

The Reduceron

Widening the von Neumann Bottleneck
for Graph Reduction using an FPGA

The Reduceron

Widening the von Neumann Bottleneck
for Graph Reduction using an FPGA



The Reduceron

Widening the von Neumann Bottleneck
for Graph **Reduction** using an FPGA

What is Reduction?

1 + 1 reduces to *2*

double (1 + 1) reduces to *4*

True && False reduces to *False*

append [1,2] [3] reduces to *[1,2,3]*

Reduction Strategies

Suppose that *double* is defined by

$$\mathit{double} \ x \quad = \quad x + x$$

and we wish to reduce

$$\mathit{double} \ (1 + 1)$$

then there are **two** main strategies.

1. Innermost Strategy

Suppose that *double* is defined by

$$\mathit{double} \ x \quad = \quad x + x$$

then, using the **innermost strategy**,

$$\begin{aligned} & \mathit{double} \ (\underline{1 + 1}) \\ \Rightarrow & \underline{\mathit{double} \ 2} \\ \Rightarrow & \underline{2 + 2} \\ \Rightarrow & 4 \end{aligned}$$

2. Outermost Strategy

Suppose that *double* is defined by

$$\mathit{double} \ x \quad = \quad x + x$$

then, using the **outermost strategy**,

$$\begin{aligned} & \mathit{double} \ (1 + 1) \\ \Rightarrow & \quad \underline{(1 + 1)} + (1 + 1) \\ \Rightarrow & \quad 2 + \underline{(1 + 1)} \\ \Rightarrow & \quad \underline{2 + 2} \\ \Rightarrow & \quad 4 \end{aligned}$$

2. Outermost Strategy

Suppose that *double* is defined by

$$\mathit{double} \ x \quad = \quad x + x$$

then, using the **outermost strategy**,

$$\begin{aligned} & \mathit{double} \ (1 + 1) \\ \Rightarrow & \quad \underline{(1 + 1)} + \underline{(1 + 1)} \\ \Rightarrow & \quad 2 + \underline{(1 + 1)} \\ \Rightarrow & \quad \underline{2 + 2} \\ \Rightarrow & \quad 4 \end{aligned}$$

Reduced twice!

Conclusion

Innermost Reduction performs fewer reductions than Outermost Reduction.

NOT NECESSARILY!

Let's play the game again, this time with a different function.

Reduction Strategies

Suppose that $\&\&$ is defined by

$$x \ \&\& \ y \ = \ \mathit{if} \ x \ \mathit{then} \ y \ \mathit{else} \ \mathit{False}$$

and we wish to reduce

$$\mathit{False} \ \&\& \ (\mathit{True} \ \&\& \ \mathit{True})$$

then there are two strategies.

1. Innermost Reduction

Suppose that $\&\&$ is defined by

$$x \ \&\& \ y \ = \ \textit{if } x \ \textit{then } y \ \textit{else } \textit{False}$$

then, using the **innermost strategy**,

$$\begin{aligned} & \textit{False} \ \&\& \ (\underline{\textit{True} \ \&\& \ \textit{True}}) \\ \Rightarrow & \ \textit{False} \ \&\& \\ & \quad (\underline{\textit{if } \textit{True} \ \textit{then } \textit{True} \ \textit{else } \textit{False}}) \\ \Rightarrow & \ \textit{False} \ \&\& \ \textit{True} \\ \Rightarrow & \ \textit{if } \textit{False} \ \textit{then } \textit{True} \ \textit{else } \textit{False} \\ \Rightarrow & \ \textit{False} \end{aligned}$$

2. Outermost Reduction

Suppose that $\&\&$ is defined by

$$x \ \&\& \ y \ = \ \mathit{if} \ x \ \mathit{then} \ y \ \mathit{else} \ \mathit{False}$$

then, using the **outermost strategy**,

$$\begin{aligned} & \underline{\mathit{False} \ \&\& \ (\mathit{True} \ \&\& \ \mathit{True})} \\ \Rightarrow & \ \mathit{if} \ \mathit{False} \ \mathit{then} \ \mathit{True} \ \&\& \ \mathit{True} \\ & \ \mathit{else} \ \mathit{False} \\ \Rightarrow & \ \mathit{False} \end{aligned}$$

2. Outermost Reduction

Suppose that $\&\&$ is defined by

$x \ \&\& \ y \ = \ \text{if } x \ \text{then } y \ \text{else } \text{False}$

then, using the **outermost strategy**,

$\text{False} \ \&\& \ (\text{True} \ \&\& \ \text{True})$
 $\Rightarrow \ \text{if } \text{False} \ \text{then } \text{True} \ \&\& \ \text{True}$
 $\qquad \qquad \qquad \text{else } \text{False}$
 $\Rightarrow \ \text{False}$

Never reduced!

Conclusion

How many times is each argument reduced?

Innermost: **Exactly once**

Outermost: **Zero or more**

Lazy Reduction

How many times is each argument reduced?

Innermost: **Exactly once**

Outermost: **Zero or more**

Lazy: **Zero or once**

Outermost, with magic!



Optimal



The Reduceron

Widening the von Neumann Bottleneck
for **Graph** Reduction using an FPGA

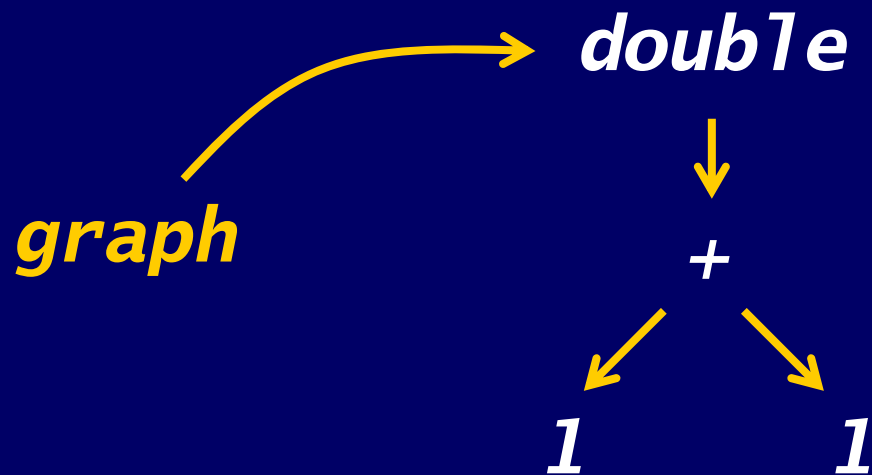


Expressions as Graphs

Suppose that *double* is defined by

$$\mathit{double} \ x \quad = \quad x + x$$

and we wish to reduce *double* (1 + 1).

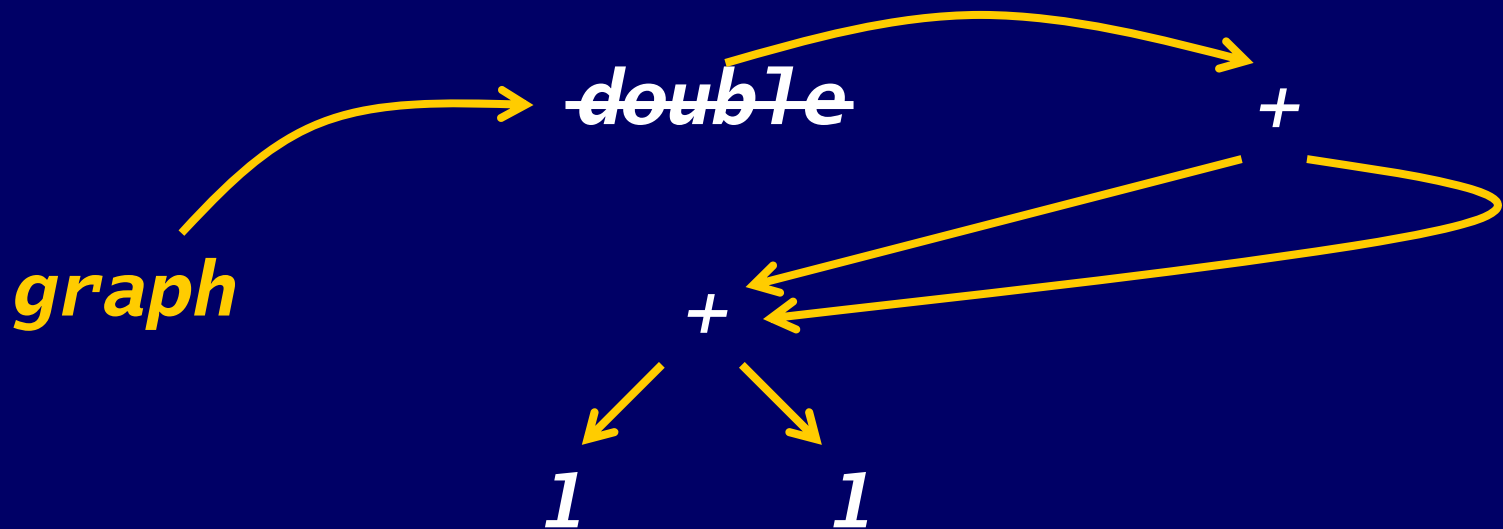


Graph Reduction

Suppose that *double* is defined by

$$\mathit{double} \ x \ = \ x + x$$

then, using the **lazy reduction** strategy,



Graph Reduction

Suppose that *double* is defined by

$$\mathit{double} \ x \quad = \quad x + x$$

then, using the **lazy reduction** strategy,



Graph Reduction

Suppose that *double* is defined by

$$\mathit{double} \ x \quad = \quad x + x$$

then, using the **lazy reduction** strategy,



The Reduceron

Widening the von Neumann Bottleneck
for Graph Reduction using an FPGA



A Graph Reduction Machine

Suppose that f is defined by

$$f\ x\ y\ z = g\ y\ (h\ z\ x)$$

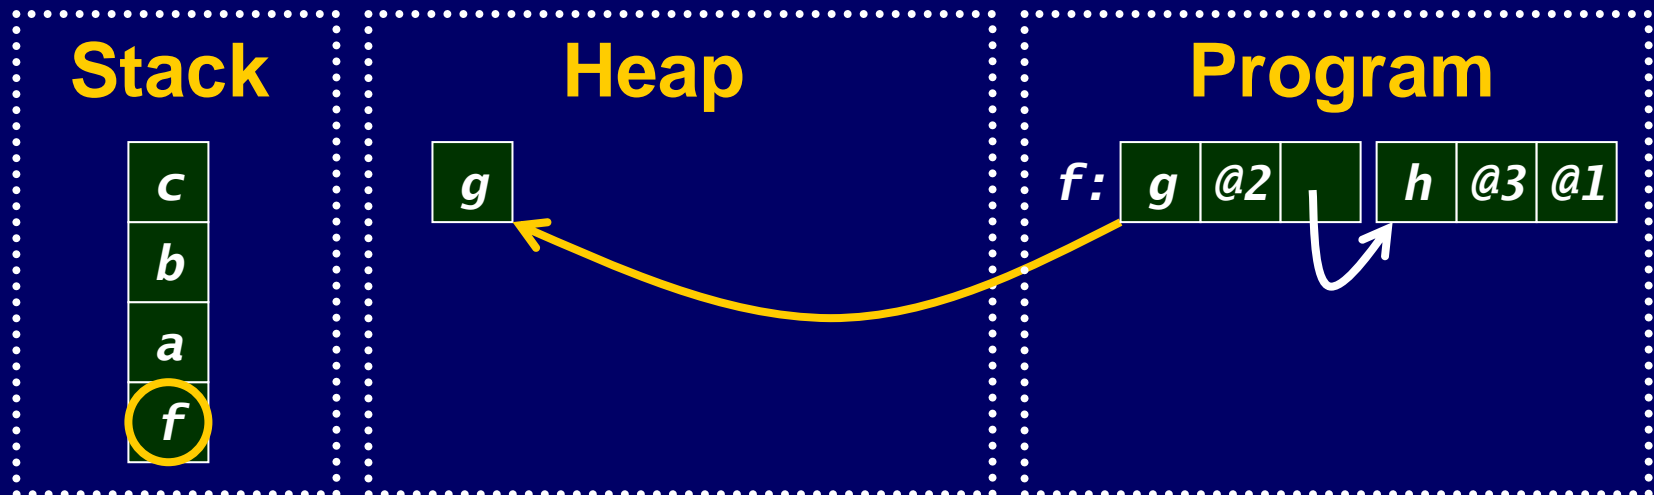
where g and h are functions and the following machine-state arises during reduction.



A Graph Reduction Machine

Operation: $f \leftarrow \text{Stack}[0]$
 $g \leftarrow \text{Code}[f]$
 $g \rightarrow \text{Heap}$

Count: 3



A Graph Reduction Machine

Operation: $arg \leftarrow Code[f+1]$
 $b \leftarrow Stack[arg]$
 $b \rightarrow Heap$

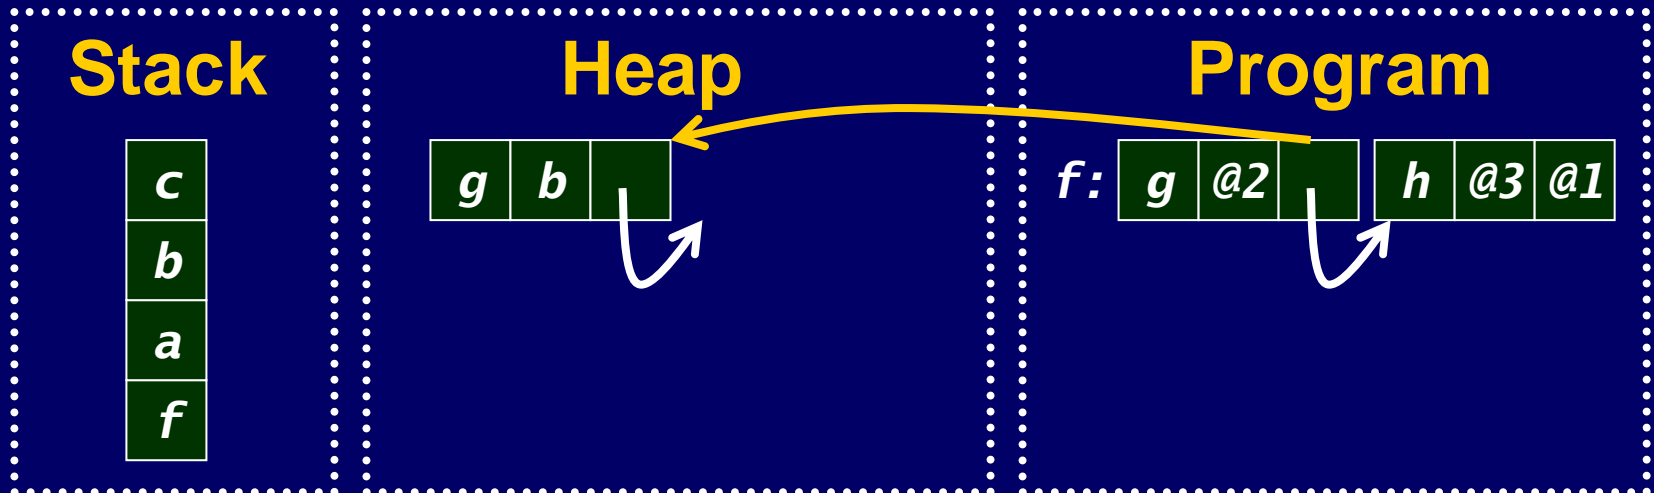
Count: 6



A Graph Reduction Machine

Operation: $ptr \leftarrow Code[f+2]$
 $ptr' \rightarrow Heap$

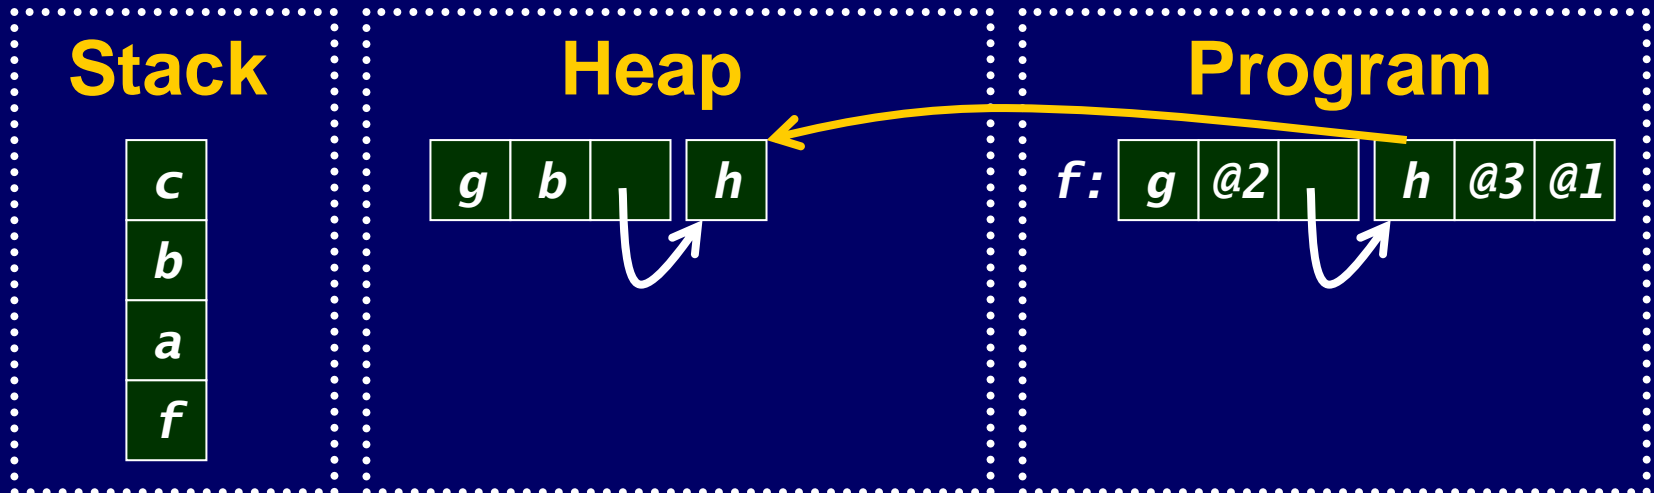
Count: 8



A Graph Reduction Machine

Operation: $h \leftarrow \text{Code}[f+3]$
 $h \rightarrow \text{Heap}$

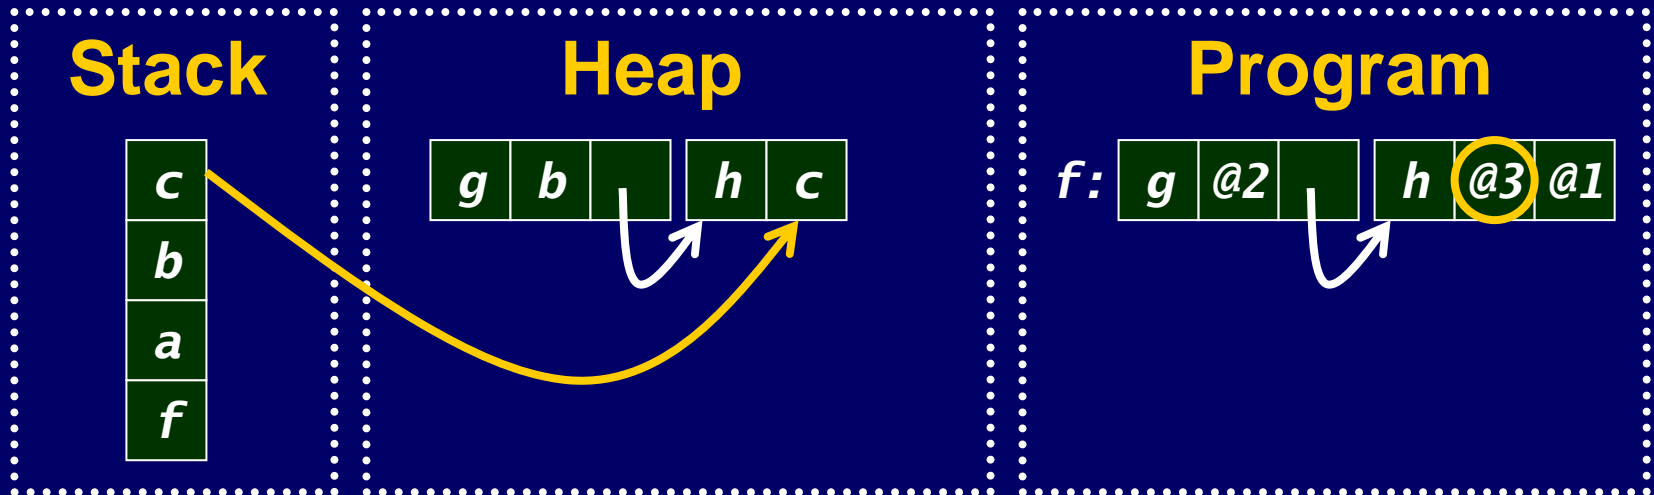
Count: 10



A Graph Reduction Machine

Operation: $arg \leftarrow Code[f+4]$
 $c \leftarrow Stack[arg]$
 $c \rightarrow Heap$

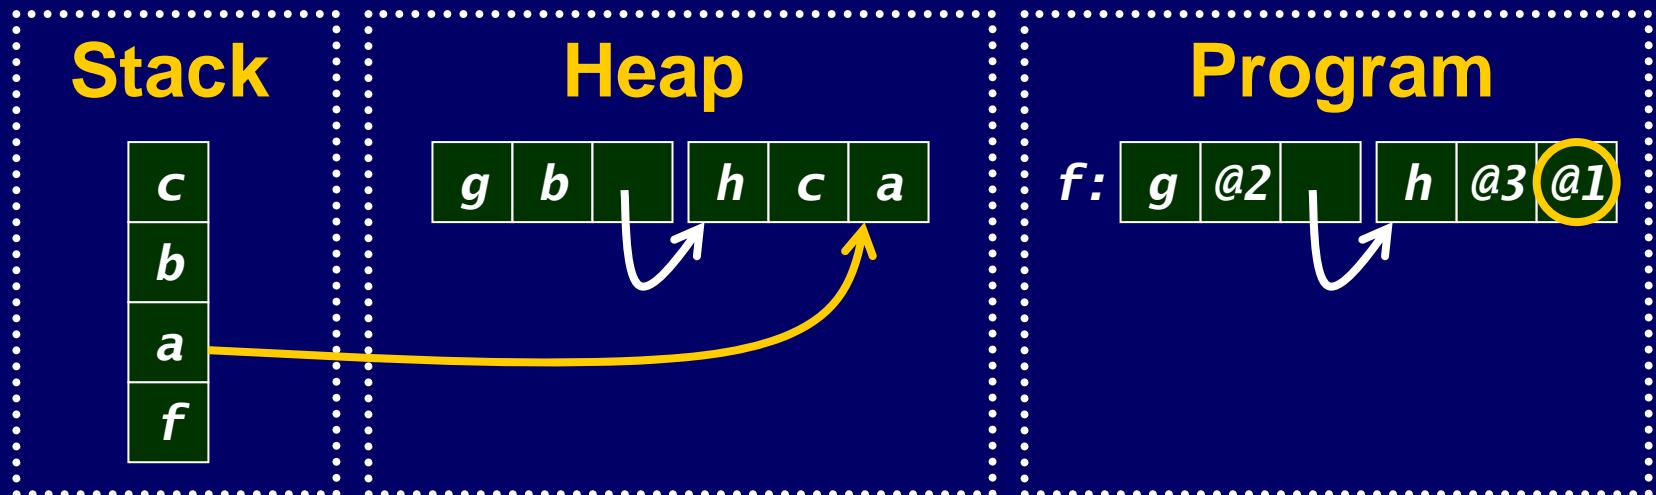
Count: 13



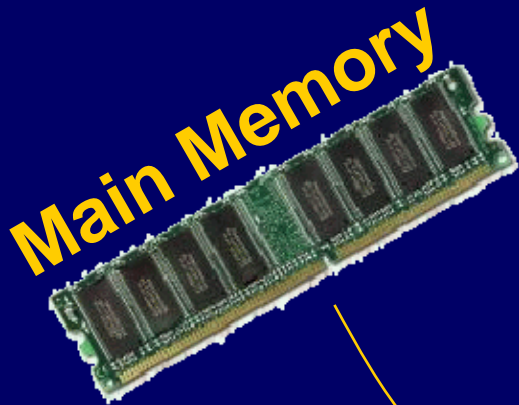
A Graph Reduction Machine

Operation: $arg \leftarrow Code[f+5]$
 $a \leftarrow Stack[arg]$
 $a \rightarrow Heap$

Count: 16



The von Neumann Bottleneck

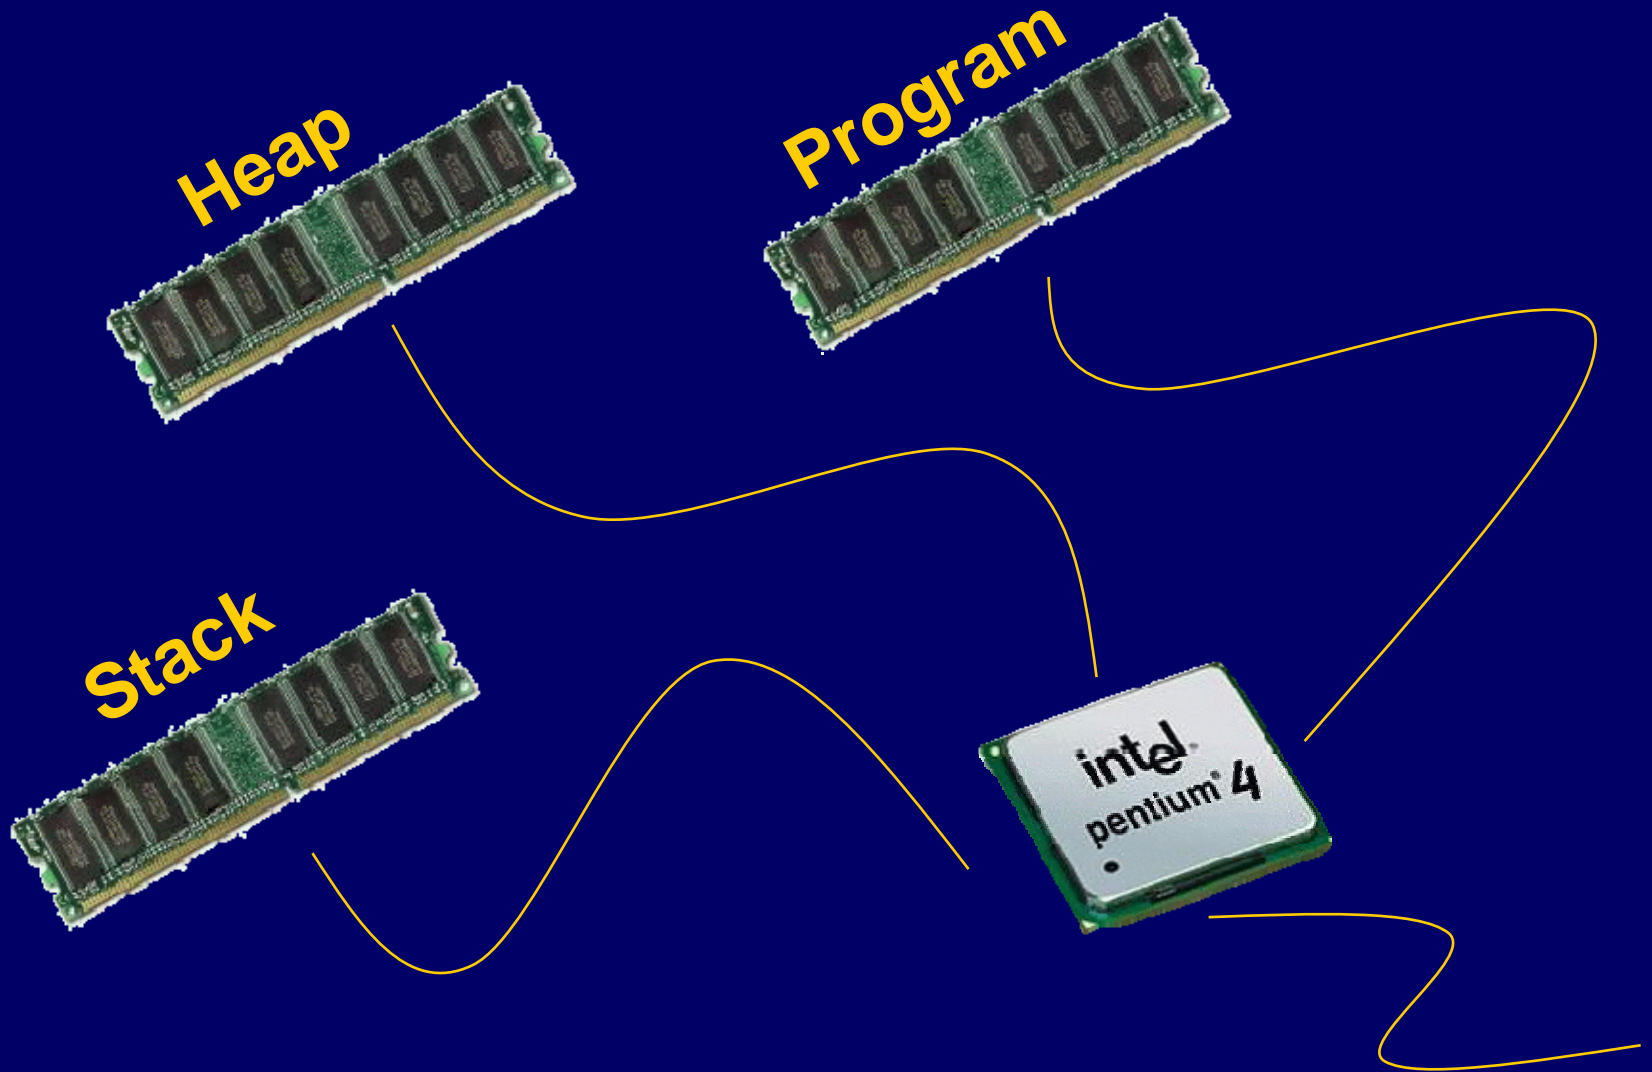


One word at a time.

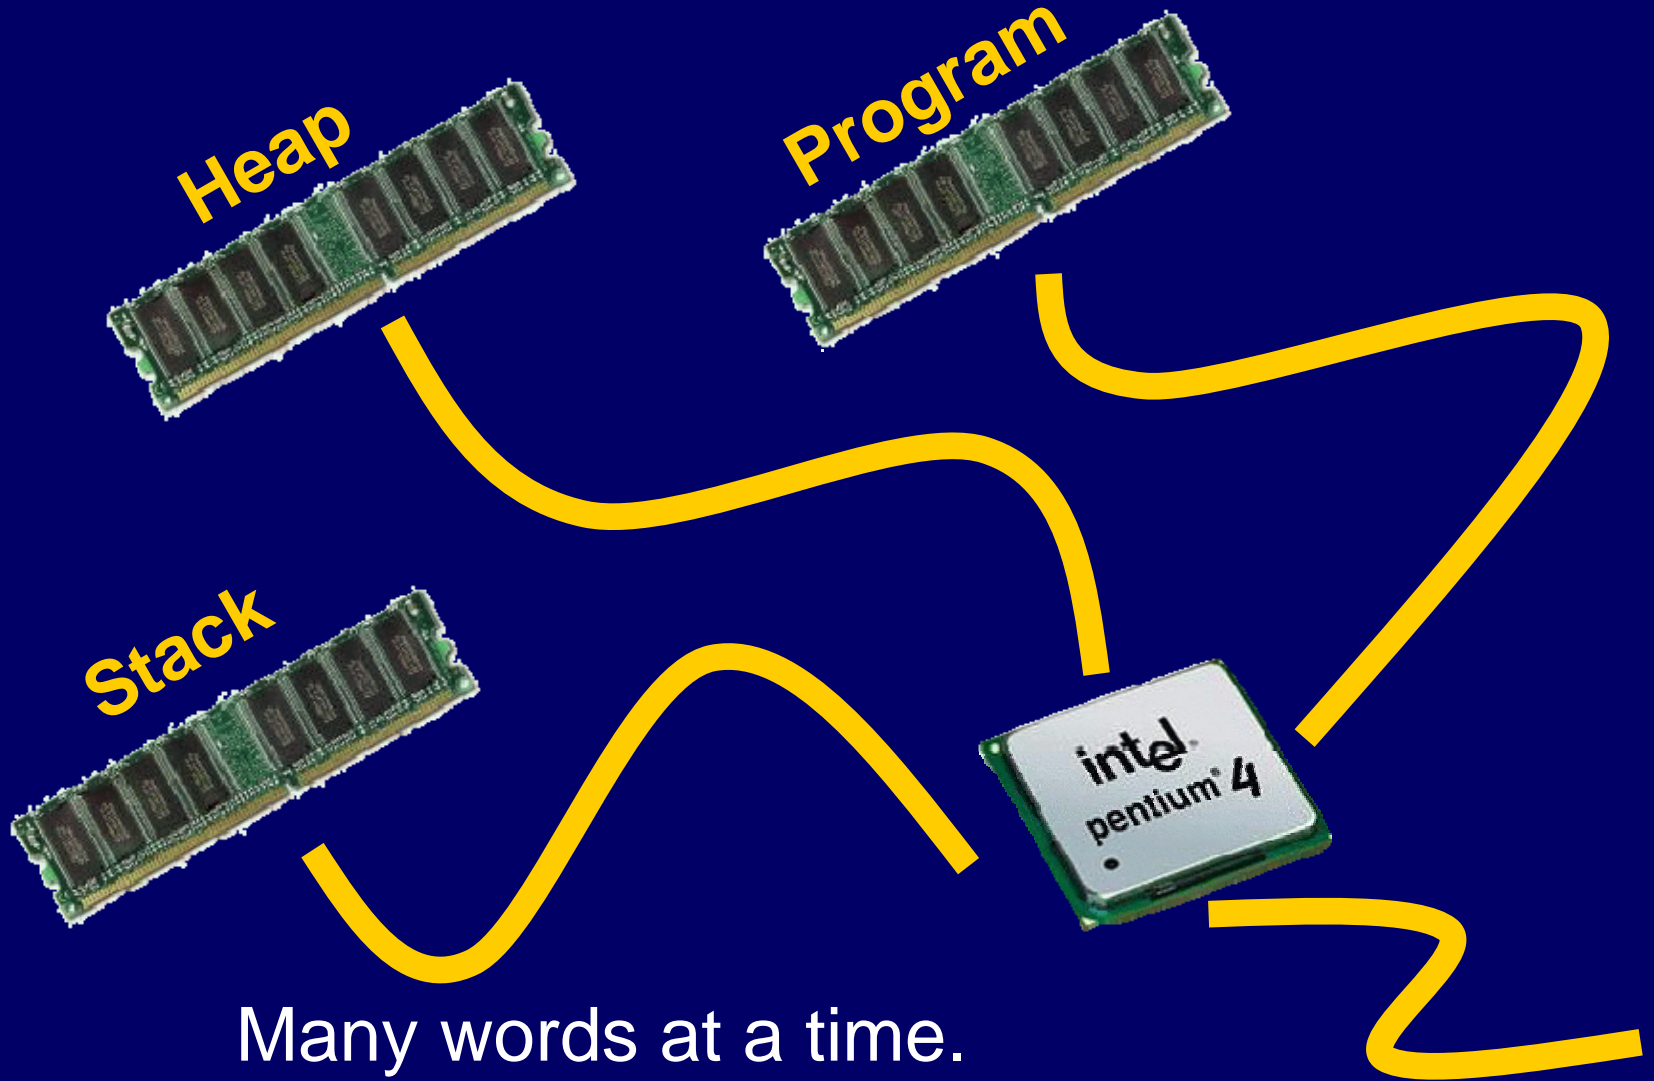
Each of the **16** memory transactions must be done sequentially.



Widening the Bottleneck

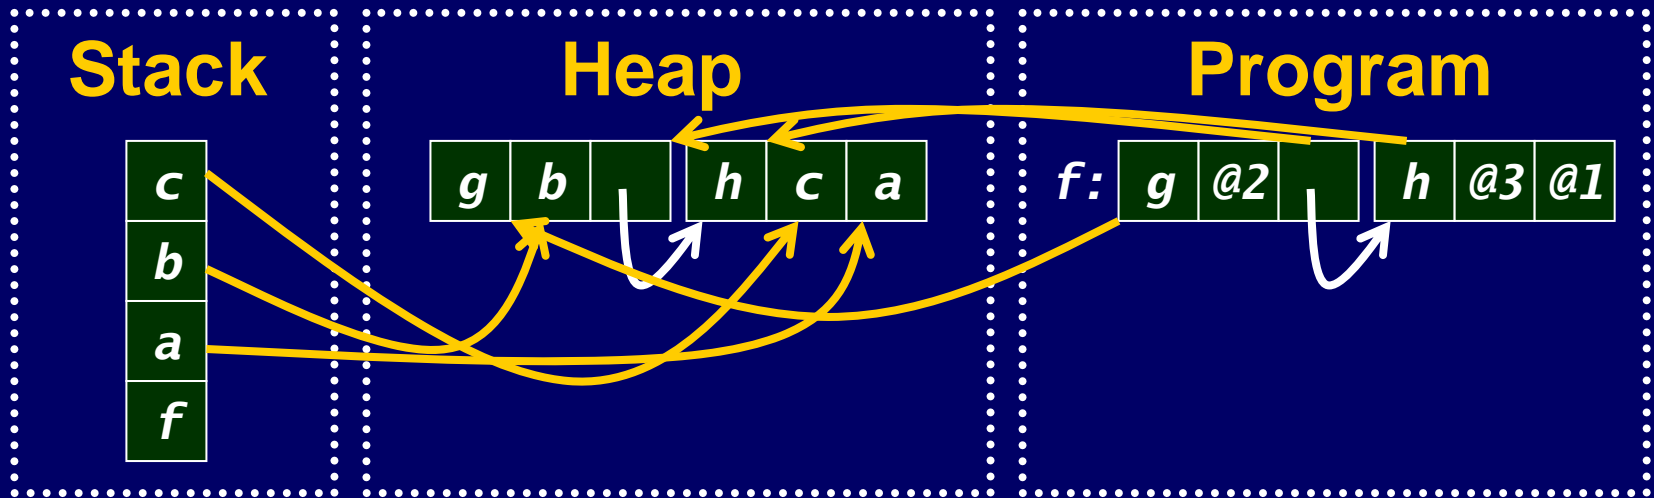


Widening the Bottleneck, again



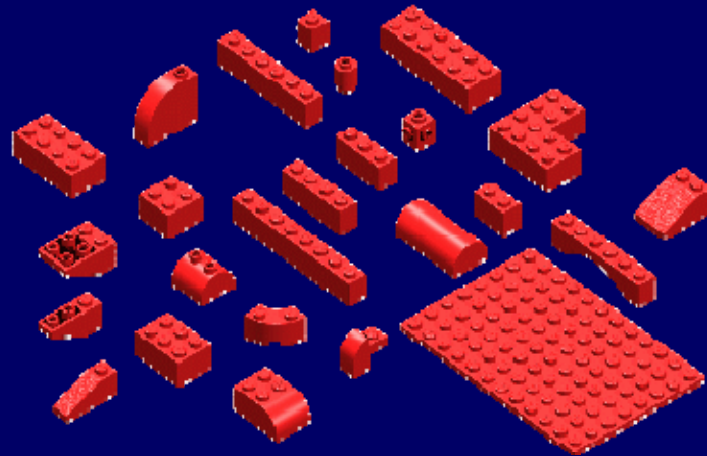
Many words at a time.

Applying a function “in one go”



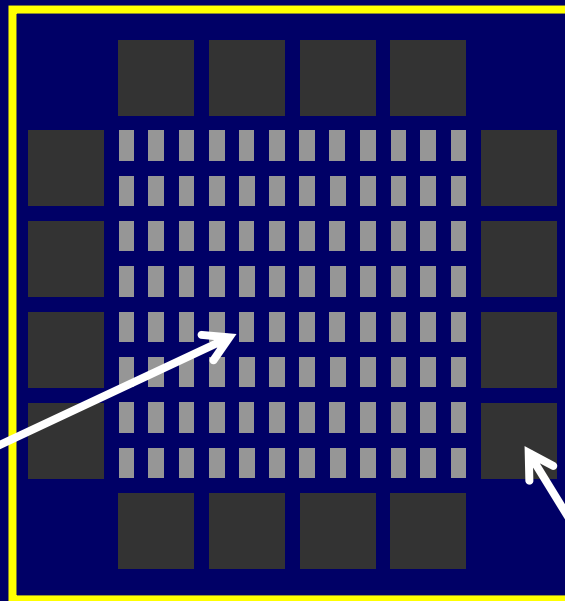
The Reduceron

Widening the von Neumann Bottleneck
for Graph Reduction **using an FPGA**



An FPGA

A large, fixed set of components that can be connected together in any desired way.



Logic cells

(and-gates, flip-flops, ...)

Memory blocks

(addressable separately)

Experimental Results

Wide Reduceron

(uses wide, parallel memories)

5x faster than

Narrow Reduceron

(single connection to memory)

Wide Reduceron

at 92MHz on Virtex-II FPGA

5x slower than

GHC -O2

(advanced optimising compiler)

at 2800MHz on Pentium-4 PC

(See thesis for theoretical results.)

Dinner Time!

