

# Separation Logic for High-Level Synthesis

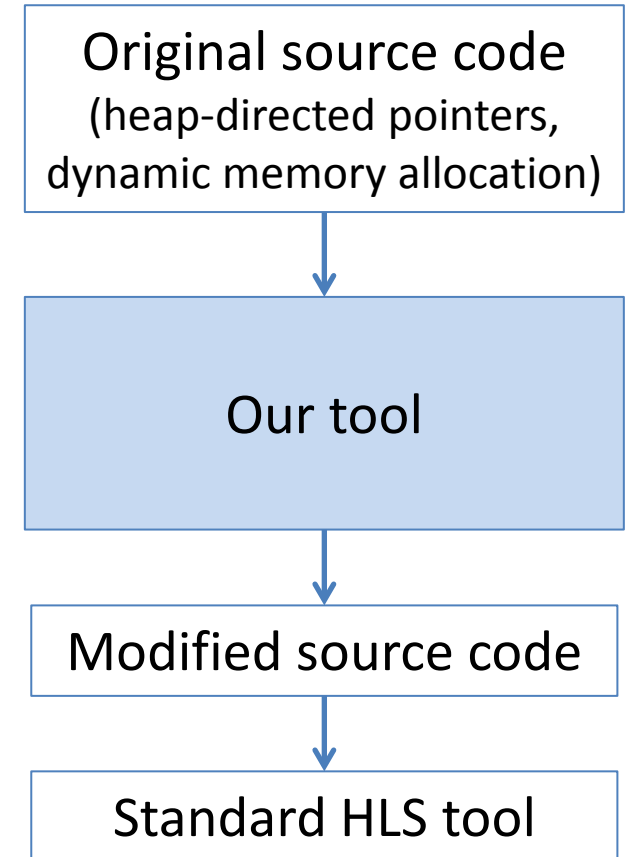
Felix Winterstein

20 March 2015


[f.winterstein12@imperial.ac.uk](mailto:f.winterstein12@imperial.ac.uk)

## Executive summary

- C-to-FPGA tools require manual source code refactoring...
  - ... to map pointer-manipulating programs efficiently into HW
- Static program analysis
  - Analyse pointer-based memory accesses and heap layout
  - Identify disjoint, independent regions in heap memory
- Source-to-source transformations
  - Partition heap across on-chip memory banks
  - Automatic loop parallelization



## Field programmable gate array



On-chip  
memory  
blocks

Programmable  
memory  
aggregation

Programmable  
look-up table (LUT):  
Implements an  
arbitrary boolean  
function

Programmable  
interconnects  
between LUTs, and  
memory blocks  
(and other  
components)

## Dedicated digital circuit

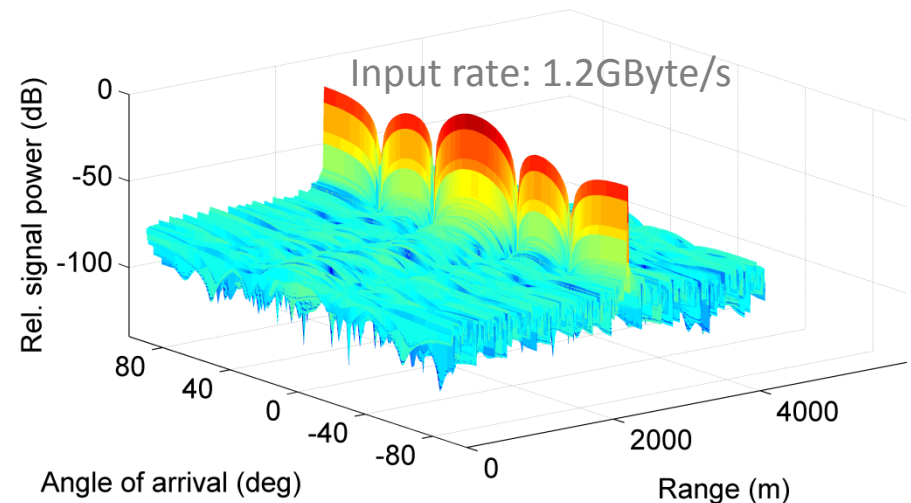
- Fast
- Energy efficient
- Need it where software is too slow

## Reconfigurable

- As flexible as software (advantage over ASICs)

Example:

Phased Array Radar: Front-end signal processing + beamforming + ranging in **real-time**



## “Software”

```
1: int sum = 0;
2: for (i=0; i<N; i++)
3:     sum += a[i];
```

## Hardware description language (e.g. VHDL)

```
1.  G0: for I in 0 to N-1 generate
2.      input_array(I) <= input_string(((I+1)*IN_BITWIDTH-1 downto I*IN_BITWIDTH));
3.  end generate G0;
4.  G2: for I in 0 to LAYERS_TREE_ADDER-1 generate
5.      G_TOP_LAYER: if I = 0 generate
6.          G2_2: for J in 0 to (N/(2**(I+1)))-1 generate
7.              addorsub_inst_2 : addorsub
8.                  generic map (
9.                      USE_DSP => USE_DSP_FOR_ADD,
10.                     A_BITWIDTH => IN_BITWIDTH,
11.                     B_BITWIDTH => IN_BITWIDTH,
12.                     RES_BITWIDTH => INT_BITWIDTH
13.                 )
14.             port map (
15.                 clk => clk,
16.                 sclr => sclr,
17.                 nd => '1',
18.                 sub => sub,
19.                 a => input_array(2*J),
20.                 b => input_array(2*J+1),
21.                 res => tmp_tree_adder_res(I,J),
22.                 rdy => open
23.             );
24.         end generate G2_2;
```



```

25.     G2_3: if N/(2**(I+1)) < integer(ceil(real(N)/real(2**(I+1)))) generate
26.         data_delay_array_proc : process(clk)
27.         begin
28.             if rising_edge(clk) then
29.                 data_delay_array(I)(0)(IN_BITWIDTH-1 downto 0) <= input_array(N/(2**(I+1))*
30.                 data_delay_array(I)(0)(INT_BITWIDTH-1 downto IN_BITWIDTH) <= (others => in
31.                 data_delay_array(I)(1 to SINGLE_ADDER_LAT-1) <= data_delay_array(I)(0 to SIN
32.             end if;
33.         end process data_delay_array_proc;
34.         tmp_tree_adder_res(I,N/(2**(I+1))) <= data_delay_array(I)(SINGLE_ADDER_LAT-1);
35.     end generate G2_3;
36. end generate G_TOP_LAYER;
37. G_OTHER_LAYER: if I > 0 AND I < LAYERS_TREE_ADDER-1 generate
38.     G2_2: for J in 0 to ((N+N-(N/(2**I))*(2**I))/(2**(I+1))-1) generate
39.         addorsub_inst_2 : addorsub
40.         generic map (
41.             USE_DSP => USE_DSP_FOR_ADD,
42.             A_BITWIDTH => INT_BITWIDTH,
43.             B_BITWIDTH => INT_BITWIDTH,
44.             RES_BITWIDTH => INT_BITWIDTH
45.         )
46.         port map (
47.             clk => clk,
48.             sclr => sclr,
49.             nd => '1',

```

```

50.         sub => sub,
51.         a => tmp_tree_adder_res(I-1,2*J),
52.         b => tmp_tree_adder_res(I-1,2*J+1),
53.         res => tmp_tree_adder_res(I,J),
54.         rdy => open
55.     );
56. end generate G2_2;
57. G2_3: if ((N+N/(2**I))*(2**I))/(2**(I+1)) < integer(ceil(real(N)/real(2**(I+1)))) genera
58.     data_delay_array_proc : process(clk)
59.     begin
60.         if rising_edge(clk) then
61.             data_delay_array(I)(0) <= tmp_tree_adder_res(I-1,N/(2**(I+1))*2);
62.             data_delay_array(I)(1 to SINGLE_ADDER_LAT-1) <= data_delay_array(I)(0 to SING
63.         end if;
64.     end process data_delay_array_proc;
65.     tmp_tree_adder_res(I,N/(2**(I+1))) <= data_delay_array(I)(SINGLE_ADDER_LAT-1);
66. end generate G2_3;
67. end generate G_OTHER_LAYER;
68. G_BOTTOM_LAYER: if I = LAYERS_TREE_ADDER-1 AND LAYERS_TREE_ADDER > 1 genera
69.     G2_2: for J in 0 to 0 generate
70.         addorsub_inst_2 : addorsub
71.         generic map (
72.             USE_DSP => USE_DSP_FOR_ADD,
73.             A_BITWIDTH => INT_BITWIDTH,
74.             B_BITWIDTH => INT_BITWIDTH,
75.             RES_BITWIDTH => INT_BITWIDTH

```

```
1: int sum = 0;
2: for (i=0; i<N; i++)
3:   sum += a[i];
```




## High-level synthesis

```

76.         )
77.         port map (
78.             clk => clk,
79.             sclr => sclr,
80.             nd => '1',
81.             sub => sub,
82.             a => tmp_tree_adder_res(I-1,2*J),
83.             b => tmp_tree_adder_res(I-1,2*J+1),
84.             res => tmp_tree_adder_res(I,J),
85.             rdy => open
86.         );
87.     end generate G2_2;
88. end generate G_BOTTOM_LAYER;
89. end generate G2;
90. ctrl_delay_line_proc : process(clk)
91. begin
92.     if rising_edge(clk) then
93.         if sclr = '1' then
94.             ctrl_delay_line <= (others => '0');
95.         else
96.             ctrl_delay_line(0) <= nd;
97.             ctrl_delay_line(1 to TREE_ADDER_LAT-1) <= ctrl_delay_line(0 to TREE_ADDER_LAT-
98.         end if;
99.     end if;
100. end process ctrl_delay_line_proc;
101. rdy <= ctrl_delay_line(TREE_ADDER_LAT-1);
102. output <= tmp_tree_adder_res(LAYERS_TREE_ADDER-1,0);

```

102



Tool	Input language
Cadence C-to-Silicon	C
Synopsys Symphony C Compiler	C
Mentor Graphics Catapult C	C
Impulse CoDeveloper	C
Xilinx Vivado HLS	C
Bluespec	BSV
National Instruments LabVIEW FPGA	LabVIEW schematic
Xilinx System Generator for DSP	Matlab/Simulink
DEFACTO	C
ROCCC	C
LegUP	C
Chisel	Scala


Tool	Input language
<ul style="list-style-type: none"><li>• State-of-the-art HLS tools don't support full featured C/C++ code</li><li>• A major restriction: Heap directed pointers and dynamic memory allocation not supported</li><li>• <b>Worth considering at all?</b></li></ul>	
Chisel	Scala


- **Case study: High-level synthesis of dynamic data structures**
- Challenge
- Motivating example
- Leveraging separation logic
- Implementation and results
- Outlook


## Case study:

- Compare computational properties of two algorithms for *K*-means clustering
- SW (C++) / RTL (VHDL) / HLS (C++) implementations
- Code available on GitHub (Vivado-KMeans)

# K-means clustering



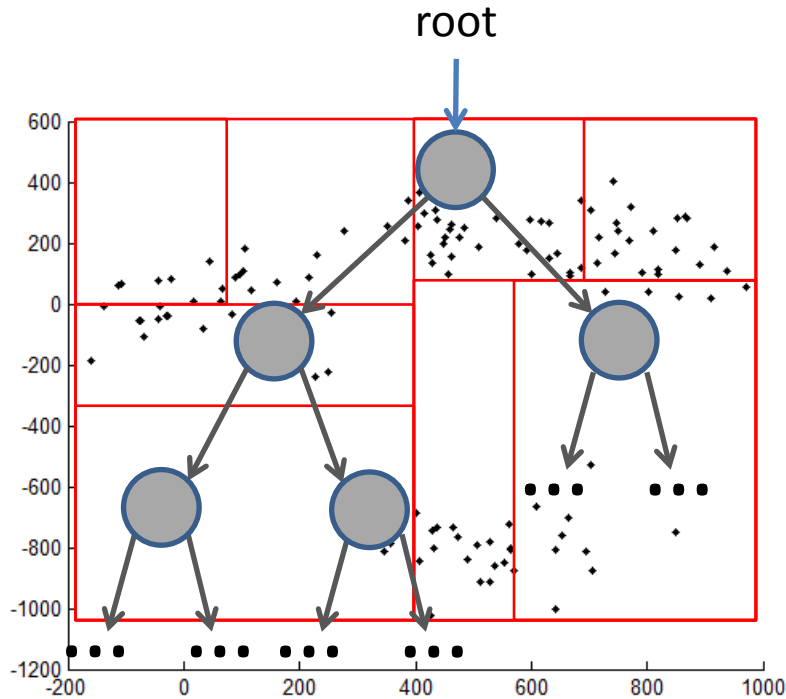




```


for all  $x_i \in$  data points do
  for all  $z_j \in$  centers do
    compute distance  $\|x_i - z_j\|^2$ 
    pick & update closest center
  end for
end for
  
```

- For each data point ...
- search among  $K$  candidates for the closest center




- Recursively split data set
- Build a pointer-linked tree data structure

\* "The Filtering Algorithm",  
Kanungo et al., 2002




- Recursively split data set
- Build a pointer-linked tree data structure
- Clustering: Tree traversal
- Acceleration through search space pruning
- Dynamically (de-)allocate memory to store intermediate results

\* "The Filtering Algorithm",  
Kanungo et al., 2002



## Brute-force algorithm

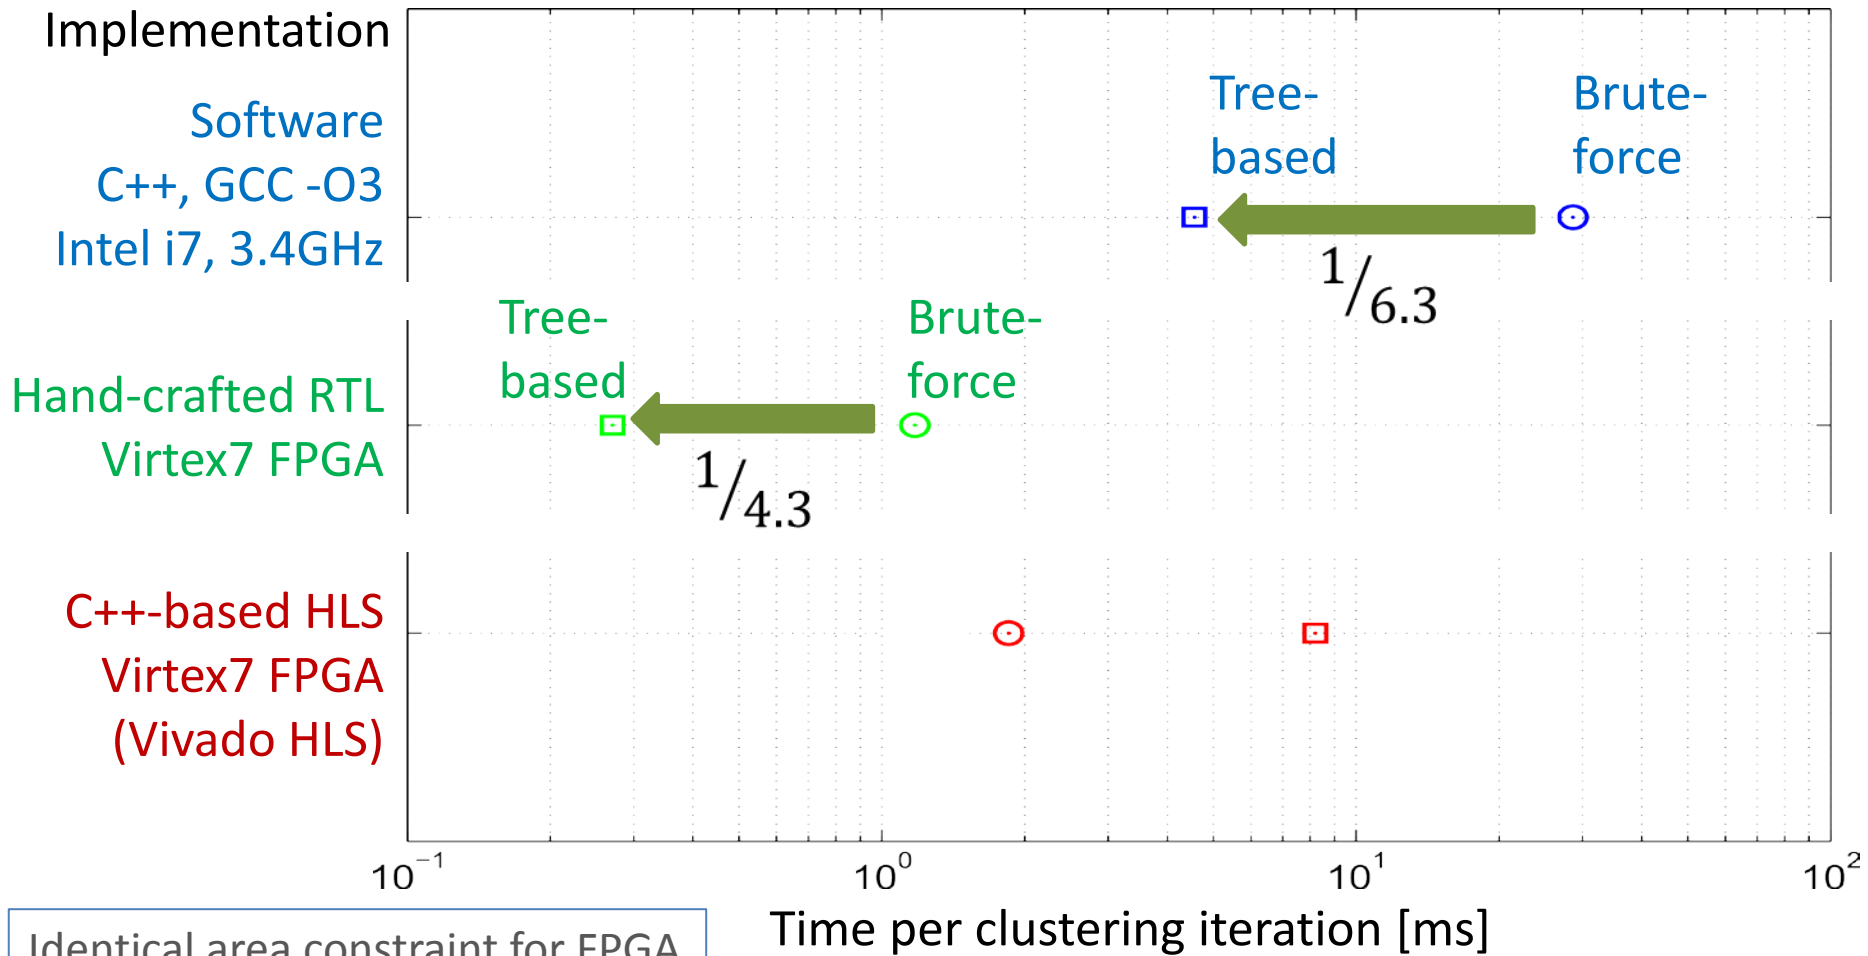
- Computationally expensive
- Simple control flow
- Embarrassingly parallel



## Tree-based algorithm


- Data-dependent control flow
- Pointer-based tree traversal
- Dynamic memory allocation

## The battlefield



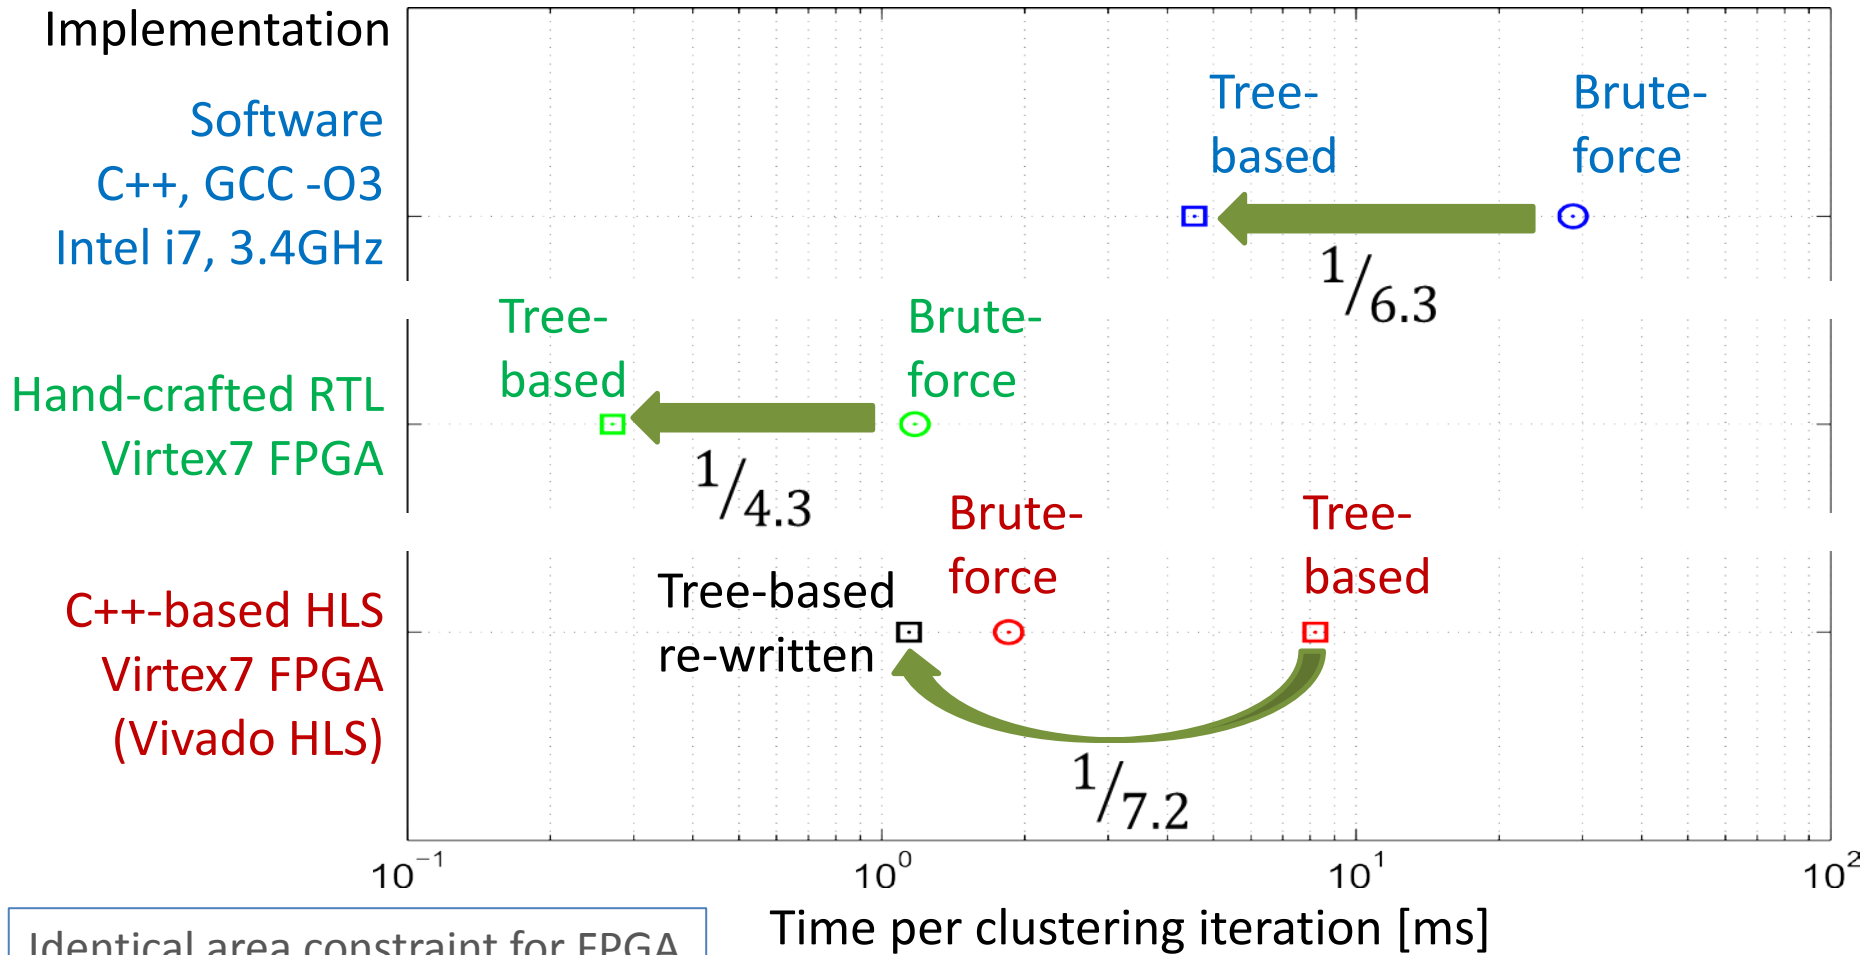
Identical area constraint for FPGA implementations: 6500 slices

## The battlefield



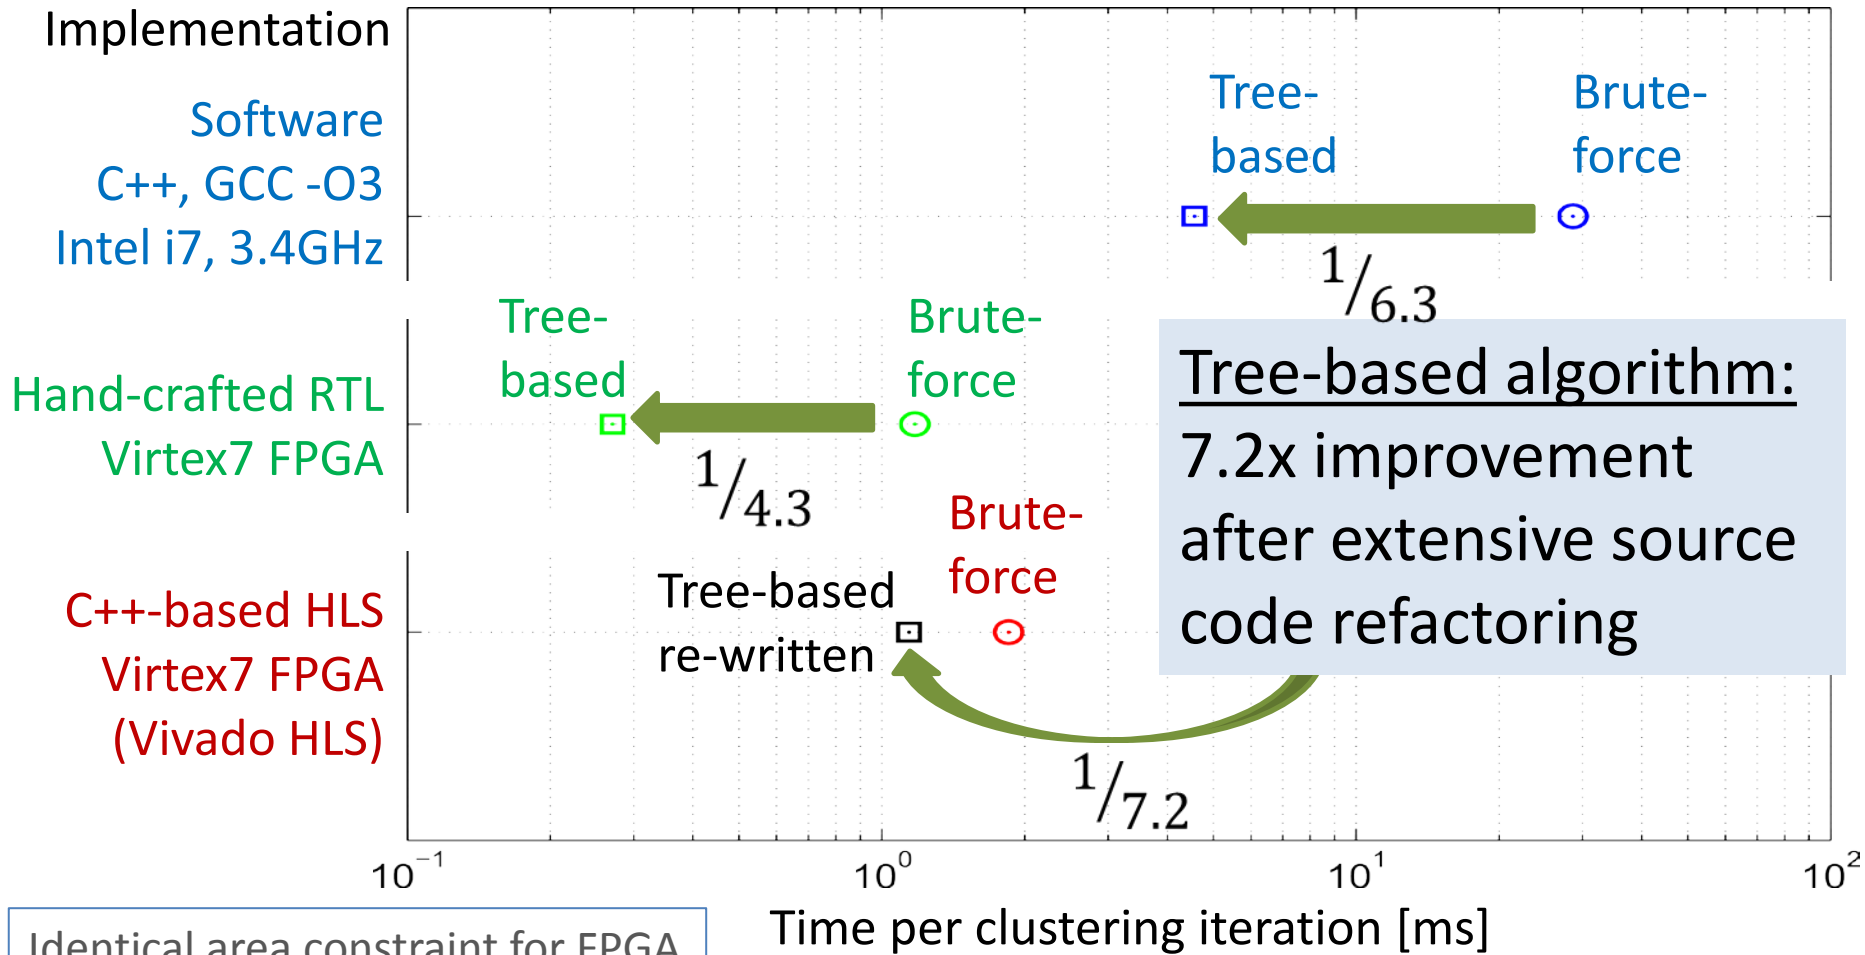
Identical area constraint for FPGA implementations: 6500 slices

## The battlefield




Identical area constraint for FPGA implementations: 6500 slices

## The battlefield




Tree-based algorithm:  
7.2x improvement  
after extensive source  
code refactoring

Identical area constraint for FPGA implementations: 6500 slices



## Brute-force algorithm

- Computationally expensive
- Simple control flow
- Embarrassingly parallel
- **Seamless C-to-FPGA implementation**




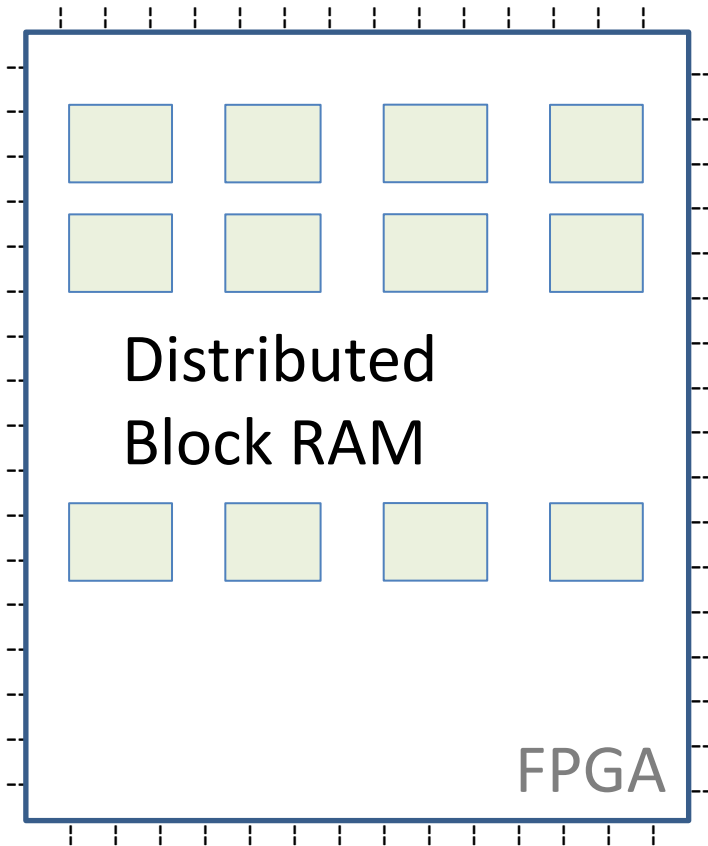
## Tree-based algorithm

- Data-dependent control flow
- Pointer-based tree traversal
- Dynamic memory allocation
- **C-to-FPGA requires substantial code modifications**

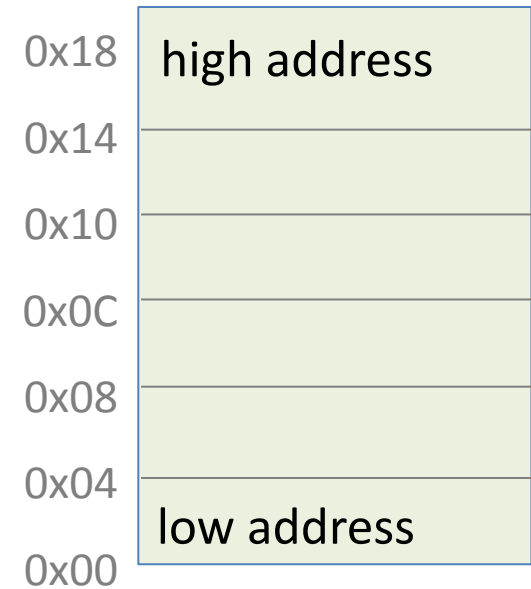
## Automate

- Memory partitioning
- Parallelization
- Custom implementation of dynamic memory allocation
- Loop flattening
- Loop distribution
- Custom bit widths
- ...

- Case study: High-level synthesis of dynamic data structures
- **Challenge**
- Motivating example
- Leveraging separation logic
- Implementation and results
- Outlook



## SW memory model




```
int main() {  
    x = A[i];  
    p = new int;  
    *p = 3;  
    ...  
}
```

Lack of automated optimizations ...

- ... for programs using pointers
- ... because pointers are difficult to analyze
- ... and memory is allocated, disposed, and reused at run-time
- Yet widely used in SW

**This work takes a step towards closing this gap**


SW memory model




```
int main() {
    x = A[i];
    p = new int;
    *p = 3;
    ...
}
```

## Our goal

- Partition heap-allocated data structures ('heaplets')
- Synthesize a parallel implementation



## SW memory model




- Ensure that heap partitions are 'private'

- Case study: High-level synthesis of dynamic data structures
- Challenge
- **Motivating example**
- Leveraging separation logic
- Implementation and results
- Outlook


Can we parallelize  
this loop?

```
s = PUSH(root, s);  
while s!=0 do  
    s = POP(&u, s);  
    ... do something  
    if (u->left!= 0) && (u->right!=0) then  
        s = PUSH(u->right, s);  
        s = PUSH(u->left, s);  
    end if  
    delete u;  
end while
```



```

s = PUSH(root, s);
while s!=0 do
  s = POP(&u, s);
  ... do something
  if (u->left!= 0) && (u->right!=0) then
    s = PUSH(u->right, s);
    s = PUSH(u->left, s);
  end if
  delete u;
end while
    
```



```
s = PUSH(root, s);
```

```
while s!=0 do
```

```
    s = POP(&u, s);
```

```
    ... do something
```

```
    if (u->left!= 0) && (u->right!=0) then
```


```
        s = PUSH(u->right, s);
```

```
        s = PUSH(u->left, s);
```

```
    end if
```


```
    delete u;
```

```
end while
```



```


s = PUSH(root, s);
while s!=0 do
    s = POP(&u, s);
    ... do something
    if (u->left!= 0) && (u->right!=0) then
        s = PUSH(u->right, s);
        s = PUSH(u->left, s);
    end if
    delete u;
end while
    
```



```


s = PUSH(root, s);
while s!=0 do
  s = POP(&u, s);
  ... do something
  if (u->left!= 0) && (u->right!=0) then
    s = PUSH(u->right, s);
    s = PUSH(u->left, s);
  end if
  delete u;
end while

```




```

s = PUSH(root, s);
while s!=0 do
    s = POP(&u, s);
    ... do something
    if (u->left!= 0) && (u->right!=0) then
        s = PUSH(u->right, s);
        s = PUSH(u->left, s);
    end if
    delete u;
end while
    
```




```

s = PUSH(root, s);
while s!=0 do
    s = POP(&u, s);
    ... do something
    if (u->left!= 0) && (u->right!=0) then
        s = PUSH(u->right, s);
        s = PUSH(u->left, s);
    end if
    delete u;
end while
    
```




```

s = PUSH(root, s);
while s!=0 do
    s = POP(&u, s);
    ... do something
    if (u->left!= 0) && (u->right!=0) then
        s = PUSH(u->right, s);
        s = PUSH(u->left, s);
    end if
    delete u;
end while
    
```




```

s = PUSH(root, s);
while s!=0 do
    s = POP(&u, s);
    ... do something
    if (u->left!= 0) && (u->right!=0) then
        s = PUSH(u->right, s);
        s = PUSH(u->left, s);
    end if
    delete u;
end while
    
```



```


s = PUSH(root, s);
while s!=0 do
    s = POP(&u, s);
    ... do something
    if (u->left!= 0) && (u->right!=0) then
        s = PUSH(u->right, s);
        s = PUSH(u->left, s);
    end if
    delete u;
end while
    
```



```


s = PUSH(root, s);
while s!=0 do
  s = POP(&u, s);
  ... do something
  if (u->left!= 0) && (u->right!=0) then
    s = PUSH(u->right, s);
    s = PUSH(u->left, s);
  end if
  delete u;
end while
    
```

After two iterations...



```


s = PUSH(root, s);
while s!=0 do
    s = POP(&u, s);
    ... do something
    if (u->left!= 0) && (u->right!=0) then
        s = PUSH(u->right, s);
        s = PUSH(u->left, s);
    end if
    delete u;
end while
    
```



```

s = PUSH(root, s);
while s!=0 do
    s = POP(&u, s);
    ... do something
    if (u->left!= 0) && (u->right!=0) then
        s = PUSH(u->right, s);
        s = PUSH(u->left, s);
    end if
    delete u;
end while
    
```

- Partition linked list and tree



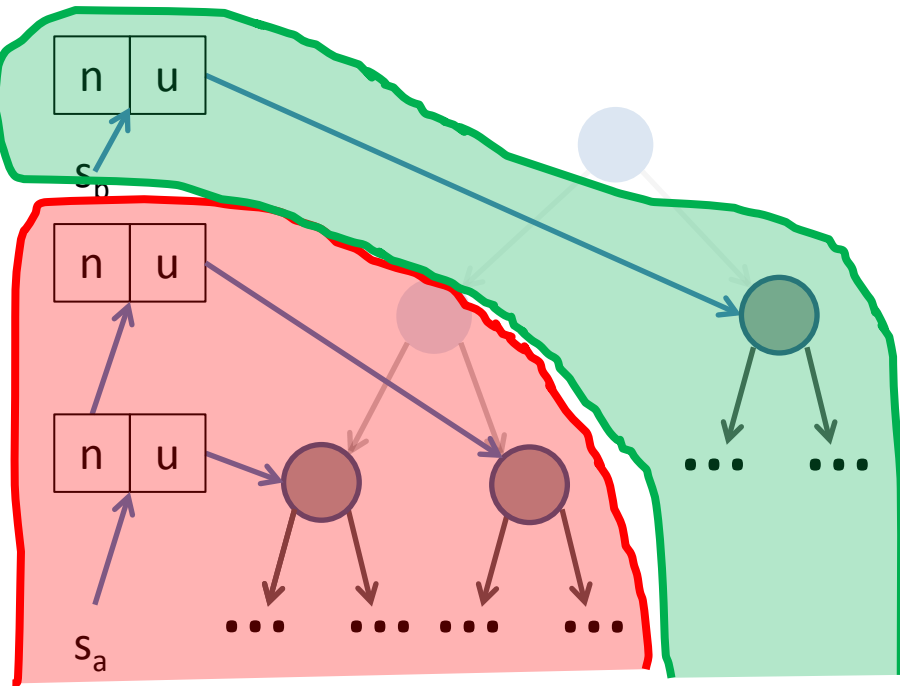
```

... preamble (accessing root node)

while  $s_a \neq 0$  do
    ... loop body (access left sub-tree)
end while

while  $s_b \neq 0$  do
    ... loop body (access right sub-tree)
end while
    
```

- Partition linked list and tree




... preamble (accessing root node)

```
while sa != 0 do
    ... loop body (access left sub-tree)
end while
```

```
while sb != 0 do
    ... loop body (access right sub-tree)
end while
```

- Partition linked list and tree
- Will the **red loop** ever access data in the **green partition**? No!
- Parallelization is legal (does not violate data dependencies)
- Why is it hard for a tool to figure this out?


Heap accessed  
in the next  
iteration



Heap accessed in  
some iteration in  
the future


Heap accessed  
in the next  
iteration

Heap accessed in  
some iteration in  
the future



- Do these iterations access the same memory cell?

Heap accessed  
in the next  
iteration




Heap accessed in  
some iteration in  
the future

- Do these iterations access the same memory cell?

heap[s]

Heap accessed  
in the next  
iteration




Heap accessed in  
some iteration in  
the future

- Do these iterations access the same memory cell?

`heap[heap[s].u]`

Heap accessed  
in the next  
iteration




Heap accessed in  
some iteration in  
the future

- Do these iterations access the same memory cell?

`heap[heap[s].u]`

`heap[s]`

Heap accessed  
in the next  
iteration




Heap accessed in  
some iteration in  
the future

- Do these iterations access the same memory cell?

`heap[heap[s].u]`

`heap[heap[s].n]`

Heap accessed  
in the next  
iteration




Heap accessed in  
some iteration in  
the future

- Do these iterations access the same memory cell?

`heap[heap[s].u]`

`heap[heap[heap[s].n].n]`

Heap accessed  
in the next  
iteration




Heap accessed in  
some iteration in  
the future

- Do these iterations access the same memory cell?

`heap[heap[s].u]`

`heap[heap[heap[heap[s].n].n].u]`

Heap accessed  
in the next  
iteration



Heap accessed in  
some iteration in  
the future


- Do these iterations access the same memory cell?

$\text{heap}[\text{heap}[s].u]$

$\text{heap}[\text{heap}[\text{heap}[s].n].n].u]$

$= ?$

Heap accessed  
in the next  
iteration



Heap accessed in  
some iteration in  
the future

Traverse a  
linked list


- Do these iterations access the same memory cell?

$\text{heap}[\text{heap}[s].u]$

$\text{heap}[\text{heap}[\text{heap}[\text{heap}[s].n].n].u]$

= ?

Heap accessed  
in the next  
iteration



Heap accessed in  
some iteration in  
the future

... which  
has links to  
sub-trees


- Do these iterations access the same memory cell?

$\text{heap}[\text{heap}[s].u]$

$\text{heap}[\text{heap}[\text{heap}[\text{heap}[s].n].n].u]$

= ?

Heap accessed  
in the next  
iteration



Heap accessed in  
some iteration in  
the future

... which  
has links to  
sub-trees

- Do these iterations access the same memory cell?

$\text{heap}[\text{heap}[s].u]$


$\text{heap}[\text{heap}[\text{heap}[\text{heap}[s].n].n].u]$

= ?

- Need to reason about structure, heap layout and disjointness
- None of this is explicit in the above representation

- Case study: High-level synthesis of dynamic data structures
- Challenge
- Motivating example
- **Leveraging separation logic**
- Implementation and results
- Outlook


Describe heap layout with formulae



$$\underbrace{s}_{\text{points to}} \rightarrow \underbrace{[u: u'_1, n: s'_1]}_{\text{record with fields } u \text{ and } n}$$


“s points to a record with fields *u* and *n*”

Describe heap  
layout with  
formulae




$$s \rightarrow [u: u'_1, n: s'_1] \wedge s'_1 \rightarrow [u: u'_2, n: s'_2]$$

Describe heap layout with formulae




$$s \rightarrow [u: u'_1, n: s'_1] \wedge s'_1 \rightarrow [u: u'_2, n: s'_2] \wedge s'_2 \rightarrow [u: u'_3, n: 0]$$

Describe heap layout with formulae



$$s \rightarrow [u: u'_1, n: s'_1] \wedge s'_1 \rightarrow [u: u'_2, n: s'_2] \wedge s'_2 \rightarrow [u: u'_3, n: 0] \\ \wedge u'_1 \rightarrow [l: u'_4, r: u'_5]$$


Describe heap  
layout with  
formulae



$$s \rightarrow [u: u'_1, n: s'_1] \wedge s'_1 \rightarrow [u: u'_2, n: s'_2] \wedge s'_2 \rightarrow [u: u'_3, n: 0]$$

$$\wedge u'_1 \rightarrow [l: u'_4, r: u'_5] \wedge u'_3 \rightarrow [l: u'_8, r: u'_9]$$


Describe heap layout with formulae



$$\begin{aligned}
 & s \rightarrow [u: u'_1, n: s'_1] \wedge s'_1 \rightarrow [u: u'_2, n: s'_2] \wedge s'_2 \rightarrow [u: u'_3, n: 0] \\
 & \wedge u'_1 \rightarrow [l: u'_4, r: u'_5] \wedge u'_3 \rightarrow [l: u'_8, r: u'_9] \wedge u'_2 \rightarrow [l: u'_6, r: u'_7]
 \end{aligned}$$




Describe heap layout with formulae



Conjunction '∧' does not rule out aliasing!

$$\begin{aligned}
 & s \rightarrow [u: u'_1, n: s'_1] \wedge s'_1 \rightarrow [u: u'_2, n: s'_2] \wedge s'_2 \rightarrow [u: u'_3, n: 0] \\
 & \wedge u'_1 \rightarrow [l: u'_4, r: u'_5] \wedge u'_3 \rightarrow [l: u'_8, r: u'_9] \wedge u'_2 \rightarrow [l: u'_6, r: u'_7] \\
 & \wedge \dots
 \end{aligned}$$

Describe heap layout with formulae



Formula below can also mean this

Conjunction

All  $u$ -pointers alias

$$u'_1 = u'_2 = u'_3 = u'_4 =$$

$$u'_5 = u'_6 = u'_7 = u'_8 = u'_9 = \dots$$


aliasing!

$$s \rightarrow [u: u'_1, n: s'_1] \wedge s'_1 \rightarrow [u: u'_2, n: s'_2] \wedge s'_2 \rightarrow [u: u'_3, n: 0]$$

$$\wedge u'_1 \rightarrow [l: u'_4, r: u'_5] \wedge u'_3 \rightarrow [l: u'_8, r: u'_9] \wedge u'_2 \rightarrow [l: u'_6, r: u'_7]$$

$$\wedge \dots$$

Describe heap layout with formulae




Could add loads of constraints

Conjunction '∧' does not rule out aliasing!

$$\begin{aligned}
 & s \rightarrow [u: u'_1, n: s'_1] \wedge s'_1 \rightarrow [u: u'_2, n: s'_2] \wedge s'_2 \rightarrow [u: u'_3, n: 0] \\
 & \wedge u'_1 \rightarrow [l: u'_4, r: u'_5] \wedge u'_3 \rightarrow [l: u'_8, r: u'_9] \wedge u'_2 \rightarrow [l: u'_6, r: u'_7] \\
 & \wedge \dots \wedge u'_1 \neq u'_2 \wedge u_1 \neq u'_3 \wedge u'_1 \neq u'_4 \wedge u'_1 \neq u'_5 \wedge u'_1 \neq u'_6 \wedge u'_1 \neq u'_7 \\
 & \wedge u'_1 \neq u'_8 \wedge u_1 \neq u'_9 \wedge u'_3 \neq u'_4 \wedge u'_3 \neq u'_5 \wedge u'_3 \neq u'_6 \wedge u'_3 \neq u'_7 \\
 & \wedge u'_3 \neq u'_8 \wedge u_2 \neq u'_3 \wedge u'_2 \neq u'_4 \wedge u'_2 \neq u'_5 \wedge u'_2 \neq u'_6 \wedge \dots
 \end{aligned}$$

Describe heap layout with formulae



- Tractable heap analysis
- Task: Split the heap formula into red and green partition

$$\begin{aligned}
 & s \rightarrow [u: u'_1, n: s'_1] * s'_1 \rightarrow [u: u'_2, n: s'_2] * s'_2 \rightarrow [u: u'_3, n: 0] \\
 & * u'_1 \rightarrow [l: u'_4, r: u'_5] * u'_3 \rightarrow [l: u'_8, r: u'_9] * u'_2 \rightarrow [l: u'_6, r: u'_7] \\
 & * \dots \wedge u'_1 \neq u'_2 \wedge u_1 \neq u'_3 \wedge u'_1 \neq u'_4 \wedge u'_1 \neq u'_5 \wedge u'_1 \neq u'_6 \wedge u'_1 \neq u'_7 \\
 & \wedge u'_1 \neq u'_8 \wedge u_1 \neq u'_9 \wedge u'_3 \neq u'_4 \wedge u'_3 \neq u'_5 \wedge u'_3 \neq u'_6 \wedge u'_3 \neq u'_7 \\
 & \wedge u'_3 \neq u'_8 \wedge u_2 \neq u'_3 \wedge u'_2 \neq u'_4 \wedge u'_2 \neq u'_5 \wedge u'_2 \neq u'_6 \wedge \dots
 \end{aligned}$$

- Symbolically execute the program using (a modified version of) *coreStar*
- *coreStarIL* syntax [Botincan et al. 2011]:

$$\begin{aligned} \text{CoreStmt} ::= & \bar{x} := \{Pre\}\{Post\} \\ & | \text{label } l \\ & | \text{goto } l_1, \dots, l_n \end{aligned}$$

*new(x)*:  $x := \{emp\}\{y' \mapsto \_ * \$retv1 = y'\}$       allocation

*[x] := 3*:  $\{x \mapsto y'\}\{x \mapsto 3\}$       heap assignment

- On-going work: LLVM-IR to *coreStarIL*

Given:

- Call site assertion before executing command  $C$ :  $\Delta_H$
- Pre and post specification of command  $C$ :  $\{P\} C \{Q\}$

Then

- coreStar's theorem prover finds the frame assertion  $\Delta_F$  such that:

$$\Delta_H \vdash P * \Delta_F$$

- New call site assertion after successful execution of  $C$ :

$$\Delta_H' = Q * \Delta_F$$

Raza et al. 2009:

Given:


- Call site assertion before executing command C:  $\Delta_H[L]$
- Pre and post specification of command C:  $\{P\} C \{Q\}$

Then

- the theorem prover finds the frame assertion  $\Delta_F$  such that:  
$$\Delta_H[L] \vdash P[L_A] * \Delta_F[L_F] \quad L_A: \text{accessed labels}$$
- New call site assertion after successful execution of C:  
$$\Delta'_H[L'] = Q[L_A \cup \{newlabel\}] * \Delta_F[L_F]$$


- Label sets describe the “heap footprint”
- We use them to partition the spatial formula
- Introduce cut-points:

*“A cut-point is a (non primed) program variable referencing a heaplet”*




- Peel-off the first loop iterations
- Insert cut-points (multiple options)

$$\begin{aligned}
 & s \rightarrow [u: u'_1, n: s'_1] \quad * \quad s'_1 \rightarrow [u: u'_2, n: s'_2] \quad * \quad s'_2 \rightarrow [u: u'_3, n: 0] \\
 & * \quad u'_1 \rightarrow [l: u'_4, r: u'_5] \quad * \quad u'_3 \rightarrow [l: u'_8, r: u'_9] \quad * \quad u'_2 \rightarrow [l: u'_6, r: u'_7] \\
 & * \quad \dots
 \end{aligned}$$




- Peel-off the first loop iterations
- Insert cut-points (multiple options)

$$\begin{aligned}
 & s \rightarrow [u: u'_1, n: s'_1] \quad * \quad s'_1 \rightarrow [u: u'_2, n: s'_2] \quad * \quad s'_2 \rightarrow [u: u'_3, n: 0] \\
 & * \quad u'_1 \rightarrow [l: u'_4, r: u'_5] \quad * \quad u'_3 \rightarrow [l: u'_8, r: u'_9] \quad * \quad u'_2 \rightarrow [l: u'_6, r: u'_7] \\
 & * \quad \dots
 \end{aligned}$$



- Peel-off the first loop iterations
- Insert cut-points (multiple options)

$$\begin{aligned}
 & s \rightarrow [u: u'_1, n: s'_1] \quad * \quad s'_1 \rightarrow [u: u'_2, n: s'_2] \quad * \quad s'_2 \rightarrow [u: u'_3, n: 0] \\
 & * \quad u'_1 \rightarrow [l: u'_4, r: u'_5] \quad * \quad u'_3 \rightarrow [l: u'_8, r: u'_9] \quad * \quad u'_2 \rightarrow [l: u'_6, r: u'_7] \\
 & * \quad \dots
 \end{aligned}$$




- Peel-off the first loop iterations
- Insert cut-points (multiple options)

$$s \rightarrow [u: u'_1, n: s'_1] \quad * \quad s'_1 \rightarrow [u: u'_2, n: s'_2] \quad * \quad s'_2 \rightarrow [u: u'_3, n: 0]$$


$$* \quad u'_1 \rightarrow [l: u'_4, r: u'_5] \quad * \quad u'_3 \rightarrow [l: u'_8, r: u'_9] \quad * \quad u'_2 \rightarrow [l: u'_6, r: u'_7]$$

\* ...




- Peel-off the first loop iterations
- Insert cut-points (multiple options)

$$\begin{aligned}
 & s \rightarrow [u: u'_1, n: s'_1] \mathbf{a} * s'_1 \rightarrow [u: u'_2, n: s'_2] \mathbf{a} * s'_2 \rightarrow [u: u'_3, n: 0] \\
 & * u'_1 \rightarrow [l: u'_4, r: u'_5] \mathbf{a} * u'_3 \rightarrow [l: u'_8, r: u'_9] \mathbf{a} * u'_2 \rightarrow [l: u'_6, r: u'_7] \\
 & * \dots
 \end{aligned}$$




- Peel-off the first loop iterations
- Insert cut-points (multiple options)

$$\begin{aligned}
 & s \rightarrow [u: u'_1, n: s'_1] \mathbf{a} * s'_1 \rightarrow [u: u'_2, n: s'_2] \mathbf{a} * s'_2 \rightarrow [u: u'_3, n: 0] \\
 & * u'_1 \rightarrow [l: u'_4, r: u'_5] \mathbf{a} * u'_3 \rightarrow [l: u'_8, r: u'_9] \mathbf{a} * u'_2 \rightarrow [l: u'_6, r: u'_7] \\
 & * \dots
 \end{aligned}$$



- Peel-off the first loop iterations
- Insert cut-points (multiple options)

$$\begin{aligned}
 & s \rightarrow [u: u'_1, n: s'_1] \mathbf{a} * s'_1 \rightarrow [u: u'_2, n: s'_2] \mathbf{a} * s'_2 \rightarrow [u: u'_3, n: 0] \mathbf{b} \\
 & * u'_1 \rightarrow [l: u'_4, r: u'_5] \mathbf{a} * u'_3 \rightarrow [l: u'_8, r: u'_9] \mathbf{a} * u'_2 \rightarrow [l: u'_6, r: u'_7] \mathbf{b} \\
 & * \dots
 \end{aligned}$$



- Peel-off the first loop iterations
- Insert cut-points (multiple options)

$$s \rightarrow [u: u'_1, n: s'_1] \mathbf{a} * s'_1 \rightarrow [u: u'_2, n: s'_2] \mathbf{a} * s'_2 \rightarrow [u: u'_3, n: 0] \mathbf{b}$$

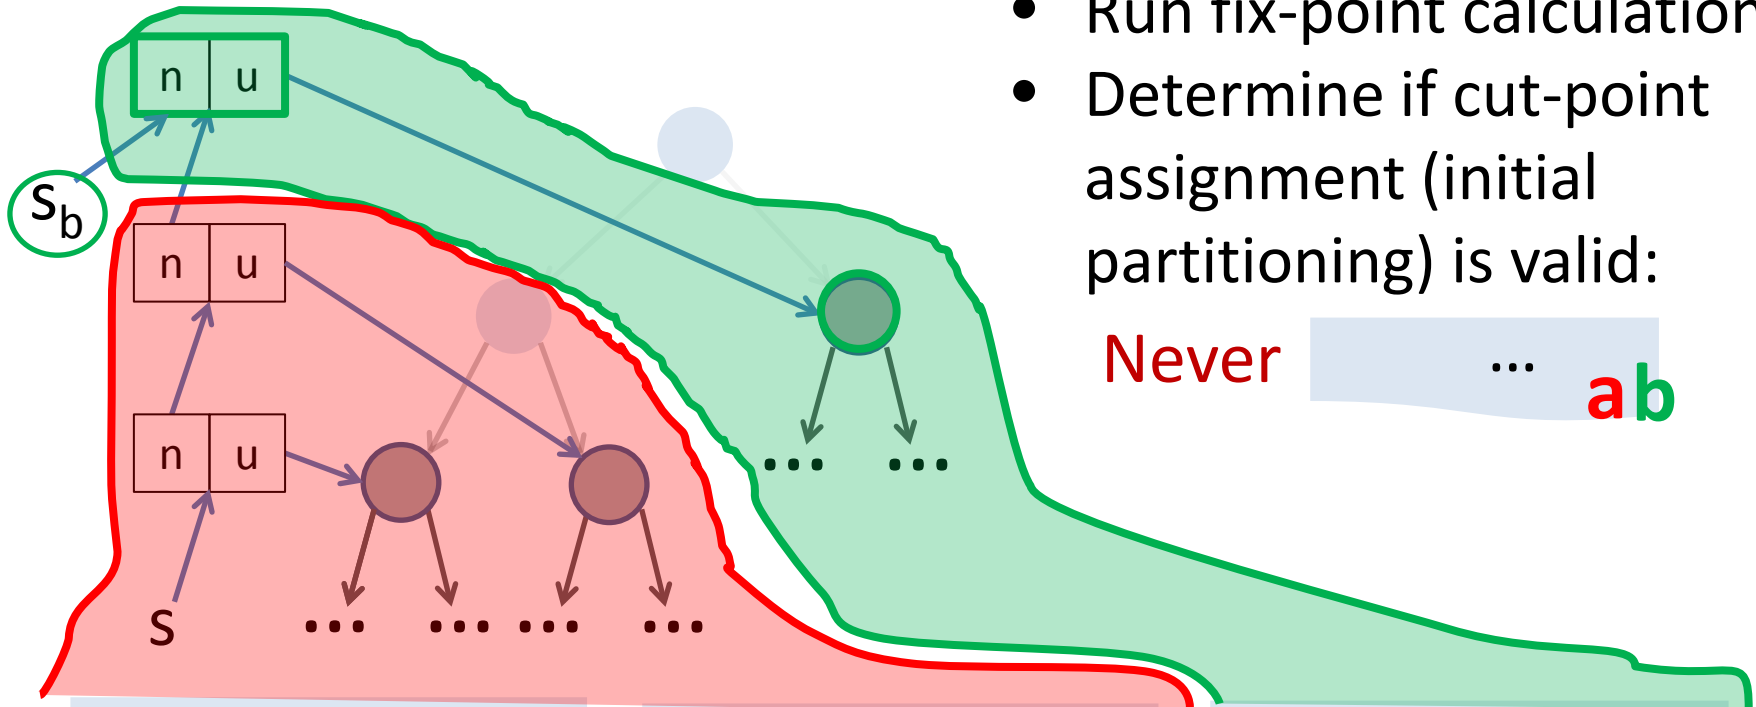
$$* u'_1 \rightarrow [l: u'_4, r: u'_5] \mathbf{a} * u'_3 \rightarrow [l: u'_8, r: u'_9] \mathbf{a} * u'_2 \rightarrow [l: u'_6, r: u'_7] \mathbf{b}$$

\* ...

- Symbolically execute loop iterations until heap partitioning can be shown to be loop invariant
- Based on abstraction:

$$x \mapsto [n: x'_1] * ls(x'_1, \text{nil}) \rightsquigarrow ls(x, \text{nil})$$

- Using a heuristic for folding [Magill et al. 2006]:  
Don't fold across program variables
- Avoids folding across cut-points



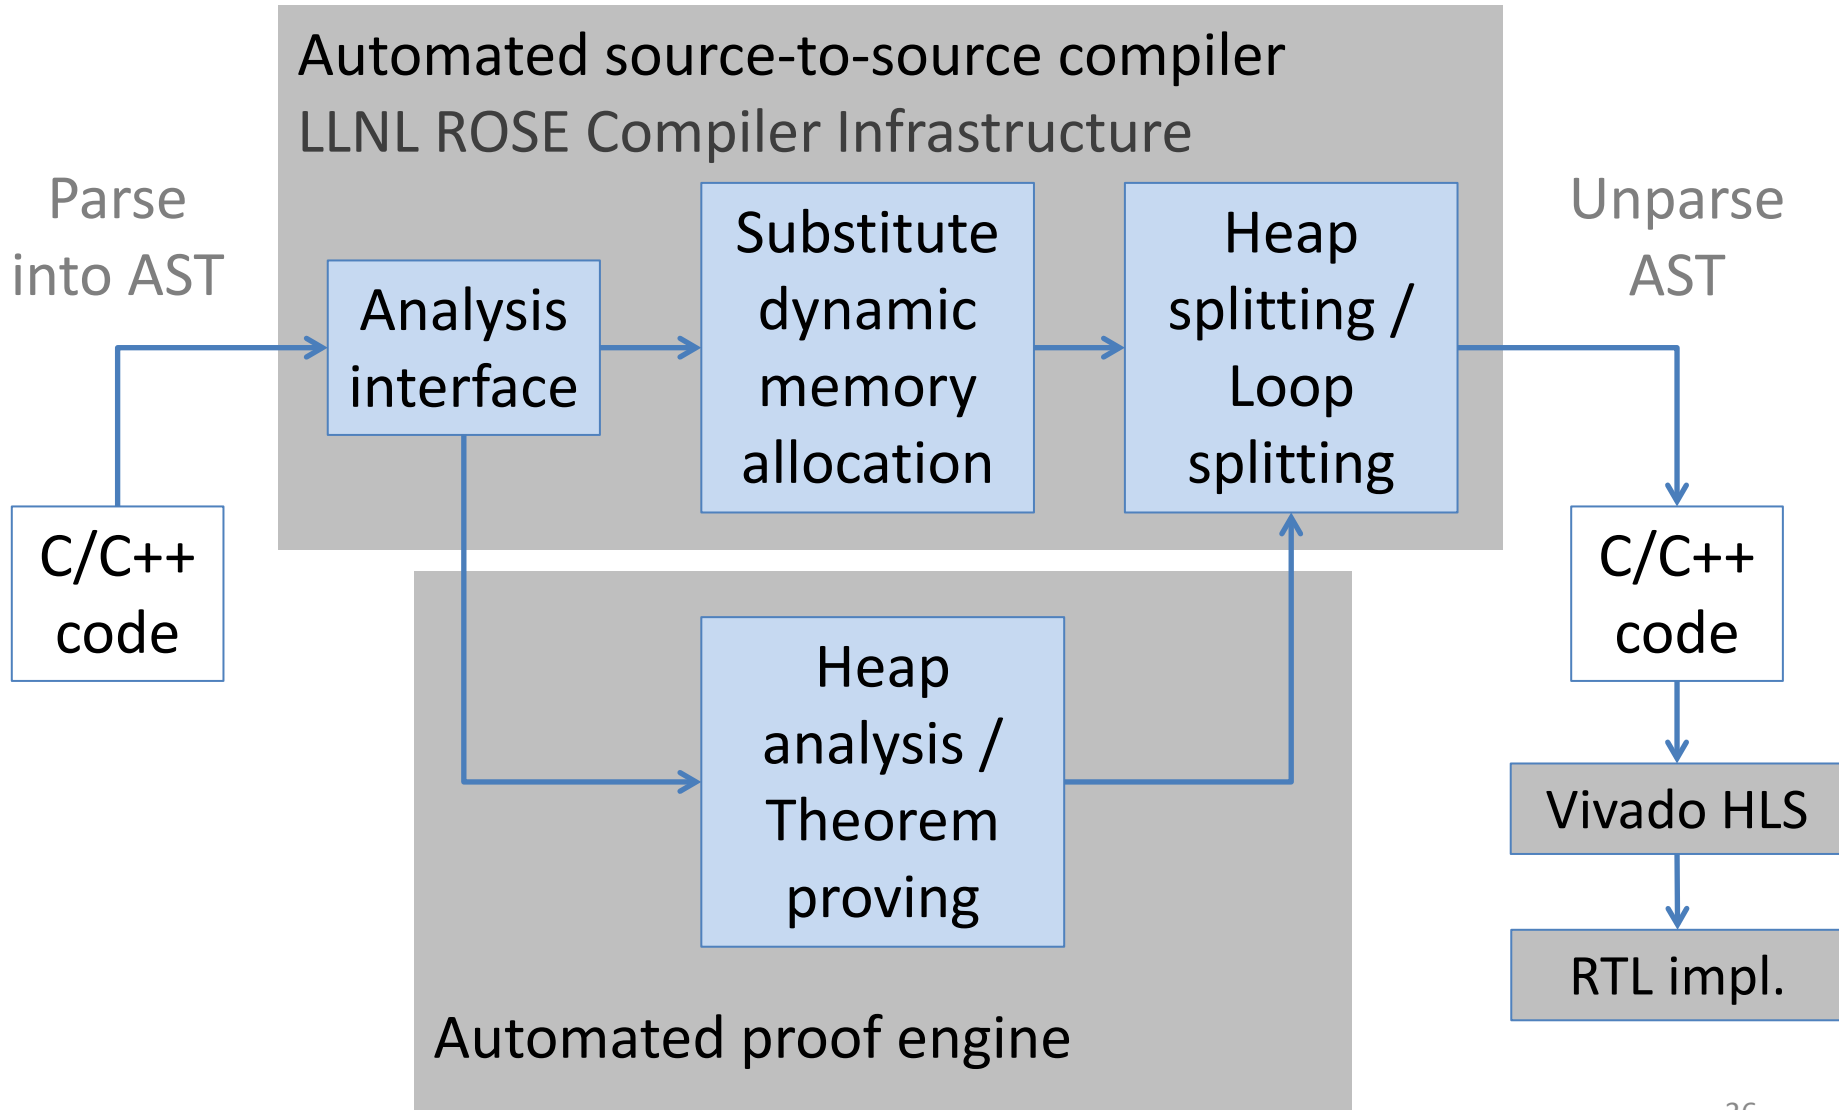
- Run fix-point calculation
- Determine if cut-point assignment (initial partitioning) is valid:

Never ... **a****b**


$$\begin{aligned}
 & s \rightarrow [u: u'_1, n: s'_1] \mathbf{a} * s'_1 \rightarrow [u: u'_2, n: s'_2] \mathbf{a} * s'_2 \rightarrow [u: u'_3, n: 0] \mathbf{b} \\
 & * u'_1 \rightarrow [l: u'_4, r: u'_5] \mathbf{a} * u'_3 \rightarrow [l: u'_8, r: u'_9] \mathbf{a} * u'_2 \rightarrow [l: u'_6, r: u'_7] \mathbf{b} \\
 & * \dots
 \end{aligned}$$




Communication free parallelism




- Case study: High-level synthesis of dynamic data structures
- Challenge
- Motivating example
- Leveraging separation logic
- **Implementation and results**
- Outlook



## Tree-based *K*-means clustering




	P	Slices	Clock	Cycles	
<b>1 Merger (linked lists)</b>					
Baseline (no par.)	1	574	9.0 ns	21167k	 x4
Autom. Parallelization	4	965	8.7 ns	5483k	
<b>2 Tree deletion (tree, linked list)</b>					
Baseline (no par.)	1	1521	5.2 ns	901k	 x2
Autom. Parallelization	2	4069	6.0 ns	487k	
<b>3 K-means (tree, linked list, single heap records)</b>					
Baseline (no par.)	1	2694	6.1 ns	1120k	 x2
Autom. Parallelization	2	5618	7.0 ns	606k	


	P	Slices	Clock	Cycles	
<b>1 Merger (linked lists)</b>					
Baseline (no par.)	1	574	9.0 ns	21167k	 x4
Autom. Parallelization	4	965	8.7 ns	5483k	
<b>2 Tree deletion (tree, linked list)</b>					
Baseline (no par.)	1	1521	5.2 ns	901k	 x2
Autom. Parallelization	2	4069	6.0 ns	487k	
<b>3 K-means (tree, linked list, single heap records)</b>					
Baseline (no par.)	1	2694	6.1 ns	1120k	 x2
Autom. Parallelization	2	5618	7.0 ns	606k	
Hand-optimized HLS	2	5492	5.5 ns	165k	

Manual loop flattening, pipelining, custom bit widths, data streaming directives, data packing, ...

	P	Slices	Clock	Cycles	
<b>1 Merger (linked lists)</b>					
Baseline (no par.)	1	574	9.0 ns	21167k	x4
Autom. Parallelization	4	965	8.7 ns	5483k	
<b>2 Tree deletion (tree, linked list)</b>					
Baseline (no par.)	1	1521	5.2 ns	901k	x2
Autom. Parallelization	2	4069	6.0 ns	487k	
<b>3 K-means (tree, linked list, single heap records)</b>					
Baseline (no par.)	1	2694	6.1 ns	1120k	x2
Autom. Parallelization	2	5618	7.0 ns	606k	
Hand-optimized HLS	2	5492	5.5 ns	165k	x3.6

Manual loop flattening, pipelining, custom bit widths, data streaming directives, data






Cache  
coherence

Communication-free parallelism:  
Caches are truly private

- Case study: High-level synthesis of dynamic data structures
- Challenge
- Motivating example
- Leveraging separation logic
- Implementation and results
- **Outlook**

1. Worst-case/ average case bounds on heap usage
2. Separation logic and temporal logic

## Profiling the tree-based clustering app



- Our abstraction gets rid of time
- Sometimes we need a notion of time
  - Rescheduling loop iterations
- Including temporal operators like “next”, “until”, etc. in our loop invariants can be useful
- What stops us from combining separation logic and temporal operators?

- Limited HLS support for heap-manipulating programs
- Static analysis of heap-manipulating programs
  - Leveraging recent advances in separation logic
  - Distribute heap across on-chip memory banks
  - Loop parallelization
- Tool implementation
  - Automated heap analyzer
  - Source-to-source transformations (synthesizability and parallelization)
  - Successful parallelization using standard HLS tool
- Future work
  - Compute worst-case/average-case bounds on heap usage
  - Integration of temporal logic in the analysis to enable more loop optimizations

**Thank you for listening.**

f.winterstein12@imperial.ac.uk

**<http://cas.ee.ic.ac.uk/people/fw1811/>**