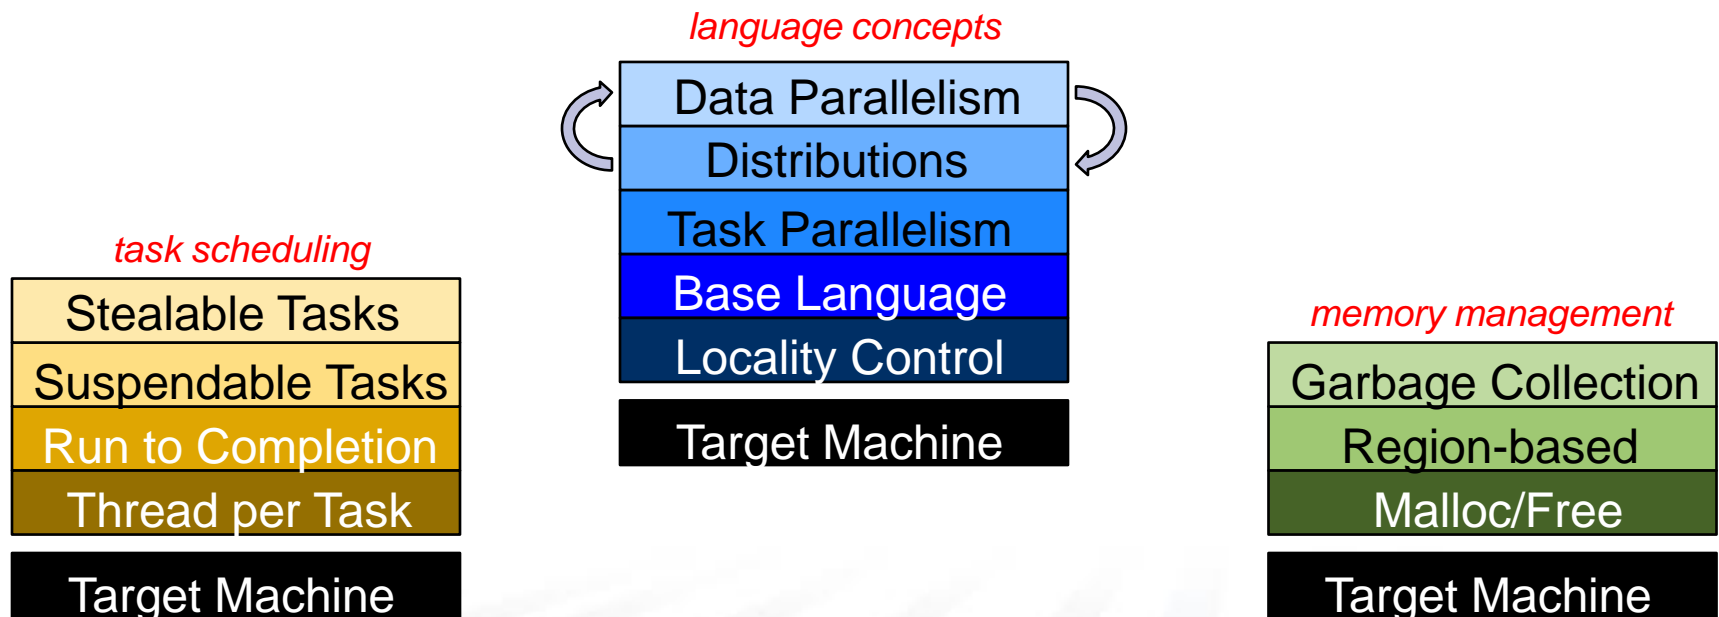
“Why do my hands feel tied?”



# Multiresolution Language Design

**Our Approach:** Permit the language to be utilized at multiple levels, as required by the problem/programmer


- provide high-level features and automation for convenience
- provide the ability to drop down to lower, more manual levels
- use appropriate separation of concerns to keep these layers clean



# Outline

- ✓ Chapel Context
- ✓ Defining my terms
- Some Chapel features
- Wrap-up


# Domains



*domain*: a first-class index set


```
var m = 4, n = 8;
```

```
var D: domain(2) = [1..m, 1..n];
```



*D*

# Domains




*domain*: a first-class index set


```
var m = 4, n = 8;
```

```
var D: domain(2) = [1..m, 1..n];
```

```
var Inner: subdomain(D) = [2..m-1, 2..n-1];
```




# Domains: Some Uses




- Declaring arrays:

```
var A, B: [D] real;
```




- Iteration (sequential or parallel):

```
for ij in Inner { ... }
or: forall ij in Inner { ... }
or: ...
```




- Array Slicing:

```
A[Inner] = B[Inner];
```





- Array reallocation:


```
D = [1..2*m, 1..2*n];
```




# Domains: Different Flavors



*dense*




*strided*




*sparse*

*graphs*





*associative*




# Locality: Locales

*locale*: architectural unit of locality

- has capacity for processing and storage
- threads within a locale have ~uniform access to local memory
- memory within other locales is accessible, but at a price
- e.g., a multicore processor or SMP node could be a locale



# Locality: Locales



- user specifies # locales on executable command-line

```
prompt> myChapelProg -nl=8
```

- Chapel programs have built-in domain/array of locales:

```
config const numLocales: int;
```

```
const LocaleSpace = [0..numLocales-1],
```

0	1	2	3	4	5	6	7
---	---	---	---	---	---	---	---

```
  Locales: [LocaleSpace] locale;
```


- Programmers can create their own locale views:

```
const CompGrid = Locales.reshape([1..GridRows,
                                  1..GridCols]);
```

0	1	2	3
4	5	6	7




# Distributions: Overview



Domains may be distributed across locales

```
var D: domain(2) distributed Block on CompGrid = ...;
```

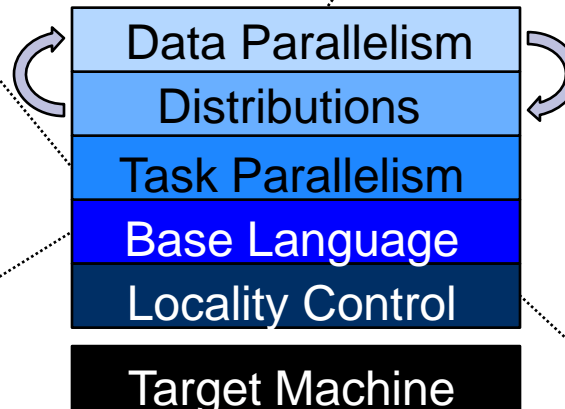


- A distribution implies...
  - ...ownership of the domain's indices (and its arrays' elements)
  - ...the default work ownership for operations on the domain/arrays
- Advanced users can write their own distributions (and standard distributions are written using the identical mechanism)
- A distribution must implement...
  - ...the mapping from indices to locales
  - ...the per-locale representation of domain indices and array elements
  - ...the compiler's target interface for lowering global-view operations

# Other Chapel Features

- **unstructured task parallelism**
  - begins
  - sync statements
- **structured task parallelism**
  - cobegins
  - coforalls
- **synchronization**
  - sync/single variables
  - atomic blocks

- subdomains and index types
- zippered and tensor iteration
- scalar function promotion
- reductions and scans



- value- and reference-based OOP (optional)
- generic programming features
- latent types / shallow static type inference
- rich compile-time language
- iterators (generators)
- tuples

- fine-grain control over task and data placement via on clauses

# Outline

- ✓ Chapel Context
- ✓ Defining my terms
- ✓ Some Chapel features
- Wrap-up

# Chapel and Mainstream Multicore

- While Chapel doesn't specifically target mainstream multicore, it could be applicable
  - removes much tedium and nitpicky details from data parallelism
  - raises level of discourse for task parallelism beyond threads
  - though not a dialect of a mainstream language, not far afield either
    - programmers today seem more multilingual than in the past
    - interoperability more crucial than extending a language
- Chapel's locales and distributions are likely overkill for today's multicore processors
  - but what about for future generations of multicore?
- Chapel team does most of our development and testing on mainstream multicore machines
  - Linux, Mac, Windows, ...
  - AMD, Intel, ...

# Takeaways

- Key themes for mainstream multicore languages
  - generality, to the extent possible
  - global-view abstractions (for Joe & Joan)
  - multiresolution design (for Steve & Stephanie)
  
- Other Lessons from HPC
  - abstract implementing mechanisms further from user's view than we traditionally have
  - SPMD: an execution case to optimize for, not the only tool in the box

# Future Directions (with an eye on multicore)

- expand locale concept
  - hierarchical locales
    - to expose hierarchy within a node
  - heterogeneous locale types
    - to describe coarse-grain HW heterogeneity
  - dynamically varying numbers of locales
  
- equivalent of “distribution” concept for task parallelism
  - (Edward Lee’s directors?)
  
- and of course, many others...

# Collaborations

**UIUC (Vikram Adve and Rob Bocchino):** Software Transactional Memory (STM) over distributed memory (PPoPP `08)

**ORNL (David Bernholdt *et al.*):** Chapel code studies – Fock matrix computations, MADNESS, Sweep3D, ... (HIPS `08)

**PNNL (Jarek Nieplocha *et al.*):** ARMCI port of comm. layer

**EPCC (Michele Weiland, Thom Haddow):** performance study of single-locale task parallelism

**CMU (Franz Franchetti):** Chapel as portable parallel back-end language for SPIRAL

(Your name here?)

# Chapel Contributors

## ■ Current Team

- Brad Chamberlain
- Steve Deitz
- Samuel Figueroa
- David Iten
- Andy Stone (not shown)

## ■ Alumni

- Robert Bocchino
- David Callahan
- James Dinan
- Roxana Diaconescu
- Shannon Hoffswell
- Mary Beth Hribar
- Mark James
- Mackale Joyner
- John Plevyak
- Wayne Wong
- Hans Zima





# Questions?

bradc@cray.com  
chapel\_info@cray.com

<http://chapel.cs.washington.edu>