

Language Improvements

Chapel Team, Cray Inc.

Chapel version 1.9 summary

April 17th, 2014 (released) / May 2014 (documented)



COMPUTE | STORE | ANALYZE

Safe Harbor Statement

The Cray logo is located in the top right corner of the slide. It consists of the word "CRAY" in a blue, sans-serif font, followed by a small graphic of a cluster of dots in various colors (red, blue, green, yellow) arranged in a hexagonal pattern.

This presentation may contain forward-looking statements that are based on our current expectations. Forward looking statements may include statements about our financial guidance and expected operating results, our opportunities and future potential, our product development and new product introduction plans, our ability to expand and penetrate our addressable markets and other statements that are not historical facts. These statements are only predictions and actual results may materially vary from those projected. Please refer to Cray's documents filed with the SEC from time to time concerning factors that could affect the Company and these forward-looking statements.



COMPUTE | STORE | ANALYZE

Copyright 2014 Cray Inc.

2

Outline



- Updated Assignment Signature
- 'noinit': Squashing default initialization
- 'ref' intents for 'this'
- Dynamic casts of 'nil'
- Break out of 'param' for-loops
- Fixing Domain Literals of Domain Values
- Expression-free Serial Statements



COMPUTE | STORE | ANALYZE

Copyright 2014 Cray Inc.

3

Updated Assignment Signature

COMPUTE | STORE | ANALYZE

Copyright 2014 Cray Inc.

4

Updated Assignment Signature

The Cray logo is located in the top right corner of the slide. It consists of the word "CRAY" in a blue, sans-serif font, followed by a decorative graphic of a cluster of small, colored circles in shades of blue, red, and yellow.

Old-style assignment:

- Took two arguments, but returned the RHS:

```
proc =(a: int(?w), b: int(w)) return b;
```
- Rationale:
 - Historically, Chapel did not have 'ref' intents
 - Therefore, couldn't modify 'a' without incurring copies via '[in]out' intent
- Compiler automatically wrapped calls to assignment in a MOVE
 - Necessary to copy the result back into the LHS arg:

```
lhs = rhs;
```


was translated into:

```
('move' lhs ("=" lhs rhs))
```
- Problems:
 - Requires creating a temporary result of the same type as the lhs.
 - Then requires internal tricks to avoid a verbatim copy-back.
 - "Cognitive dissonance" with other languages
 - also with 'op=' assignments which had previously switched to using 'ref' intents
 - Larger internal representation



COMPUTE | STORE | ANALYZE

Copyright 2014 Cray Inc.

5

Updated Assignment Signature

CRAY

New-style assignment:

- Now, assignment takes LHS by reference:

```
proc =(ref a: int(?w), b: int(w)) {  
    __primitive("=", a, b); // implement via a primitive ("C-level") assignment  
}  
  
proc =(ref a: myR, b: myR) {  
    a.x = b.x;  
    a.y = b.y;  
}
```
- Eliminates the need for the "wrapper MOVE".
- Easier to understand (no magic).
- Uses same signature as other *op=* assignments.

Status (as of v1.9 release):

- Assignment for all internal types now uses the new signature
- A different solution needed for default assignment of extern types



COMPUTE | STORE | ANALYZE

Copyright 2014 Cray Inc.

6

- The `__primitive("=", ...)` call is an implementation-level hook for getting access to "primitive" (below Chapel-level) capabilities. In this case it corresponds to a C-level assignment or `memcpy()` or potentially a remote `put/get`. In subsequent slides, it's referred to as `PRIM_ASSIGN` (the compiler-internal name for it).
- The "MOVE" referenced on this slide (and the previous) is another Chapel primitive (`PRIM_MOVE`). It plays a similar role as `PRIM_ASSIGN`, but is considered a bit bloated / dated at present, so the approach being taken with this work is to introduce `PRIM_ASSIGN`, do it right/well, and then slowly phase out all references to `PRIM_MOVE`.

Updated Assignment Signature



Additional Work Completed Since 1.9 Release:

- Provided default assignment for extern types
- Removed catch-all assignment definition
- Removed compiler insertion of wrapper MOVES
- Upgraded PRIM_ASSIGN to handle operands of wide class types

Next Steps:

- Port data motion optimizations to PRIM_ASSIGN
- Eliminate primitive MOVE entirely



COMPUTE | STORE | ANALYZE

Copyright 2014 Cray Inc.

7



'noinit': Squashing default initialization

COMPUTE | STORE | ANALYZE

Copyright 2014 Cray Inc.

8

'noinit': Motivation and Background



- **Default initialization occurs on all declared variables**
 - But in some cases this is fairly expensive – for instance, arrays
 - For a variable written before it's read, default initialization is overhead
 - Compiler can automatically optimize many simple cases
 - Yet arrays can easily thwart such analysis
- **'noinit' permits users to specify “no default initialization”**



COMPUTE | STORE | ANALYZE

Copyright 2014 Cray Inc.



- Arrays present a challenge due to indexing. For example, given “ $A[0..n/2] = 1$; $A[m..n] = 2$;”, is A initialized safely enough for the compiler to squash the default initialization?
- Limits of interprocedural analysis/alias analysis are other limiting factors

'noinit': Implementation Today



- **'noinit' is specified as follows:**

```
var foo: int = noinit;
```

- **When present, default initialization is skipped**

- current support limited to simple types

- **supports cleaner compilation**

- also sets the stage for further compiler clean-up in the future



COMPUTE | STORE | ANALYZE

Copyright 2014 Cray Inc.

10

- At present, this only occurs for primitive types such as integers and strings; enums; tuples; and ranges
- Certain temporary variables necessary for default initialization are not inserted into the compiler when this keyword is specified, requiring less clean up in simple cases

'noinit': Limitations (by design)



- **types, consts, params do not accept 'noinit'**
 - rationale: by definition, can't be assigned at a later point
- **no static/dynamic def-before-use safety checks for 'noinit'**
 - users read unassigned 'noinit' variables at their own risk
 - could add conservative compiler-based warnings in the future
- **formal arguments do not accept 'noinit'**
 - could reconsider this for 'out'/'in'/'inout' arguments if desirable
 - the resulting error message is not particularly clear today



COMPUTE | STORE | ANALYZE

Copyright 2014 Cray Inc.

11

There is also no dynamic user query to see whether a 'noinit' variable has been initialized or not

'noinit': Next Steps

CRAY

- **Implement the remaining types**
 - records and classes
 - domains and arrays
 - syncs, singles, and atomics
- **Support Chapel-level specification of default initialization**
 - Move current default initialization support from compiler to modules
 - Potentially give end-users a hook for specifying default values
 - Simplify compiler code related to initialization as a result
- **Improve error message for 'noinit' arguments**



COMPUTE | STORE | ANALYZE

Copyright 2014 Cray Inc.

12

- Generally speaking, for the remaining types, the work involves separating the mechanisms used to allocate space from those used to initialize it – currently, these are entangled.
- This work also sets us up for better “first-touch” and parallel initialization policies for arrays in a way that will better match subsequent iterations than the current very naïve scheme.



'ref' intents for 'this'

COMPUTE | STORE | ANALYZE

Copyright 2014 Cray Inc.

13

Allow 'this' to be marked with a 'ref' intent

CRAY

Background: Allow modifications of 'this' for value types

- Previously, this was possible for developers via 'pragma "ref this"'

```
pragma "ref this"
proc int.triple() {
    this *= 3;
}
var x = 2;
x.triple();    // x == 6 after this call
```

- We also had language support for marking 'this' with a **param** intent

```
proc param string.length param {
    ...
}
"hello there!".length    // results in 12 at compile time
```



COMPUTE | STORE | ANALYZE

Copyright 2014 Cray Inc.

14

Allow 'this' to be marked with a 'ref' intent



This Effort: Bring 'pragma "ref this"' into the language

- We now reuse the syntax for 'param this' to allow for **ref** or **param**

```
proc ref int.triple() {  
    this *= 3;  
}  
  
var x = 2;  
x.triple();    // x == 6 after this call
```

Impact: Reduced complexity within the compiler

- Removed another pragma and FLAG_REF_THIS
- The implicit 'this' argument is now marked with the proper intent
 - This was the source of several special cases



COMPUTE | STORE | ANALYZE

Copyright 2014 Cray Inc.

15

Allow 'this' to be marked with a 'ref' intent

CRAY

Next Steps:

- Make this feature illegal for reference types
 - Currently there are no restrictions for marking 'this' as ref

```
class Foo { ... }  
proc ref Foo.change() {  
    this = new Foo();    //??  
}
```
 - In the future this should be an error
- Consider adding other intent types to 'this' as well
 - Including more formally defining what a 'blank' this intent implies?



COMPUTE | STORE | ANALYZE

Copyright 2014 Cray Inc.

16



Dynamic casts of 'nil'

COMPUTE | STORE | ANALYZE

Copyright 2014 Cray Inc.

17

Dynamic casts of 'nil'

The Cray logo is located in the top right corner of the slide. It consists of the word "CRAY" in a blue, sans-serif font, followed by a decorative graphic of a hexagonal grid with some cells highlighted in red, blue, and yellow.

Background: Dynamic casts of 'nil' caused a segfault

```
class Node { ... }  
class SpecialNode : Node { ... }  
proc walk(tree: Node) {  
    walk(tree.left);  
    walk(tree.right);  
    var x = tree: SpecialNode;    // segfaults if tree == nil  
    ...  
}
```

- This aspect of the language was never defined

Improvement: Dynamic casts of 'nil' always produce 'nil'

- Could be considered a cast failure, which also returns 'nil'

```
...  
var x = tree: SpecialNode;    // x == nil or x == valid SpecialNode  
...
```



COMPUTE | STORE | ANALYZE

Copyright 2014 Cray Inc.

18



Break out of 'param' for-loops

COMPUTE | STORE | ANALYZE

Copyright 2014 Cray Inc.

19

Break out of 'param' for-loops

CRAY

Background:

- The break statement can be used to exit a loop early

```
for i in 1..10 do
  if shouldExitEarly(i) then break; else f(i);
```
- 'param' for-loops are fully unrolled so each index is a 'param' value

```
for param i in 1..4 do
  f(i);
```

// this loop is equivalent to: f(1); f(2); f(3); f(4);
- But break has not been supported in 'param' for-loops

This Effort:

- Support break to allow param for loops to exit early:

```
for param i in 1..10 do
  if shouldExitEarly(i) then break; else f(i);
```

Impact:

- 'param' for-loops now treated more uniformly with other loops



COMPUTE | STORE | ANALYZE

Copyright 2014 Cray Inc.

20

Fixing Domain Literals of Domain Values

COMPUTE | STORE | ANALYZE

Copyright 2014 Cray Inc.

21

Fixing Domain Literals of Domain Values

CRAY

Background: For domain D , we have interpreted ' $\{D\}$ ' as ' D '

- for example, the following two loops have been equivalent:

```
forall i in D ...           forall i in {D} ...
```
- yet, $\{D\}$ syntactically specifies an associative domain with one index, D
- reasons for this are historical

This Effort: Treat $\{D\}$ as an associative domain of domains

```
forall i in {D} do ...
```

 // now, *i* would take on the single index value, D
(noting, however, that domains of domains are not currently supported)

Impact: More consistent language syntax

- required rewrites of some tests that relied on previous syntax

Next Steps: Add support for domains of domains

- Need to implement a hash function for domains
- In v1.9 using $\{D\}$ generates a warning to help codes transition



COMPUTE | STORE | ANALYZE

Copyright 2014 Cray Inc.

22

- Prior to this work, both the loops ' $\text{forall } i \text{ in } D$ ' and ' $\text{forall } i \text{ in } \{D\}$ ' would have caused ' i ' to take on the indices of ' D '.
- This change sets us up to treat the former the same, and the latter as an iteration over an associative domain of domains.
- The historical reasons for this relate to the fact that we used to use square brackets for domain literals, combined with some holdover from ZPL design decisions
- We're not aware of any deep reason that associative domains of domains are not currently supported – simply that nobody has gotten to them / needed them yet. Writing a hash function for domains is the first step (and possibly the only one).



Expression-free Serial Statements

COMPUTE | STORE | ANALYZE

Copyright 2014 Cray Inc.

23

Expression-free Serial Statements

CRAY

Background: Chapel's serial statements squash parallelism

- serial statements take a boolean condition

```
serial (here.runningTasks() > here.numCores) do
  forall i in D do ...
```
- to statically squash parallelism, we've traditionally used:

```
serial true do ...
```

This Effort: can now write serial statements without conditions

```
serial do                                // equivalent to 'serial true do'
  forall i in D do ...
```

Impact: used to squash nested parallelism (e.g., spectral-norm):

```
forall i in Dav do
  serial do Av[i] = + reduce [j in Dv] (A[i,j] * v[j]);
```

Next Steps: explore other methods for squashing parallelism

- e.g., compiler heuristics, domain maps, serial expressions, sugar, ...



COMPUTE | STORE | ANALYZE

Copyright 2014 Cray Inc.

24

It's worth noting that while some Chapel constructs, like forall loops, are fairly easily to statically serialize (by changing to a for loop, for example), cases like reductions and promoted operations across arrays are more cumbersome to do (requiring converting either case to an explicit loop rather than relying on the more direct expression), and so in these cases 'serial true' tends to be more useful.

Legal Disclaimer

The Cray logo is located in the top right corner of the page. It consists of the word "CRAY" in a bold, blue, sans-serif font. To the right of the text is a decorative graphic of a network or molecular structure, composed of small circles and lines in various colors (blue, green, red, yellow).

Information in this document is provided in connection with Cray Inc. products. No license, express or implied, to any intellectual property rights is granted by this document.

Cray Inc. may make changes to specifications and product descriptions at any time, without notice.

All products, dates and figures specified are preliminary based on current expectations, and are subject to change without notice.

Cray hardware and software products may contain design defects or errors known as errata, which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Cray uses codenames internally to identify products that are in development and not yet publically announced for release. Customers and other third parties are not authorized by Cray Inc. to use codenames in advertising, promotion or marketing and any use of Cray Inc. internal codenames is at the sole risk of the user.

Performance tests and ratings are measured using specific systems and/or components and reflect the approximate performance of Cray Inc. products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance.

The following are trademarks of Cray Inc. and are registered in the United States and other countries: CRAY and design, SONEXION, URIKA, and YARCDATA. The following are trademarks of Cray Inc.: ACE, APPRENTICE2, CHAPEL, CLUSTER CONNECT, CRAYPAT, CRAYPORT, ECOPHLEX, LIBSCI, NODEKARE, THREADSTORM. The following system family marks, and associated model number marks, are trademarks of Cray Inc.: CS, CX, XC, XE, XK, XMT, and XT. The registered trademark LINUX is used pursuant to a sublicense from LMI, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis. Other trademarks used in this document are the property of their respective owners.

Copyright 2014 Cray Inc.



COMPUTE | STORE | ANALYZE

Copyright 2014 Cray Inc.

25



<http://chapel.cray.com>

chapel_info@cray.com

<http://sourceforge.net/projects/chapel/>