The Future Is Parallel:
What's a Programmer to Do?

Breaking Sequential Habits of Thought

Guy Steele

Sun Fellow

Sun Microsystems Laboratories

April 2009

With Multicore, a Profound Shift

- Parallelism is here, now, and in our faces
  - > Academics have been studying it for 50 years
  - > Serious commercial offerings for 25 years
  - > But now it's in desktops and laptops
- Specialized expertise for science codes and databases and networking
- But soon general practitioners must go parallel
- An opportunity to make parallelism easier for everyone

This Talk Is about Performance

The bag of programming tricks

that has served us so well

for the last 50 years

is

the wrong way to think

going forward and

must be thrown out.

Why?

- Good sequential code minimizes total number of operations.
  - > Clever tricks to reuse previously computed results.
  - > Good parallel code often performs redundant operations to reduce communication.
- Good sequential algorithms minimize space usage.
  - > Clever tricks to reuse storage.
  - > Good parallel code often requires extra space to permit temporal decoupling.
- Sequential idioms stress linear problem decomposition.
  - > Process one thing at a time and accumulate results.
  - > Good parallel code usually requires multiway problem decomposition and multiway aggregation of results.

Let's Add a Bunch of Numbers

```
DO I = 1, 1000000
   SUM = SUM + X(I)
END DO
```

Can it be parallelized?

Let's Add a Bunch of Numbers

```
SUM = 0                // Oops!

DO I = 1, 1000000
  SUM = SUM + X(I)
END DO
```

Can it be parallelized?

This is already bad!
Clever compilers have to undo this.

What Does a Mathematician Say?

$$\sum_{i=1}^{1000000} x_i \quad \text{or maybe just} \quad \sum x$$

Compare Fortran 90 `SUM(X)`.

What, not how.

No commitment yet as to strategy. This is good.

Sequential Computation Tree

```
SUM = 0
DO I = 1, 1000000
  SUM = SUM + X(I)
END DO
```

Atomic Update Computation Tree (a)

```
SUM = 0
PARALLEL DO I = 1, 1000000
  SUM = SUM + X(I)
END DO
```

Atomic Update Computation Tree (b)

```
SUM = 0
PARALLEL DO I = 1, 1000000
  ATOMIC SUM = SUM + X(I)
END DO
```

Parallel Computation Tree

What sort of code
should we write
to get a computation
tree of this shape?

What sort of code
would we *like*
to write?

Finding the Length of a LISP List

first   rest

Recursive:

```
(define length (list)
  (cond ((null list) 0)
        (else (+ 1 (length (rest list)))))))
```

Append Lists (1)

Append Lists (2)

Append Lists (3)
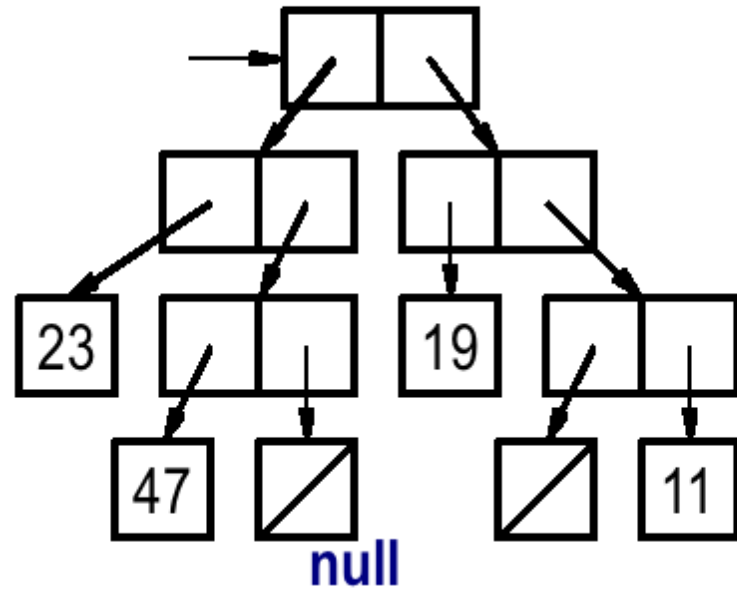
Append Lists (4)

Defining Lists Using `car cdr cons`

```scheme
(define (first x)
  (cond ((null? x) '())
        (else (car x))))

(define (rest x)
  (cond ((null? x) '())
        (else (cdr x))))

(define (append xs ys)
  (cond ((null? xs) ys)
        (else (cons (car xs) (append (cdr xs) ys)))))

(define (addleft a xs) (cons a xs))

(define (addright xs a)
  (cond ((null? xs) (list a))
        (else (cons (car xs) (addright (cdr xs) a)))))
```
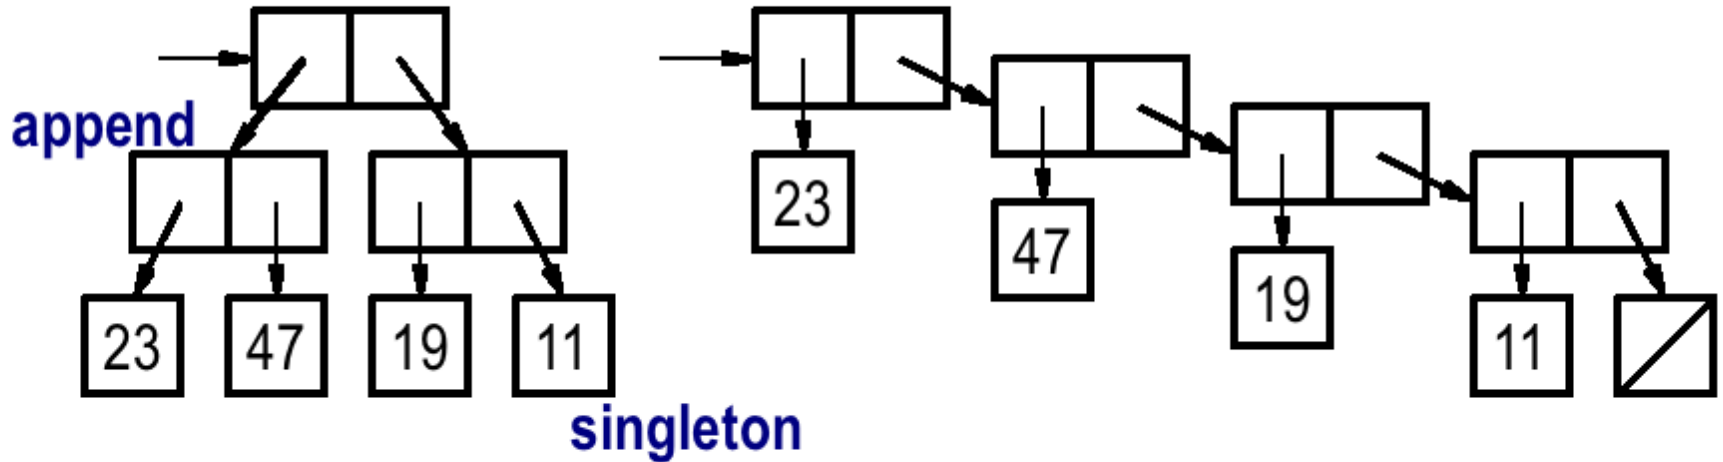
Defining Lists Using `car cdr cons`

```
(define (first x)                       ;Constant time
  (cond ((null? x) '())
        (else (car x))))

(define (rest x)                        ;Constant time
  (cond ((null? x) '())
        (else (cdr x))))

(define (append xs ys)                  ;Linear in (length xs)
  (cond ((null? xs) ys)
        (else (cons (car xs) (append (cdr xs) ys)))))

(define (addleft a xs) (cons a xs))     ;Constant time

(define (addright xs a)                 ;Linear in (length xs)
  (cond ((null? xs) (list a))
        (else (cons (car xs) (addright (cdr xs) a)))))
```

map reduce mapreduce

```
(map (λ (x) (* x x)) '(1 2 3))  =>  (1 4 9)

(reduce + 0 '(1 4 9))  =>  14

(mapreduce (λ (x) (* x x)) + 0 '(1 2 3))  =>  14


(define (map f xs)                    ;Linear in (length xs)
  (cond ((null? xs) '())
        (else (cons (f (car xs)) (map f (cdr xs))))))

(define (reduce g id xs)              ;Linear in (length xs)
  (cond ((null? xs) id)
        (else (g (car xs) (reduce g id (cdr xs))))))

(define (mapreduce f g id xs)         ;Linear in (length xs)
  (cond ((null? xs) id)
        (else (g (f (car xs)) (mapreduce f g id (cdr xs))))))
```

length filter

```
(define (length xs)                    ;Linear in (length xs)
  (mapreduce (λ (q) 1) + 0 xs))

(define (filter p xs)                  ;Linear in (length xs)
  (cond ((null? xs) '())
        ((p (car xs)) (cons p (filter p (cdr xs))))
        (else (filter p (cdr x)))))

(define (filter p xs)                  ;Linear in (length xs)??
  (apply append
         (map (λ (x) (if (p x) (list x) '())) xs)))

(define (filter p xs)                  ;Linear in (length xs)!!
  (mapreduce (λ (x) (if (p x) (list x) '()))
             append '() xs))
```

reverse

```scheme
(define (reverse xs)                            ;QUADRATIC in (length xs)
  (cond ((null? xs) '())
        (else (addright (reverse (cdr xs)) (car xs)))))

(define (revappend xs ys)                       ;Linear in (length xs)
  (cond ((null? xs) ys)
        (else (revappend (cdr xs) (cons (car xs) ys)))))

(define (reverse xs)                            ;Linear in (length xs)
  (revappend xs '()))
```

Linear versus Multiway Decomposition

- Linearly linked lists are inherently sequential.
  - > Compare Peano arithmetic: $5 = (((0+1)+1)+1)+1)+1$
  - > Binary arithmetic is much more efficient than unary!
- We need a *multiway decomposition* paradigm:

  ```
  length [ ] = 0
  length [a] = 1
  length (a++b) = (length a) + (length b)
  ```

  This is just a summation problem: adding up a bunch of 1's!
  (More generally: a bunch of 0's and 1's.)

Defining Lists Using item list split conc (1)

```
(define (first xs)                           ;Depth of left path
  (cond ((null? xs) '())
        ((singleton? xs) (item xs))
        (else (split xs (λ (ys zs) (first ys))))))

(define (rest xs)                            ;Depth of left path
  (cond ((null? xs) '())
        ((singleton? xs) '())
        (else (split xs (λ (ys zs) (append (rest ys) zs))))))

(define (append xs ys)                       ;Constant time
  (cond ((null? xs) ys)
        ((null? ys) xs)
        (else (conc xs ys))))
```

Defining Lists Using item list split conc (2)

```
(define (first xs)                     ;Depth of left path
  (cond ((null? xs) '())
        ((singleton? xs) (item xs))
        (else (split xs (λ (ys zs) (first ys))))))

(define (rest xs)                      ;Depth of left path
  (cond ((null? xs) '())
        ((singleton? xs) '())
        (else (split xs (λ (ys zs) (append (rest ys) zs))))))

(define (append xs ys)                 ;???
  (cond ((null? xs) ys)
        ((null? ys) xs)
        (else (rebalance (conc xs ys)))))
```

Defining Lists Using
`item list split conc` (3)

```
(define (addleft a xs)
  (cond ((null? xs) (list a))
        ((singleton? xs) (append (list a) xs))
        (else (split xs (λ (ys zs) (append (addleft a ys) zs))))))

(define (addright xs a)
  (cond ((null? xs) (list a))
        ((singleton? xs) (append xs (list a)))
        (else (split xs (λ (ys zs) (append ys (addright a zs)))))))


(define (addleft a xs) (append (list a) xs))

(define (addright xs a) (append xs (list a)))
```

Parallel `map reduce mapreduce`

```
(define (mapreduce f g id xs)           ;Logarithmic in (length xs)??
  (cond ((null? xs) id)
        ((singleton? xs) (f (item xs)))
        (else (split xs (λ (ys zs)
                  (g (mapreduce f g id ys)
                     (mapreduce f g id zs)))))))

(define (map f xs)
  (mapreduce (λ (x) (list (f x))) append '() xs))

(define (reduce g id xs)
  (mapreduce (λ (x) x) g id xs))
```

length filter reverse

```
(define (length xs)                     ;Logarithmic in (length xs)??
  (mapreduce (λ (q) 1) + 0 xs))

(define (filter p xs)                   ;Logarithmic in (length xs)??
  (mapreduce (λ (x) (if (p x) (list x) '()))
             append '() xs))

(define (reverse xs)                    ;Logarithmic in (length xs)??
  (mapreduce list (λ (yx zs) (append zs ys)) '() xs))
```

Exercise: Write Mergesort and Quicksort in This Binary-split Style

- Mergesort: structural induction on input
  - > Cheaply split input in half
  - > Recursively sort the two halves
  - > Carefully merge the two sorted sublists (tricky)
- Quicksort: structural induction on output
  - > Carefully split input into lower and upper halves (tricky)
  - > Recursively sort the two halves
  - > Cheaply append the two sorted sublists

A Modest Example: Filter (1)

$$sequentialFilter[\![E]\!](a \colon \mathrm{List}[\![E]\!], p \colon E \to \mathrm{Boolean}) \colon \mathrm{List}[\![E]\!] = \texttt{do}$$

$$result \colon \mathrm{List}[\![E]\!] := \langle\,\rangle$$

$$\texttt{for } k \leftarrow seq(0 \mathbin{\#} a.size()) \texttt{ do}$$

$$\texttt{if } p(a_k) \texttt{ then } result := result.addRight(a_k) \texttt{ end}$$

$$\texttt{end}$$

$$result$$

$$\texttt{end}$$

Example of use:

$$odd(x \colon \mathbb{Z}) = ((x \texttt{ MOD } 2) \neq 0)$$

$$sequentialFilter(\langle\, 1, 4, 7, 2, 5 \,\rangle, odd)$$

So what language is this? **Fortress**.

A Modest Example: Filter (2)

$recursiveFilter[\![E]\!](a\colon\mathrm{List}[\![E]\!], p\colon E \to \mathrm{Boolean})\colon\mathrm{List}[\![E]\!] =$

    `if` $a.isEmpty()$ `then` $\langle\,\rangle$

    `else`

      $(first, rest) = a.extractLeft().get()$

      $rest' = recursiveFilter(rest, p)$

      `if` $p(first)$ `then` $rest'.addLeft(first)$ `else` $rest'$ `end`

    `end`

A Modest Example: Filter (3)

$parallelFilter[\![E]\!](a\!: \text{List}[\![E]\!], p\!: E \rightarrow \text{Boolean})\!: \text{List}[\![E]\!] =$

    `if` $|a| = 0$ `then` $\langle\,\rangle$

    `elif` $|a| = 1$ `then`

      $(\mathit{first}, \_) = a.extractLeft().get()$

      `if` $p(\mathit{first})$ `then` a `else` $\langle\,\rangle$ `end`

    `else`

      $(x, y) = a.split()$

      $parallelFilter(x, p) \parallel parallelFilter(y, p)$

    `end`

A Modest Example: Filter (4)

$$reductionFilter[\![E]\!](a\colon \mathrm{List}[\![E]\!], p\colon E \to \mathrm{Boolean})\colon \mathrm{List}[\![E]\!] =$$
$$\underset{x \leftarrow a}{\|} \; (\texttt{if } p(x) \texttt{ then } \langle\, x\,\rangle \texttt{ else } \langle\,\rangle \texttt{ end})$$

$$\langle\, x \mid x \leftarrow a, p(x) \,\rangle$$

Oh, yes: $\displaystyle\sum_{i \leftarrow 1:1000000} x_i$

Splitting a String into Words (1)

- Given: a string
- Result: List of strings, the words separated by spaces
  - > Words must be nonempty
  - > Words may be separated by more than one space
  - > String may or may not begin (or end) with spaces

Splitting a String into Words (2)

- Tests:

  $println\ words($"`This is a sample`"$)$

  $println\ words($"` Here  is  another  sample `"$)$

  $println\ words($"`JustOneWord`"$)$

  $println\ words($"`  `"$)$

  $println\ words($""$)$

- Expected output:

  $\langle$ `This`, `is`, `a`, `sample` $\rangle$

  $\langle$ `Here`, `is`, `another`, `sample` $\rangle$

  $\langle$ `JustOneWord` $\rangle$

  $\langle\ \rangle$

  $\langle\ \rangle$

Splitting a String into Words (3)

$$words(s\colon \text{String}) = \texttt{do}$$

$\qquad result\colon \text{List}[\![\text{String}]\!] := \langle\,\rangle$

$\qquad word\colon \text{String} := \text{``''}$

$\qquad \texttt{for } k \leftarrow seq(0\,\#\,length(s)) \texttt{ do}$

$\qquad\qquad char = substring(s, k, k+1)$

$\qquad\qquad \texttt{if } (char = \text{`` ''}) \texttt{ then}$

$\qquad\qquad\qquad \texttt{if } (word \neq \text{``''}) \texttt{ then } result := result \,\|\, \langle\, word\,\rangle \texttt{ end}$

$\qquad\qquad\qquad word := \text{``''}$

$\qquad\qquad \texttt{else}$

$\qquad\qquad\qquad word := word \,\|\, char$

$\qquad\qquad \texttt{end}$

$\qquad \texttt{end}$

$\qquad \texttt{if } (word \neq \text{``''}) \texttt{ then } result := result \,\|\, \langle\, word\,\rangle \texttt{ end}$

$\qquad result$

$\texttt{end}$

Splitting a String into Words (4a)

Here is a sesquipedalian string of words

Here is a sesquipeda|lian string of words

Here is a |sesquipeda|lian strin|g of words

Splitting a String into Words (4b)

Here is a sesquipedalian string of words

Chunk("sesquipeda")

Splitting a String into Words (4c)

Here is a sesquipedalian string of words

Segment("g", ⟨"of"⟩, "words")

Splitting a String into Words (4d)

Here is a sesquipedalian string of words

Segment("Here", ⟨"is", "a"⟩, "")

Splitting a String into Words (4e)

Here is a sesquipeda|lian| |strin|g of words

Segment("lian", ⟨ ⟩, "strin")

Splitting a String into Words (4f)

Here is a sesquipeda

Here is a sesquipeda

Segment("Here", ⟨"is", "a"⟩, "")   Chunk("sesquipeda")

Segment("Here", ⟨"is", "a"⟩, "sesquipeda")

Splitting a String into Words (4g)

lian | string | of | words

lian | string | of | words

Segment("lian", ⟨ ⟩, "strin")

Segment("g", ⟨"of"⟩, "words")

Segment("lian", ⟨"string", "of"⟩, "words")

Splitting a String into Words (4h)

Here is a sesquipedalian string of words

Segment(
  "Here",
  ⟨"is", "a", "sesquipedialian", "string", "of"⟩,
  "words")

Splitting a String into Words (5)

$maybeWord(s : \text{String}) : \text{List}[\![\text{String}]\!] =$
  `if` $s = $ "" `then` $\langle\,\rangle$ `else` $\langle\, s \,\rangle$ `end`

`trait` WordState
      `extends` $\{\,\text{Associative}[\![\text{WordState}, \oplus]\!]\,\}$
      `comprises` $\{\,\text{Chunk}, \text{Segment}\,\}$
  `opr` $\oplus(\texttt{self}, other : \text{WordState}) : \text{WordState}$
`end`

Splitting a String into Words (6)

object Chunk(s:String) extends WordState
   opr $\oplus$(self, $other$:Chunk):WordState =
     Chunk($s \parallel other.s$)
   opr $\oplus$(self, $other$:Segment):WordState =
     Segment($s \parallel other.l, other.A, other.r$)
end

Splitting a String into Words (7)

$$\texttt{object} \ \mathrm{Segment}(l\colon \mathrm{String}, A\colon \mathrm{List}[\![\mathrm{String}]\!], r\colon \mathrm{String})$$
$$\texttt{extends} \ \mathrm{WordState}$$
$$\texttt{opr} \ \oplus(\texttt{self}, \mathit{other}\colon \mathrm{Chunk})\colon \mathrm{WordState} =$$
$$\mathrm{Segment}(l, A, r \parallel \mathit{other}.s)$$
$$\texttt{opr} \ \oplus(\texttt{self}, \mathit{other}\colon \mathrm{Segment})\colon \mathrm{WordState} =$$
$$\mathrm{Segment}(l, A \parallel \mathit{maybeWord}(r \parallel \mathit{other}.l) \parallel \mathit{other}.A, \mathit{other}.r)$$
$$\texttt{end}$$

Splitting a String into Words (8)

$processChar(c\colon \text{String})\colon \text{WordState} =$
   `if` $(c = \text{`` ''})$ `then` $\text{Segment}(\text{``''}, \langle\rangle, \text{``''})$
   `else` $\text{Chunk}(c)$
   `end`

$words(s\colon \text{String}) =$ `do`
   $g = \bigoplus_{k \leftarrow 0 \# length(s)} processChar(substring(s, k, k+1))$
   `typecase` g `of`
     $\text{Chunk} \Rightarrow maybeWord(g.s)$
     $\text{Segment} \Rightarrow maybeWord(g.l) \parallel g.A \parallel maybeWord(g.r)$
   `end`
`end`

Splitting a String into Words (9)

$(\ast$ The mechanics of BIG OPLUS $\ast)$

opr BIG $\oplus$ $[\![T]\!](g:(\text{Reduction}[\![\text{WordState}]\!],$
$$T \rightarrow \text{WordState})$$
$$\rightarrow \text{WordState}): \text{WordState} =$$
$g(\text{GlomReduction}, identity[\![\text{WordState}]\!])$

object GlomReduction extends $\text{Reduction}[\![\text{WordState}]\!]$
  getter $toString() =$ "GlomReduction"
  $empty(): \text{WordState} = \text{Chunk}(``")$
  $join(a: \text{WordState}, b: \text{WordState}): \text{WordState} = a \oplus b$
end

What's Going On Here?

Instead of linear induction

with one base case (empty),

we have multiway induction

with two base cases (empty and unit).

Why are these two base cases important?

Representation of Abstract Collections

Binary operator $\lozenge$

Leaf operator ("unit") $\square$

Optional empty collection ("zero") ε

that is the identity for $\lozenge$

Algebraic Properties of ◇

| Associative	Commutative	Idempotent	
no	no	no	binary trees
no	no	yes	weird
no	yes	no	"mobiles"
no	yes	yes	weird
yes	no	no	lists (arrays)
yes	no	yes	weird
yes	yes	no	multisets (bags)
yes	yes	yes	sets

The "Boom hierarchy"

# Associativity

These are all considered
to be equivalent.

52
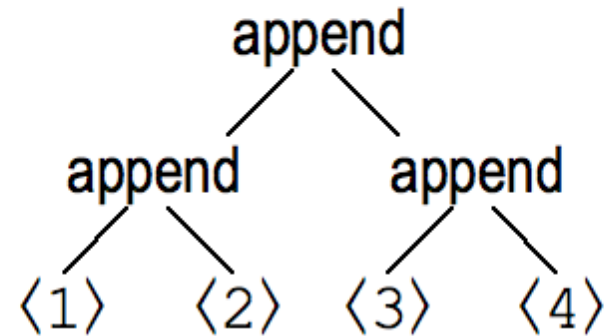
Catamorphism: Summation

Replace $\diamond$ $\square$ ε with $+$ identity 0

Computation: Summation

Replace $\diamond$ $\square$ ε with $+$ identity 0

Catamorphism: Lists

Replace $\diamondsuit \quad \square \quad \varepsilon$ with append $\langle - \rangle \quad \langle \rangle$

Computation: Lists

Replace ◊ □ ε with append ⟨–⟩ ⟨⟩

Representation: Lists

Replace $\Diamond$ $\square$ ε with append $\langle - \rangle$ $\langle \rangle$

Catamorphism: Loops

Replace $\lozenge$ $\square$ ε with seq identity () or par identity ()
where seq: (),() $\rightarrow$ () and par: (),() $\rightarrow$ ()

for $i \leftarrow seq(1:4)$ do $print\ i$ end
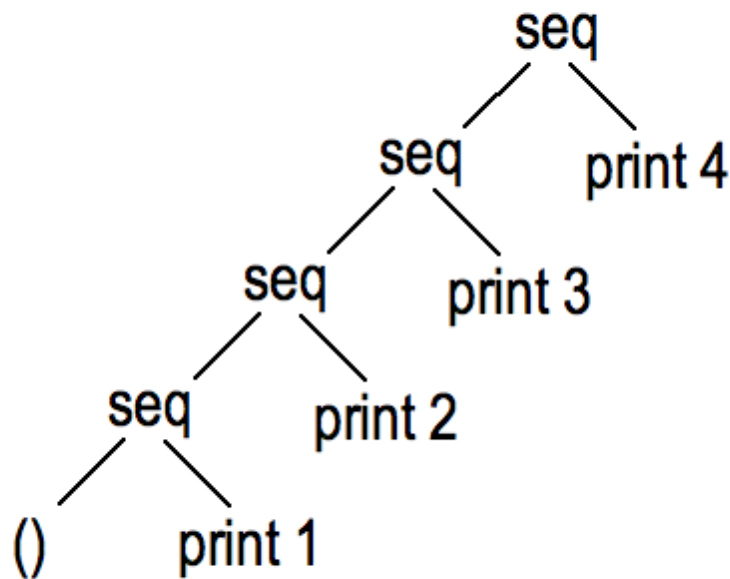
for $i \leftarrow 1:4$ do $print\ i$ end

Generators can modify the catamorphism
and so control the parallelism. 58

To Summarize: A Big Idea

- Loops and summations and list constructors are alike!

$$\texttt{for } i \leftarrow 1:1000000 \texttt{ do } x_i := x_i^2 \texttt{ end}$$

$$\sum_{i \leftarrow 1:1000000} x_i^2$$

$$\langle\, x_i^2 \mid i \leftarrow 1:1000000 \,\rangle$$

  - > Generate an abstract collection
  - > The *body* computes a function of each item
  - > Combine the results (or just synchronize)

- Whether to be sequential or parallel is a separable question
  - > That's why they are especially good abstractions!
  - > Make the decision on the fly, to use available resources

Another Big Idea

- Formulate a sequential loop as successive applications of state transformation functions f_i

- Find an *efficient* way to compute and represent compositions of such functions (this step requires ingenuity)

- Instead of computing

  $$s := s_0; \texttt{for } i \leftarrow seq(1:1000000) \texttt{ do } s := f_i(s) \texttt{ end},$$

  compute $s := \left(\underset{i \leftarrow 1:1000000}{\circ} f_i \right) s_0$

- Because function composition is associative, the latter has a parallel strategy

- In the "words in a string" problem, each character can be regarded as defining a state transformation function

Splitting a String into Words (3, again)

$$words(s\text{:}\,\mathrm{String}) = \texttt{do}$$

$result\text{:}\,\mathrm{List}[\![\mathrm{String}]\!] := \langle\,\rangle$

$word\text{:}\,\mathrm{String} := \text{``''}$

$\texttt{for}\ k \leftarrow seq(0\,\#\,length(s))\ \texttt{do}$

$char = substring(s, k, k + 1)$

$\texttt{if}\ (char = \text{`` ''})\ \texttt{then}$

$\texttt{if}\ (word \neq \text{``''})\ \texttt{then}\ result := result \,\|\, \langle\,word\,\rangle\ \texttt{end}$

$word := \text{``''}$

$\texttt{else}$

$word := word \,\|\, char$

$\texttt{end}$

$\texttt{end}$

$\texttt{if}\ (word \neq \text{``''})\ \texttt{then}\ result := result \,\|\, \langle\,word\,\rangle\ \texttt{end}$

$result$

$\texttt{end}$

61

MapReduce Is a Big Deal!

- Associative combining operators are a VERY BIG DEAL!
  - \> Google MapReduce requires that combining operators also be commutative.
  - \> There are ways around that.
- Inventing new combining operators is a very, very big deal.
  - \> Creative catamorphisms!
  - \> We need programming languages that encourage this.
  - \> We need assistance in proving them associative.

We Need a New Mindset

- DO loops are so 1950s! (Literally: Fortran is now 50 years old.)

- So are linear linked lists! (Literally: Lisp is now 50 years old.)

- Java™-style iterators are **so** last millennium!

- Even arrays are suspect!

- As soon as you say "first, `SUM = 0`" you are hosed. Accumulators are BAD.

- If you say, "process subproblems in order," you lose.

- The great tricks of the sequential past DON'T WORK.

- The programming idioms that have become second nature to us as everyday tools DON'T WORK.

The Parallel Future

- We need new strategies for problem decomposition.
  - \> Data structure design/object relationships
  - \> Algorithmic organization
  - \> Don't split a problem into "the first" and "the rest."
  - \> Do split a problem into roughly equal pieces.
    Then figure out how to combine general subsolutions.
  - \> Often this makes combining the results a bit harder.
- We need programming languages and runtime
  implementations that support parallel strategies
  and hybrid sequential/parallel strategies.
- We must learn to manage new space-time tradeoffs.

Conclusion

- A program organized according to linear problem decomposition principles can be really hard to parallelize.

- A program organized according to parallel problem decomposition principles is easily run either in parallel or sequentially, according to available resources.

- The new strategy has costs and overheads. They will be reduced over time but will not disappear.

- In a world of parallel computers of wildly varying sizes, this is our only hope for program portability in the future.

- Better language design can encourage better parallel programming.

guy.steele@sun.com
http://research.sun.com/projects/plrg
http://projectfortress.sun.com