

The Fortress Parallel Programming Language

Guy Steele

Sun Microsystems Laboratories

1 Network Drive UBUR02-311 Burlington, MA 02183 USA

+1 781 442 2620 guy.steele@sun.com

October 31, 2006

GSPx Multi-core Applications Conference

Copyright © 2006 Sun Microsystems, Inc. ("Sun"). All rights are reserved by Sun except as expressly stated as follows. Permission is given to Global Technology Conferences, Inc., to distribute this work in any form as part of the conference proceeding materials distributed by it for GSPx Multicore Applications 2006, provided that the work is distributed in its entirety, including this notice. Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted, provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers, or to redistribute to lists, requires prior specific written permission of Sun.

Fortress: “To Do for Fortran What Java™ Did for C”

- Catch “stupid mistakes” (like array bounds errors)
- Extensive libraries (e.g., for network environment)
- Security model (including type safety)
- Dynamic compilation
- Platform independence
- Multithreading

- Make programmers more productive

The Context of the Research

- Improving programmer productivity for scientific and engineering applications
- Research funded in part by the DARPA IPTO (Defense Advanced Research Projects Agency Information Processing Technology Office) through their High Productivity Computing Systems program
- Goal is economically viable technologies for both government and industrial applications by the year 2010 and beyond
- Targeted to a range of hardware, from petaflop supercomputers to single/multiple multicore chips

Key Ideas

- Don't build the language—grow it
- Make programming notation closer to math
- Ease use of parallelism

Growing a Language

- Languages have gotten much bigger
- You can't build one all at once
- Therefore it must grow over time
- What happens if you design it to grow?
- How does the need to grow affect the design?
- Need to grow a user community, too

See Steele, “Growing a Language” keynote talk, OOPSLA 1998;
Higher-Order and Symbolic Computation **12**, 221–236 (1999)

Interesting Language Design Strategy

Wherever possible,
consider whether a proposed language feature
can be provided by a library
rather than having it built into the compiler.

To make this work, library designers need
substantial control over syntax and semantics—
not just the ability to code new functions or methods
invoked by a single method call syntax `foo(a, b, c)`.

Minimalist Approach

- As few primitive types as possible (cf. Bacon's Kava)
 - > Objects with methods and fields
 - > Binary words of many different sizes
 - > Linear sequences (fixed length known at compile time)
 - > Heap sequences (fixed length known at allocation time)
- User-defined parameterized types
- User-defined polymorphic (overloaded) operators
- Aggressive type inference to reduce clutter
 - > Many variables require no type declarations
- Aggressive static and dynamic optimization

Types Defined by Libraries

- Lists, vectors, sets, multisets, and maps
 - > Like C Standard Template Library, but better notation

$$\langle 1, 2, 4, 3, 4 \rangle \quad A \cup \{1, 2, 3, 4\}$$

$$[3 \ 4 \ 5] \times [1 \ 0 \ 0]$$

- Matrices and multidimensional arrays
- Integers, floats, rationals, **with physical units**

m : \mathbb{R} Mass = 3.7 kg

\mathbf{v} : \mathbb{R}^3 Velocity = [3.5 0 1] m/s

\mathbf{p} : \mathbb{R}^3 Momentum = $m \mathbf{v}$

- Data structures may be local or distributed

ASCII (“Wiki-like markup”) Notation

- Lists, vectors, sets, multisets, and maps
 - > Like C Standard Template Library, but better notation

<|1, 2, 3, 4|> **A UNION {1, 2, 3, 4}**
[3 4 5] CROSS [1 0 0]

- Matrices and multidimensional arrays
- Integers, floats, rationals, with physical units

m: RR Mass = 3.7 kg_

v: RR^3 Velocity = [3.5 0 1] m/s_

_p: RR^3 Momentum = m _v

- Data structures may be local or distributed

Replaceable Components

- Avoid a monolithic “Standard Library”
- Replaceable components with version control
- Encourage alternate implementations
 - > Performance choices
 - > Test them against each other
- Encourage experimentation
 - > Framework for alternate language designs

Data and Control Models

- Data model: shared global address space
- Control model: multithreaded
 - > Basic primitive is “spawn”
 - > We hope application code seldom uses it
- Declared distribution of data and threads
 - > Managing aggregates integrated into type system
 - > Policies programmed as libraries, not wired in
- Transactional access to shared variables
 - > Atomic blocks (implicit or explicit retry)
 - > Lock-free (no blocking, no deadlock)

Simple Example: NAS CG Kernel

```

conjGrad || Elt extends Number, nat N,
          Mat extends Matrix || Elt, N × N ||,
          Vec extends Vector || Elt, N ||
          || (A : Mat, x : Vec) : (Vec, Elt)

```

```
cgitmax = 25
```

```
z : Vec := 0
```

```
r : Vec := x
```

```
p : Vec := r
```

```
ρ : Elt := rT r
```

```
for j ← seq(1 : cgitmax) do
```

```
  q = A p
```

$$\alpha = \frac{\rho}{p^T q}$$

```
z := z + α p
```

```
r := r - α q
```

```
ρ0 = ρ
```

```
ρ := rT r
```

$$\beta = \frac{\rho}{\rho_0}$$

```
p := r + β p
```

```
end
```

```
(z, ||x - A z||)
```

It's not new or surprising that code written in a programming language might be displayed in a conventional math-like format. The point of this example is how similar the code is to the math notation: the gap between the two syntaxes is relatively small. We want to see what will happen if a principal goal of a new language design is to minimize this gap.

Comparison: NAS NPB 1 Specification

```

z = 0
r = x
ρ = rT r
p = r
DO i = 1, 25
    q = A p
    α = ρ / (pT q)
    z = z + α p
    ρ0 = ρ
    r = r - α q
    ρ = rT r
    β = ρ / ρ0
    p = r + β p
ENDDO
compute residual norm explicitly: ||r|| = ||x - A z||
  
```

```

z : Vec := 0
r : Vec := x
p : Vec := r
ρ : Elt := rT r
for j ← seq(1 : cgitmax) do
    q = A p
    α =  $\frac{\rho}{p^T q}$ 
    z := z + α p
    r := r - α q
    ρ0 = ρ
    ρ := rT r
    β =  $\frac{\rho}{\rho_0}$ 
    p := r + β p
end
(z, ||x - A z||)
  
```

Comparison: NAS NPB 2.3 Serial Code

```

do j=1,naa+1
  q(j) = 0.0d0
  z(j) = 0.0d0
  r(j) = x(j)
  p(j) = r(j)
  w(j) = 0.0d0
enddo
sum = 0.0d0
do j=1,lastcol-firstcol+1
  sum = sum + r(j)*r(j)
enddo
rho = sum
do cgit = 1,cgitmax
  do j=1,lastrow-firstrow+1
    sum = 0.d0
    do k=rowstr(j),rowstr(j+1)-1
      sum = sum + a(k)*p(colidx(k))
    enddo
    w(j) = sum
  enddo
  do j=1,lastcol-firstcol+1
    q(j) = w(j)
  enddo
enddo

```

```

do j=1,lastcol-firstcol+1
  w(j) = 0.0d0
enddo
sum = 0.0d0
do j=1,lastcol-firstcol+1
  sum = sum + p(j)*q(j)
enddo
d = sum
alpha = rho / d
rho0 = rho
do j=1,lastcol-firstcol+1
  z(j) = z(j) + alpha*p(j)
  r(j) = r(j) - alpha*q(j)
enddo
sum = 0.0d0
do j=1,lastcol-firstcol+1
  sum = sum + r(j)*r(j)
enddo
rho = sum
beta = rho / rho0
do j=1,lastcol-firstcol+1
  p(j) = r(j) + beta*p(j)
enddo
enddo

```

```

do j=1,lastrow-firstrow+1
  sum = 0.d0
  do k=rowstr(j),rowstr(j+1)-1
    sum = sum + a(k)*z(colidx(k))
  enddo
  w(j) = sum
enddo
do j=1,lastcol-firstcol+1
  r(j) = w(j)
enddo
sum = 0.0d0
do j=1,lastcol-firstcol+1
  d = x(j) - r(j)
  sum = sum + d*d
enddo
d = sum
rnorm = sqrt( d )

```

Should Parallelism Be the Default?

- “Loop” can be a misleading term
 - > A set of executions of a parameterized block of code
 - > Whether to order or parallelize those executions should be a separate question
- Fortress “loops” are parallel by default
 - > This is actually a library convention about generators
 - > You get sequential execution by asking for it specifically

In Fortress, Parallelism Is the Default

```
for i←1:m, j←1:n do
  a[i,j] := b[i] c[j]
end
```

1:n is a generator

```
for i←seq(1:m) do
  for j←seq(1:n) do
    print a[i,j]
  end
end
```

seq(1:n) is a sequential generator

```
for i←1:m, j←i:n do
  a[i,j] := b[i] c[j]
end
```

a.indices is a generator for the indices of the array **a**

a.indices.rowMajor is a sequential generator of indices

```
for (i,j)←a.indices do a[i,j] := b[i] c[j] end
```

```
for (i,j)←a.indices.rowMajor do print a[i,j] end
```

- Generators (defined by libraries) manage parallelism and the assignment of threads to processors

Generators

- Aggregates
 - > Lists $\langle 1, 2, 4, 3, 4 \rangle$ and vectors $[1\ 2\ 4\ 3\ 4]$
 - > Sets $\{1, 2, 3, 4\}$ and multisets $\{|1, 2, 3, 4, 4|\}$
 - > Arrays (including multidimensional)
- Ranges $1:10$ and $1:99:2$ and $5\#20$
- Index sets $a.indices$ and $a.indices.rowMajor$
- Reducers (also defined by libraries) such as Σ (or **SUM**) and **MAX** may have serial/parallel implementations
- Reducers are driven by generators
- Distribution of generator guides parallelism of reducer

Loops, Reducers, Comprehensions

for $k \leftarrow 1:n$ **do** *print k* **end**

$$y = \sum_{k \leftarrow 1:n} a_k x^k$$

$$w = \sum S \quad (* \text{ same as } \sum_{x \leftarrow S} x *)$$

$$v = \bigcap_{\substack{k \leftarrow S \\ \text{prime } k}} \text{arrayOfSets}_k$$

$$z = \text{MAX}_{(j,k) \leftarrow a.\text{indices}} |a_{j,k} - b_{j,k}|$$

$$B = \{ f(x, y) \mid x \leftarrow S, y \leftarrow A, x \neq y \}$$

$$l_{\text{triangle}} = \left\langle \frac{x(x+1)}{2} \mid x \leftarrow 1:100 \right\rangle$$

Loops, Reducers, Comprehensions

for $k \leftarrow 1:n$ do print i end

$y = \Sigma[\mathbf{k \leftarrow 1:n}] \mathbf{a[k]} \mathbf{x^k}$

$w = \Sigma S$ (* same as $\Sigma[\mathbf{x \leftarrow S}] \mathbf{x}$ *)

$v = \cap[\mathbf{k \leftarrow S, \text{prime } k}] \mathbf{arrayOfSets[k]}$

$z = \mathbf{MAX}[\mathbf{(j,k) \leftarrow a.indices}] \mathbf{|a[j,k]-b[j,k]|}$

$B = \{ \mathbf{f(x,y)} \mid \mathbf{x \leftarrow S, y \leftarrow A, x \neq y} \}$

$\mathbf{1_triangle} = \langle \mathbf{x(x+1)/2} \mid \mathbf{x \leftarrow 1:100} \rangle$

Parallelism Is Not a Feature!


- Parallel programming is not a goal, but a pragmatic compromise.
- It would be a lot easier to program a single processor chip running at 1 PHz than a million processors running at 20 GHz.
 - > We don't know how to build a 1 PHz processor.
 - > Even if we did, someone would still want to strap a bunch of them together!
- Parallel programming is difficult and error-prone.

Research Questions

Can we encapsulate parallelism in libraries?


Will this separation be effective?

Abstract Collections




Regions

- Hierarchical data structure describes CPU and memory resources and their properties
 - > Allocation heaps
 - > Parallelism
 - > Memory coherence
- A running thread can find out its resources
- Threads may be explicitly spawned in specified regions



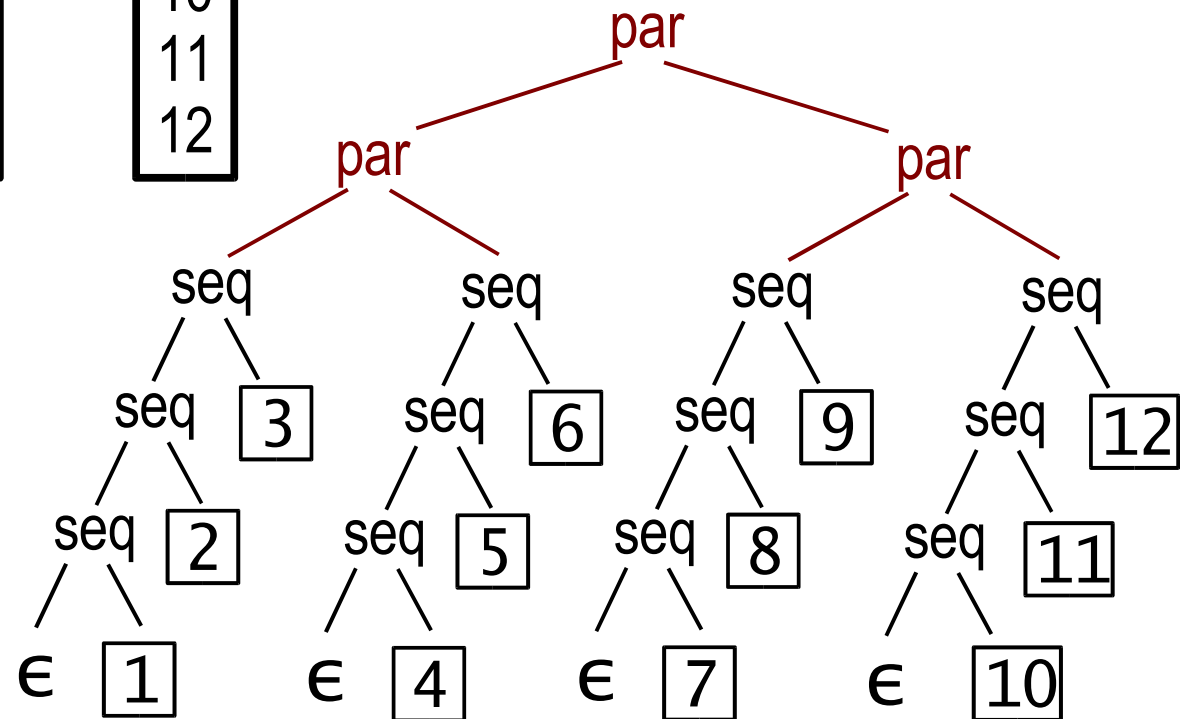

Distributions

- Describe how to map a data structure onto a region
 - > Block, cyclic, block-cyclic, Morton order ...
 - > Map an array into a chip? Use a local heap.
 - > Map an array onto a cluster? Break it up.



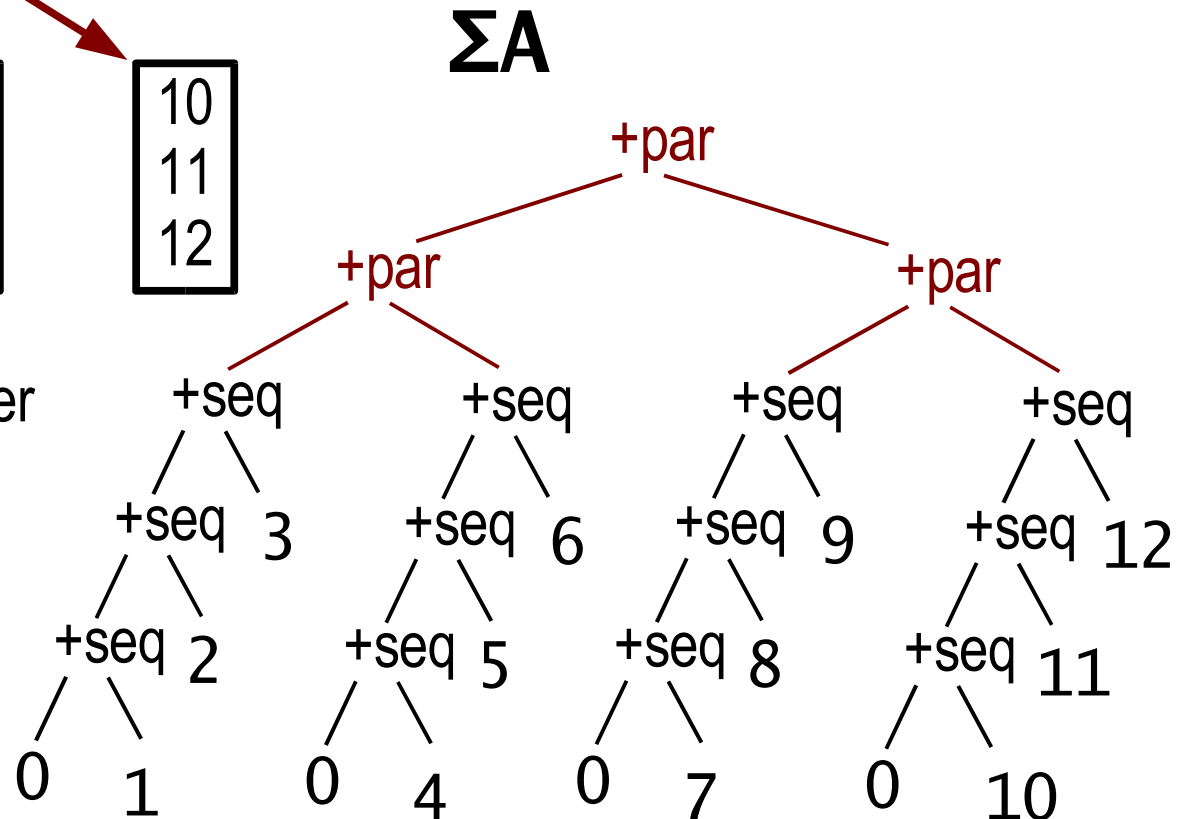

- Defined entirely by libraries!
 - > User-extensible

Generators Drive Parallelism



When a data structure (or its index set) is used as a generator, the parallelism of the generator reflects the distribution of the data structure.

Generators Modify Reducers: Parallelism




This works because integer addition is associative.

The generator knows this

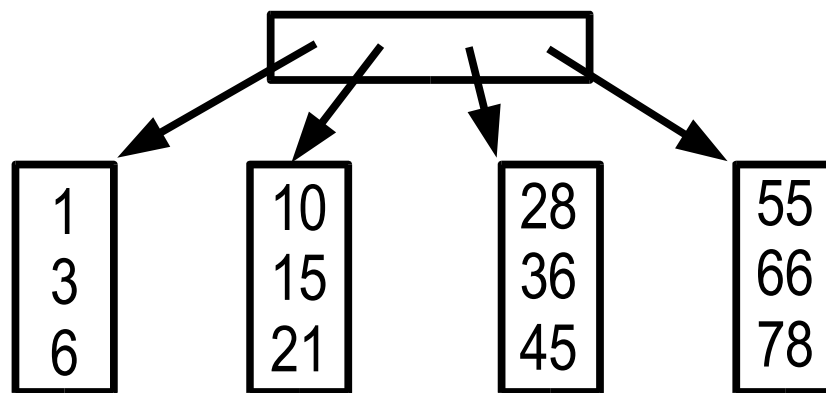
because the trait \mathbb{Z} extends the trait

Associative $[\mathbb{Z}, +]$.

Generators Modify Reducers: Distribution



$$[\ x(x+1)/2 \ | \ x \leftarrow A \]$$



Generators and reducers may agree to use a specialized protocol that, for example, communicates array shapes and distribution information.

Summary: Parallelism in Fortress

- Regions describe machine resources.
- Distributions map aggregates onto regions.
- Aggregates used as generators drive parallelism.
- Algebraic properties drive implementation strategies.
- Algebraic properties are described by traits.
- Properties are verified by automated unit testing.
- Traits allow sharing of code, properties, and test data.
- Reducers and generators negotiate through overloaded method dispatch keyed by traits to achieve mix-and-match parallel code selection.

Our Key Design Themes

- Make stupid mistakes impossible
 - And make clever mistakes relatively unlikely
- Design the language to be grown by (expert) users
 - Rich library language enables simple application languages
- Make abstraction efficient
 - Aggressive static and dynamic optimization
- Make parallelism tractable
 - Appropriate abstractions for managing thread and data distribution
- Emulate standard mathematical notation
 - Reduce the effort of translating from science to computation

guy.steele@sun.com



<http://research.sun.com/projects/plrg>



Carl Eastlund, Guy Steele, Jan-Willem Maessen, Yossi Lev, Eric Allen, Joe Hallett, Sukyoung Ryu, Sam Tobin-Hochstadt, David Chase, João Dias