

The Cascade High-Productivity Language

Chapel Motivation

Overview

Chapel: a new parallel programming language being developed by Cray Inc. as part of DARPA's High Productivity Computing Systems program (HPCS)

Motivating Themes:

- **general parallel programming**
 - to support data-, task-, and nested parallelism
 - to support general granularities of software parallelism
 - to target general levels of hardware parallelism
- **global-view abstractions**
 - to reduce the number of programmer-managed details
- **multiresolution design**
 - to support abstraction without giving up low-level tuning
- **reduce mainstream-to-parallel language gap**
 - to engage the next generation of programmers
 - to take advantage of advances in language design

Productivity

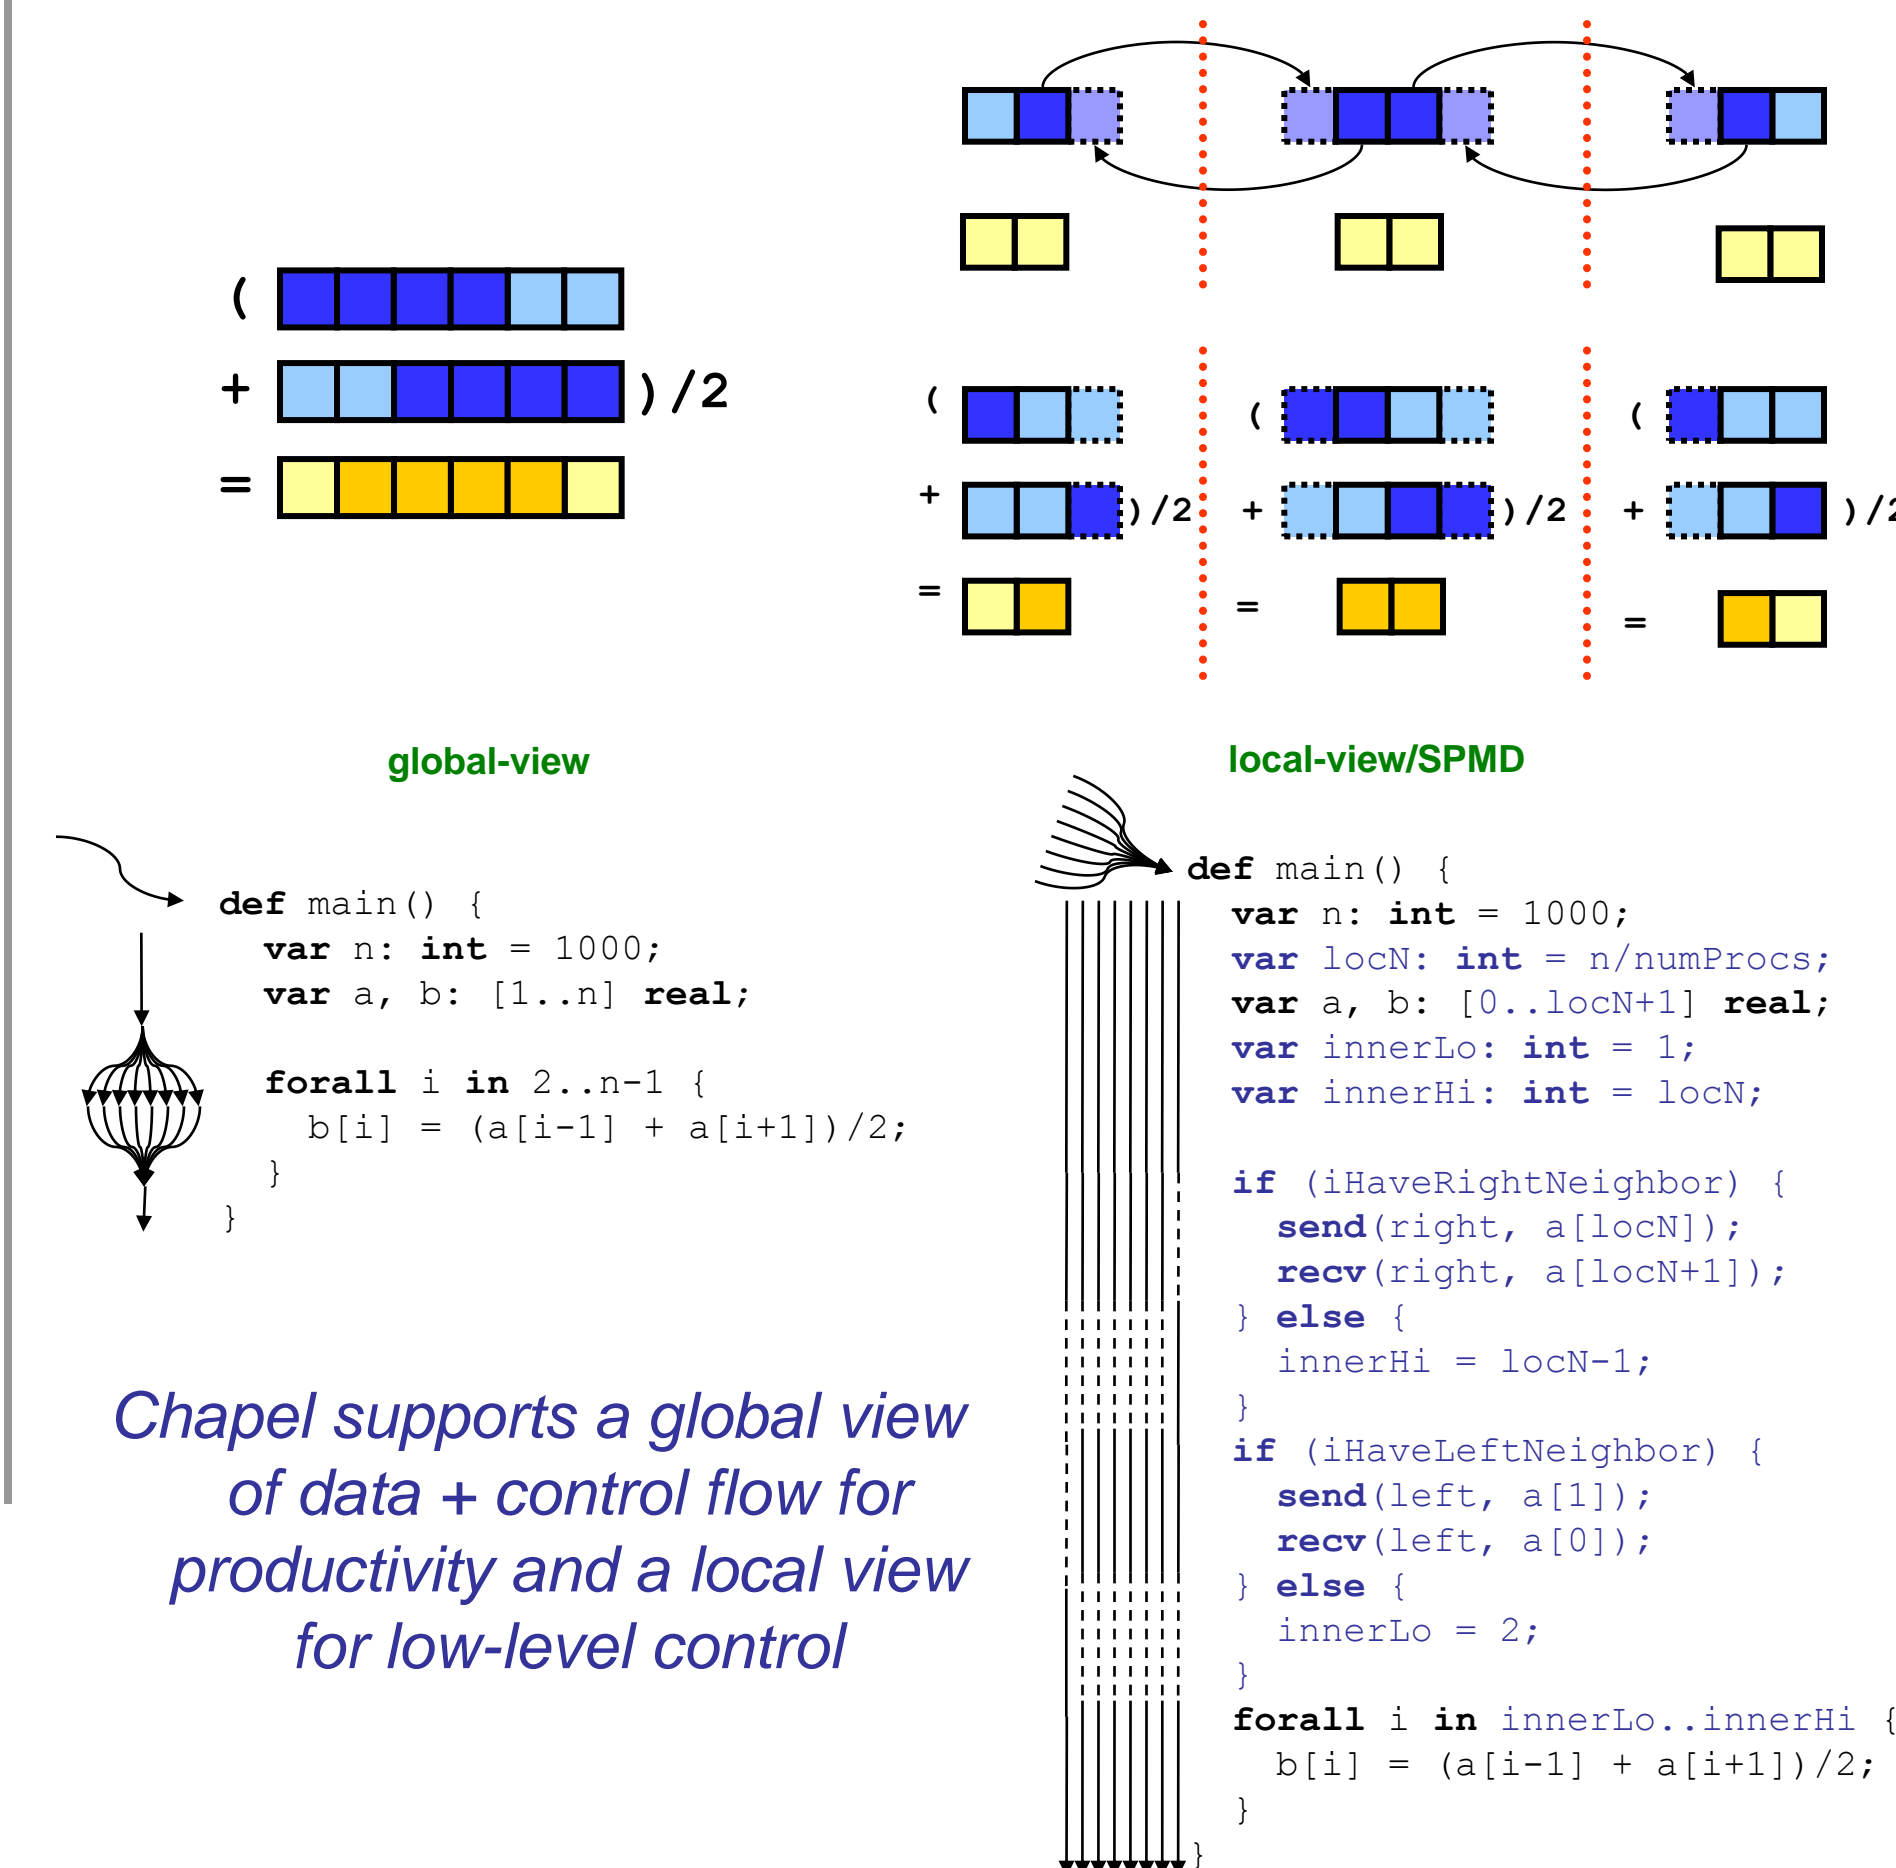
HPCS Goal: Raise user productivity by 10x

Chapel's productivity goals:

- Vastly improve **programmability**
 - writing parallel codes
 - reading, modifying, porting, tuning, maintaining them
- Support **performance** at least as good as MPI
 - competitive with MPI on generic clusters
 - better than MPI on more capable architectures
- Improve **portability**
 - as ubiquitous as MPI, but with fewer ties to architecture
 - more portable than other parallel programming models
- Improve **code robustness**
 - eliminate common error cases
 - support better abstractions to help avoid other errors

Global View of Data and Control

Example: "Apply a 3-point stencil to a vector"



Task Parallelism

begin: creates a task for concurrent evaluation

```

begin DoThisTask();
WhileContinuing();
TheOriginalThread();
    
```

sync: waits on all tasks created within its dynamic scope

```

sync {
  begin treeSearch(root);
}

def treeSearch(node) {
  if node == nil then return;
  begin treeSearch(node.right);
  begin treeSearch(node.left);
}
    
```

sync/single variables: store full/empty state with their value

```

var result$: sync real; // result is initially empty
sync {
  begin ... = result$; // block until full, leave empty
  begin result$ = ...; // block until empty, leave full
}
result$.readXX(); // non-blocking, leaves state unchanged;
// other variations also supported
    
```

cobegin statement: creates a task per component statement

```

computePivot(lo, hi, data);
cobegin {
  Quicksort(lo, pivot, data);
  Quicksort(pivot, hi, data);
} // implicit join here
    
```

coforall loop: creates a task per loop iteration

```

coforall e in Edges {
  exploreEdge(e);
} // implicit join here
    
```

Data Parallelism

domain: a first-class index set

```

const m = 4, n = 8;
const D: domain(2) = [1..m, 1..n];
var Inner: subdomain(D) = [2..m-1, 2..n-1];
    
```

Domain uses:

- Declaring arrays:

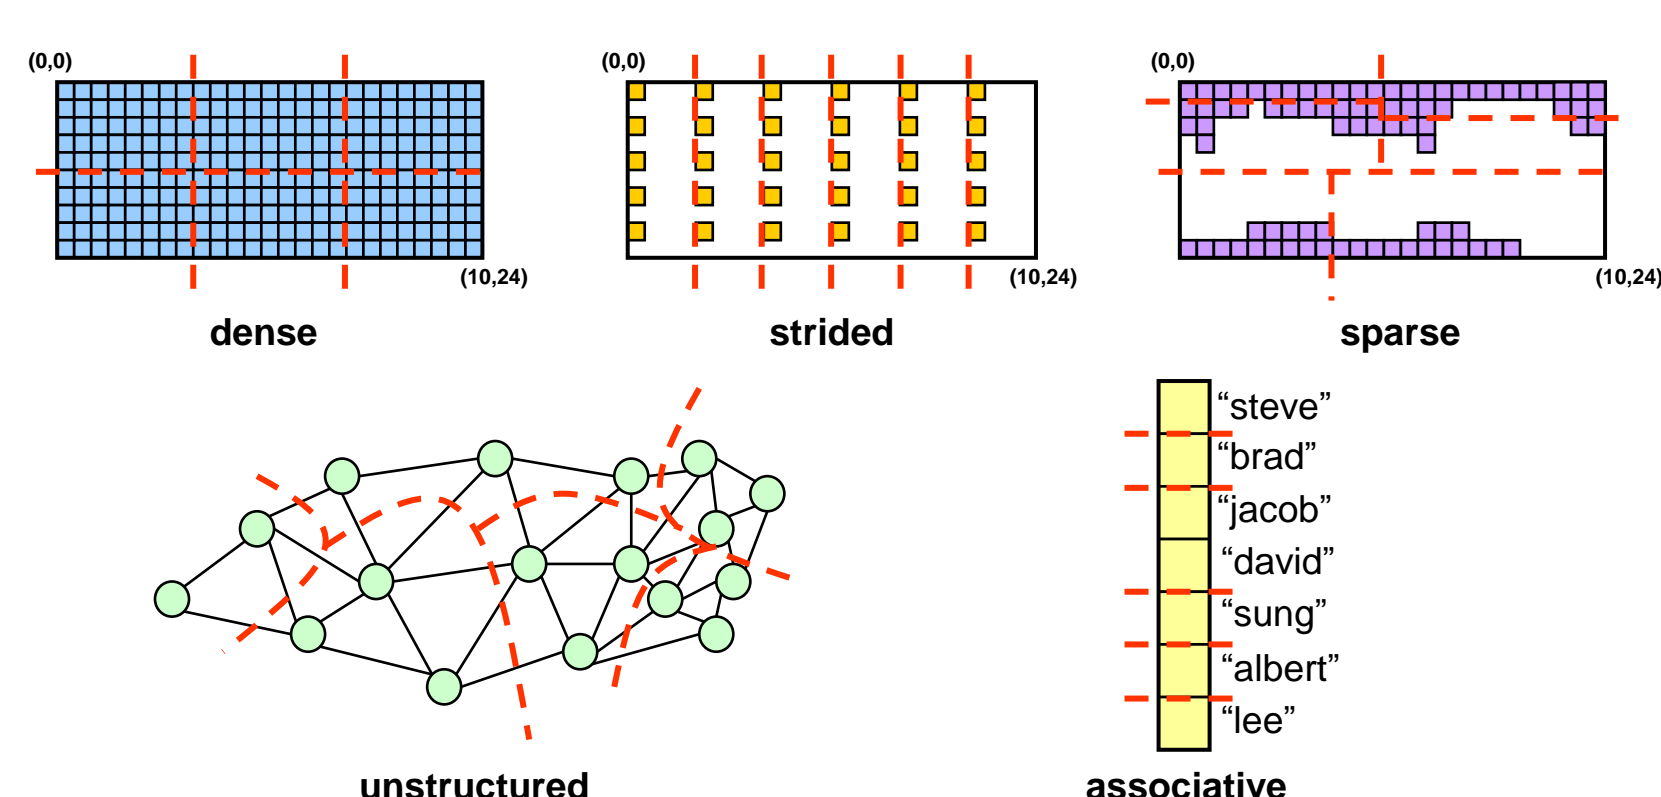

```
var A, B: [D] real;
```
- Iteration (sequential or parallel):


```
for ij in Inner { ... }
forall ij in Inner { ... }
```
- Array Slicing:


```
A[Inner] = B[Inner];
```
- Array reallocation:


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

Domain types:



Locality

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

user specifies # locales on executable command-line

```
prompt> myChapelProg -n1=8
```

Chapel programs have built-in locale variables:

```

config const numLocales: int;
const LocaleSpace = [0..numLocales-1];
Locales: [LocaleSpace] locale;
    
```

Programmers can create their own locale views:

```

var CompGrid = Locales.reshape([1..GridRows, 1..GridCols]);
var TaskALocs = Locales[1..numTaskALocs];
var TaskBlocs = Locales[numTaskALocs+1..];
    
```

on clauses: indicate where tasks should execute

```

cobegin {
  on TaskALocs do computeTaskA(...);
  on TaskBlocs do computeTaskB(...);
  on Locales(0) do computeTaskC(...);
}
    
```

Domains may be mapped across locales

```

var D: domain(2) dmapped Block(...) = ...;
    
```

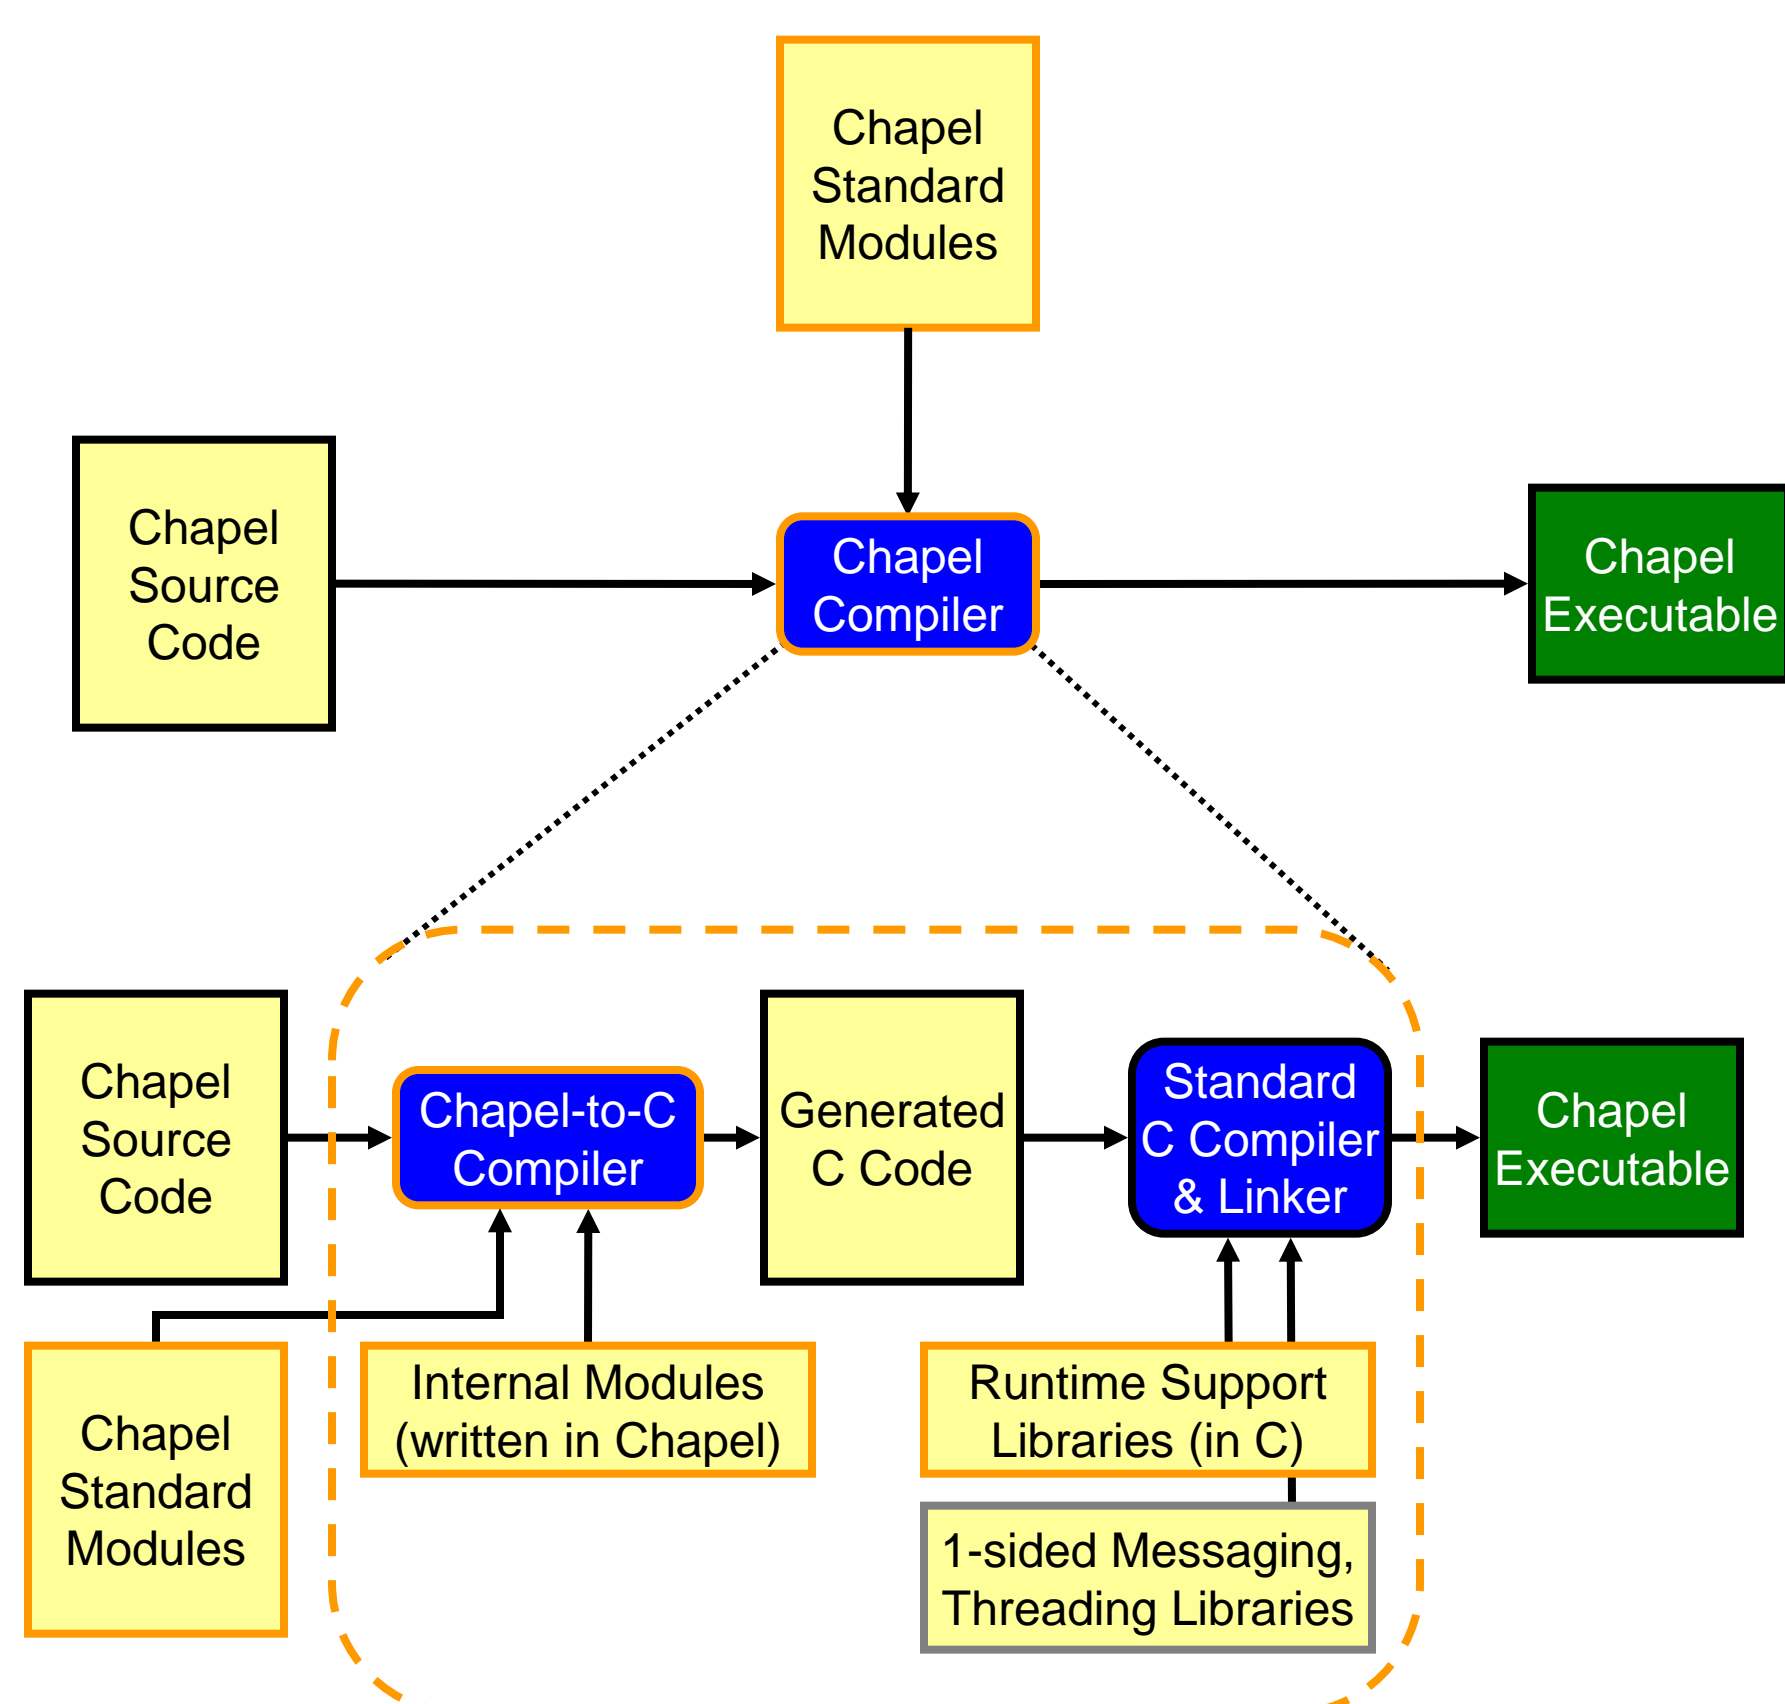
A domain map defines...

- ...ownership of the domain's indices (and its arrays' elements)
- ...the default work ownership for operations on the domains/arrays

Chapel provides...

- ...a standard library of distributions (Block, Block-Cyclic, Recursive Bisection, ...)
- ...the means for advanced users to write their own distributions

Prototype Compiler



Implementation Status

Public Release

- **Latest release:** version 1.2 (October 21, 2010)
- **Supported platforms:**
 - desktop Linux/UNIX, Mac OS X, Cygwin, ...
 - commodity clusters of same
 - Cray platforms
 - Systems from other vendors
- Download from: <http://sourceforge.net/projects/chapel>
- Your feedback desired!
- Keep in mind this is a work-in-progress:
 - ⇒ it's likely that you will find problems with the implementation
 - ⇒ this is still a reasonable time to influence the language

- **Base language:** stable, yet incomplete: I/O, OOP, exceptions
- **Task parallelism:**
 - multi-threaded, multi-locale implementation of tasks, sync variables
- **Data parallelism:**
 - multi-threaded, multi-locale regular domains/arrays
 - multi-threaded, single-locale irregular domains/arrays
- **Locality:**
 - locale types and arrays
 - richer set of distributions to be developed: Block-Cyclic, MRD, ...
- **Performance:**
 - has received much attention in designing the language
 - work remains in communication optimizations; serial code

Notable Collaborations

- Notre Dame/ORNL** (Peter Kogge, Srinivas Sridharan, Jeff Vetter): Asynchronous STM over distributed memory
- UIUC** (David Padua, Maria Garzarán, Albert Sidelnik): Chapel for hybrid CPU-GPU computing
- Universitat Politècnica de Catalunya** (Alex Duran): Chapel over Nanos++ user-level tasking
- Argonne** (Rusty Lusk, Rajeev Thakur, Pavan Balaji): Chapel over MPICH one-sided communication
- Sandia** (Richard Murphy, Kyle Wheeler): Chapel over Qthreads user-level tasking
- UT Austin** (Calvin Lin, Karthik Murthy): Productive memory consistency models
- CU Boulder** (Jeremy Siek, Jonathan Turner): interfaces, concepts, and generics
- U. Oregon/Paratools Inc.** (Sameer Shende): Performance analysis with Tau

Interested in collaborating?
See our website for possible topics; or propose your own.

