

Lecture #5 – Memory Management; Intro MIPS

2005-09-14

There is one handout today at the front and back of the room!



Lecturer PSOE, new dad Dan Garcia

www.cs.berkeley.edu/~ddgarcia

iPod nano ⇒

Thinner than a pencil,
the newest iPod release again
benefits from small hard drives.
We'll talk about how drives work!



Review

- **C has 3 pools of memory**
 - **Static storage**: global variable storage, basically permanent, entire program run
 - **The Stack**: local variable storage, parameters, return address
 - **The Heap** (dynamic storage): `malloc()` grabs space from here, `free()` returns it.
Nothing to do with heap data structure!
- `malloc()` handles free space with freelist.
Three different ways:
 - **First fit** (find first one that's free)
 - **Next fit** (same as first, start where ended)
 - **Best fit** (finds most “snug” free space)
- One problem with all three is **small fragments!**



Slab Allocator

- **A different approach to memory management (used in GNU libc)**
- **Divide blocks in to “large” and “small” by picking an arbitrary threshold size. Blocks larger than this threshold are managed with a freelist (as before).**
- **For small blocks, allocate blocks in sizes that are powers of 2**
 - **e.g., if program wants to allocate 20 bytes, actually give it 32 bytes**





Slab Allocator

- **Bookkeeping for small blocks is relatively easy: just use a *bitmap* for each range of blocks of the same size**
- **Allocating is easy and fast: compute the size of the block to allocate and find a free bit in the corresponding bitmap.**
- **Freeing is also easy and fast: figure out which slab the address belongs to and clear the corresponding bit.**



Slab Allocator



16 byte block bitmap: 11011000

32 byte block bitmap: 0111

64 byte block bitmap: 00



Slab Allocator Tradeoffs

- **Extremely fast for small blocks.**
- **Slower for large blocks**
 - **But presumably the program will take more time to do something with a large block so the overhead is not as critical.**
- **Minimal space overhead**
- **No fragmentation (as we defined it before) for small blocks, but still have wasted space!**



Internal vs. External Fragmentation

- With the slab allocator, difference between requested size and next power of 2 is wasted
 - e.g., if program wants to allocate 20 bytes and we give it a 32 byte block, 12 bytes are unused.
- We also refer to this as fragmentation, but call it *internal fragmentation* since the wasted space is actually within an allocated block.
- **External fragmentation**: wasted space between allocated blocks.




Buddy System

- **Yet another memory management technique (used in Linux kernel)**
- **Like GNU's "slab allocator", but only allocate blocks in sizes that are powers of 2 (internal fragmentation is possible)**
- **Keep separate free lists for each size**
 - **e.g., separate free lists for 16 byte, 32 byte, 64 byte blocks, etc.**



Buddy System

- If no free block of size n is available, find a block of size $2n$ and split it in to two blocks of size n
- When a block of size n is freed, if its neighbor of size n is also free, combine the blocks in to a single block of size $2n$
- **Buddy** is block in other half larger block



- Same speed advantages as slab allocator



Allocation Schemes

- **So which memory management scheme (K&R, slab, buddy) is best?**
 - **There is no single best approach for every application.**
 - **Different applications have different allocation / deallocation patterns.**
 - **A scheme that works well for one application may work poorly for another application.**



Administrivia

- **We will strive to give grades back quickly**
 - You will have one week to ask for regrade
 - After that one week, the grade will be frozen
- **Regrading projects/exams: possible to go up or down; we'll regrade whole thing**
 - **Beware: no complaints if grade goes down**
- **Others?**



Automatic Memory Management

- Dynamically allocated memory is difficult to track – why not track it **automatically**?
- If we can keep track of what memory is in use, we can reclaim everything else.
 - Unreachable memory is called *garbage*, the process of reclaiming it is called *garbage collection*.
- So how do we track what is in use?



Tracking Memory Usage

- Techniques depend heavily on the programming language and rely on help from the compiler.
- Start with all pointers in global variables and local variables (root set).
- Recursively examine dynamically allocated objects we see a pointer to.
 - We can do this in **constant space** by reversing the pointers on the way down
- How do we recursively find pointers in dynamically allocated memory?



Tracking Memory Usage

- Again, it depends heavily on the programming language and compiler.
- Could have only a single type of dynamically allocated object in memory
 - E.g., simple Lisp/Scheme system with only cons cells (61A's Scheme not "simple")
- Could use a *strongly typed* language (e.g., Java)
 - Don't allow conversion (casting) between arbitrary types.
 - C/C++ are not strongly typed.
- Here are 3 schemes to collect garbage



Scheme 1: Reference Counting

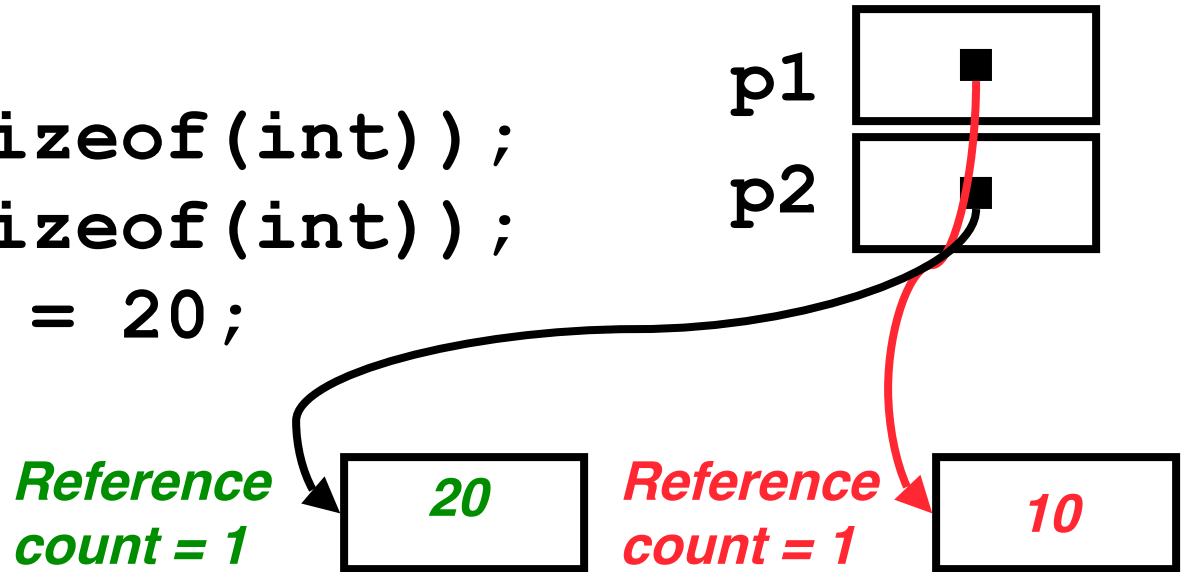
- **For every chunk of dynamically allocated memory, keep a count of number of pointers that point to it.**
- **When the count reaches 0, reclaim.**
- **Simple assignment statements can result in a lot of work, since may update reference counts of many items**



Reference Counting Example

- For every chunk of dynamically allocated memory, keep a count of number of pointers that point to it.
 - When the count reaches 0, reclaim.

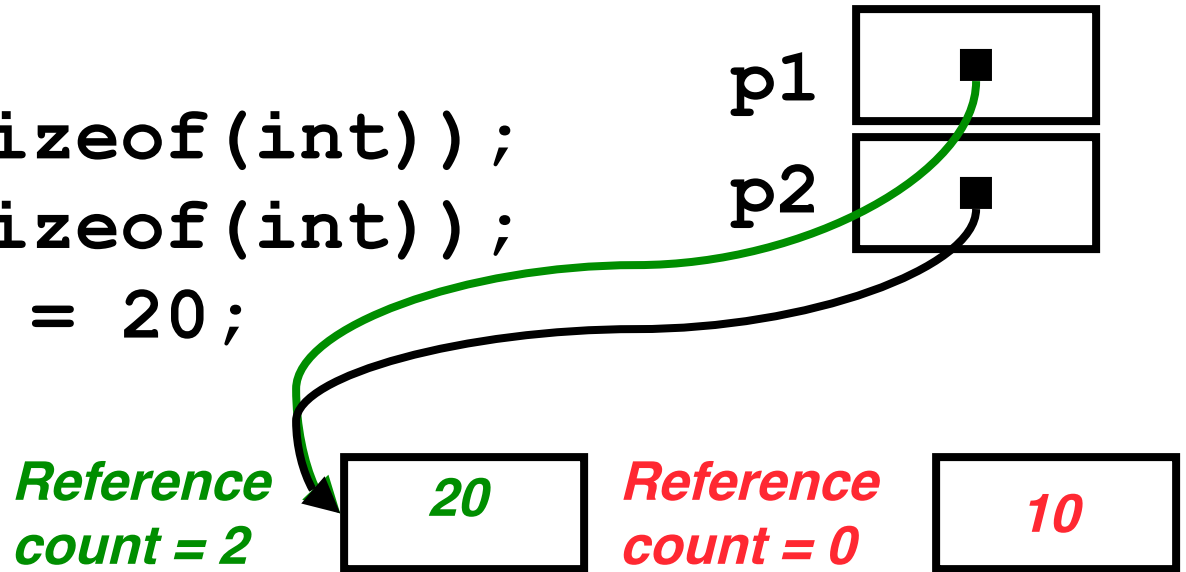
```
int *p1, *p2;  
p1 = malloc(sizeof(int));  
p2 = malloc(sizeof(int));  
*p1 = 10; *p2 = 20;
```



Reference Counting Example

- For every chunk of dynamically allocated memory, keep a count of number of pointers that point to it.
 - When the count reaches 0, reclaim.

```
int *p1, *p2;  
p1 = malloc(sizeof(int));  
p2 = malloc(sizeof(int));  
*p1 = 10; *p2 = 20;  
p1 = p2;
```



Reference Counting (p1, p2 are pointers)


`p1 = p2;`

- Increment reference count for p2
- If p1 held a valid value, decrement its reference count
- If the reference count for p1 is now 0, reclaim the storage it points to.
 - If the storage pointed to by p1 held other pointers, decrement all of their reference counts, and so on...
- Must also decrement reference count when local variables cease to exist.



Reference Counting Flaws

- **Extra overhead added to assignments, as well as ending a block of code.**
- **Does not work for circular structures!**
 - **E.g., doubly linked list:**



Scheme 2: Mark and Sweep Garbage Col.

- **Keep allocating new memory until memory is exhausted, then try to find unused memory.**
- **Consider objects in heap a graph, chunks of memory (objects) are graph nodes, pointers to memory are graph edges.**
 - **Edge from A to B \Rightarrow A stores pointer to B**
- **Can start with the root set, perform a graph traversal, find all usable memory!**
- **2 Phases: (1) Mark used nodes;(2) Sweep free ones, returning list of free nodes**



Mark and Sweep

- **Graph traversal is relatively easy to implement recursively**

```
void traverse(struct graph_node *node) {  
    /* visit this node */  
    foreach child in node->children {  
        traverse(child);  
    }  
}
```

- **But with recursion, state is stored on the execution stack.**

- **Garbage collection is invoked when not much memory left**

- **As before, we could traverse in constant space (by reversing pointers)**



Scheme 3: Copying Garbage Collection

- **Divide memory into two spaces, only one in use at any time.**
- **When active space is exhausted, traverse the active space, copying all objects to the other space, then make the new space active and continue.**
 - **Only reachable objects are copied!**
- **Use “forwarding pointers” to keep consistency**
 - **Simple solution to avoiding having to have a table of old and new addresses, and to mark objects already copied (see bonus slides)**



Peer Instruction

- A. Of {K&R, Slab, Buddy}, there is no best (it depends on the problem).
- B. Since automatic garbage collection can occur any time, it is **more difficult to measure the execution time** of a Java program vs. a C program.
- C. We don't have automatic garbage collection in C because of **efficiency**.

	ABC
1:	FFF
2:	FFT
3:	FTF
4:	FTT
5:	TFF
6:	TFT
7:	TF
8:	TTT




“And in semi-conclusion...”


- **Several techniques for managing heap via malloc and free: best-, first-, next-fit**
 - 2 types of memory fragmentation: internal & external; all suffer from some kind of frag.
 - Each technique has strengths and weaknesses, none is definitively best
- **Automatic memory management relieves programmer from managing memory.**
 - All require help from language and compiler
 - **Reference Count:** not for circular structures
 - **Mark and Sweep:** complicated and slow, works
 - **Copying:** Divides memory to copy good stuff



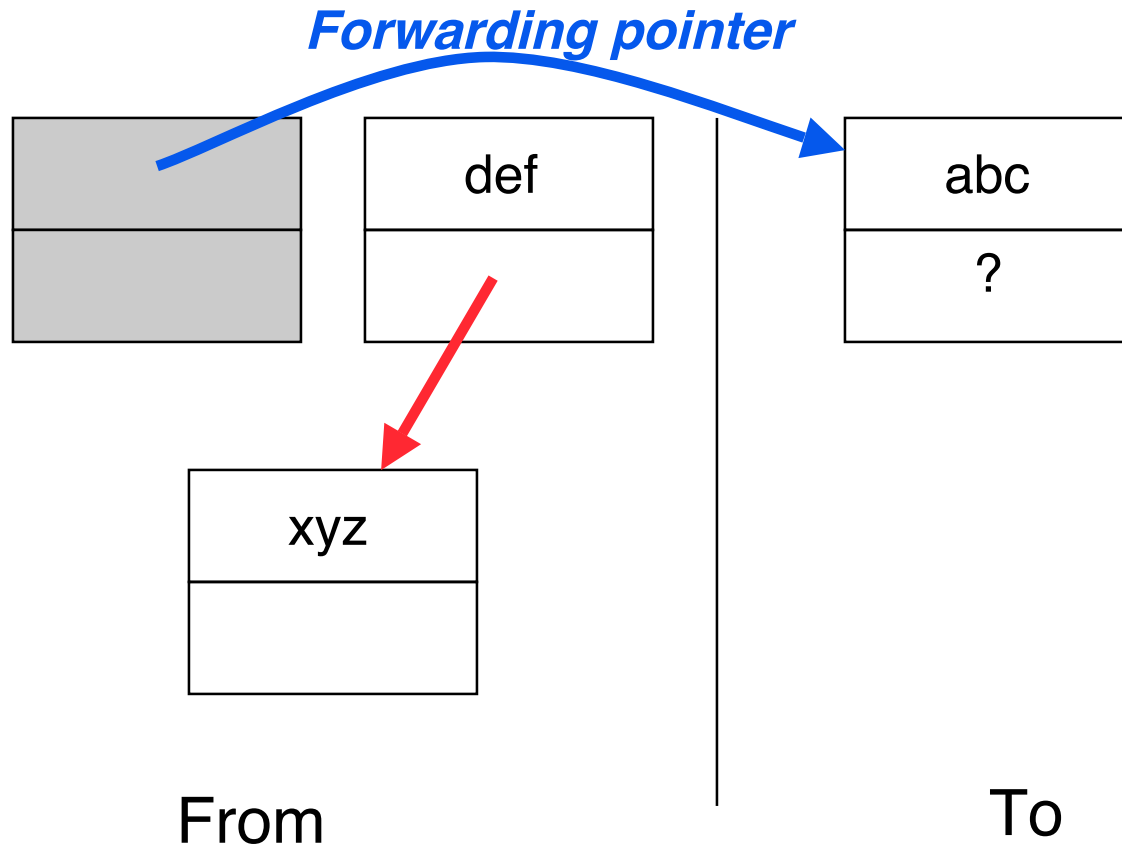
Forwarding Pointers: 1st copy “abc”



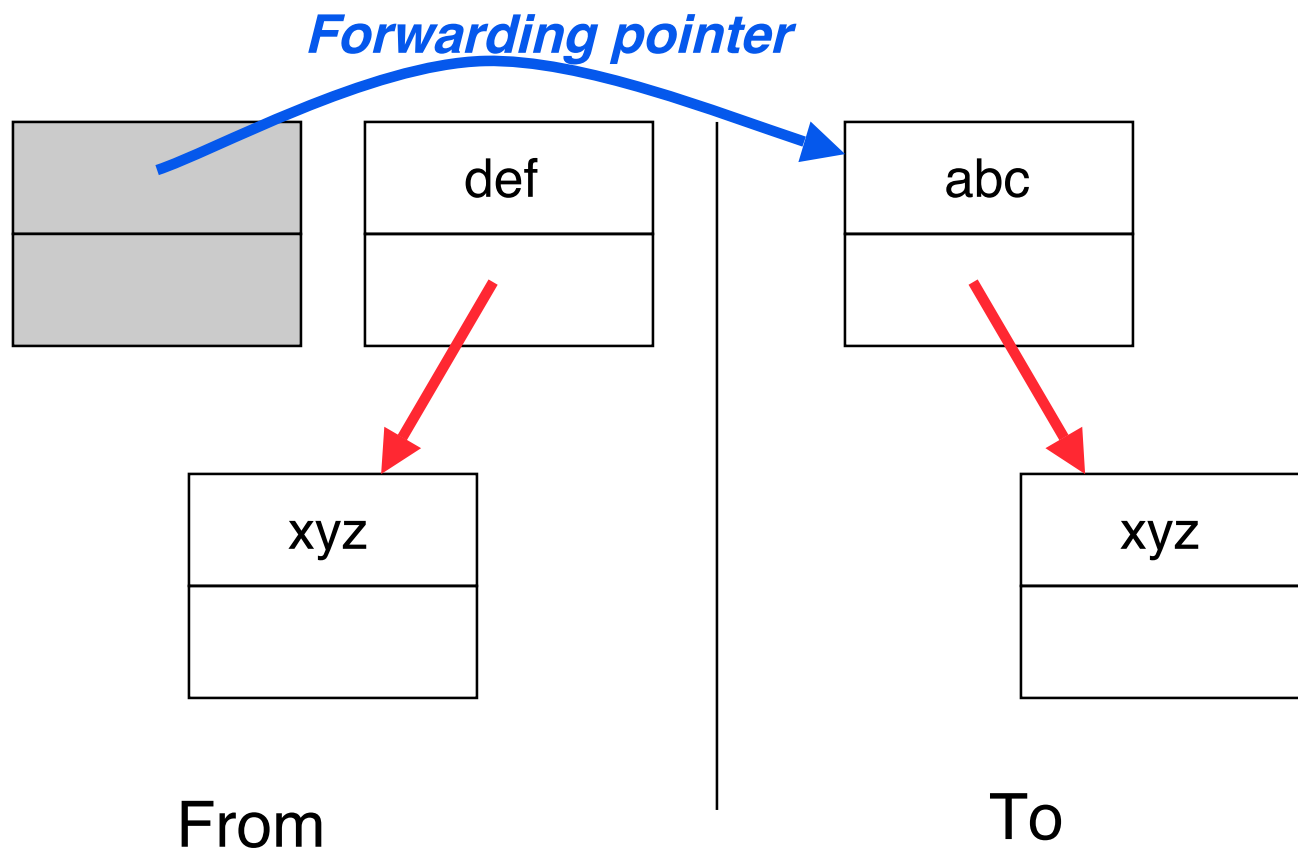
Forwarding Pointers: leave ptr to new abc



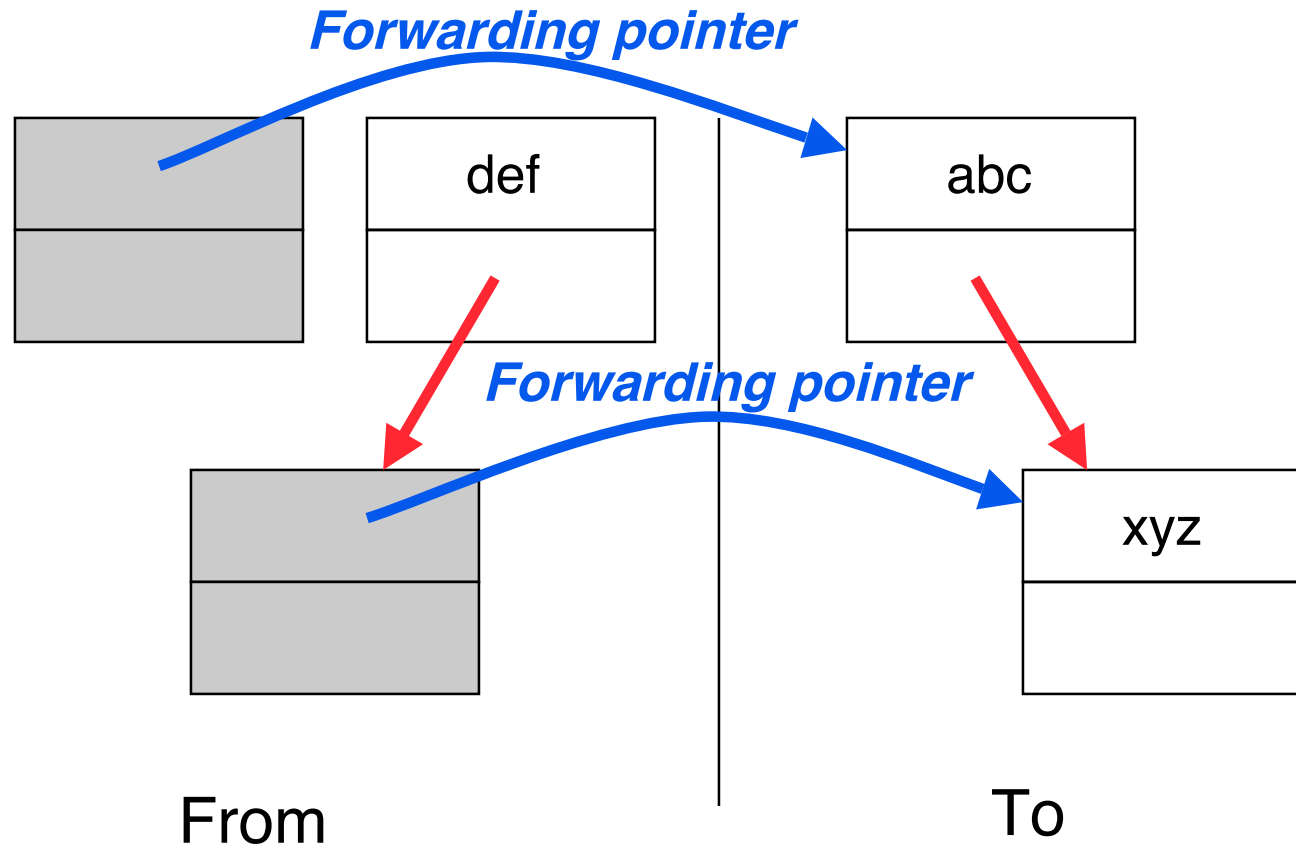
Forwarding Pointers : now copy “xyz”



Forwarding Pointers: leave ptr to new xyz



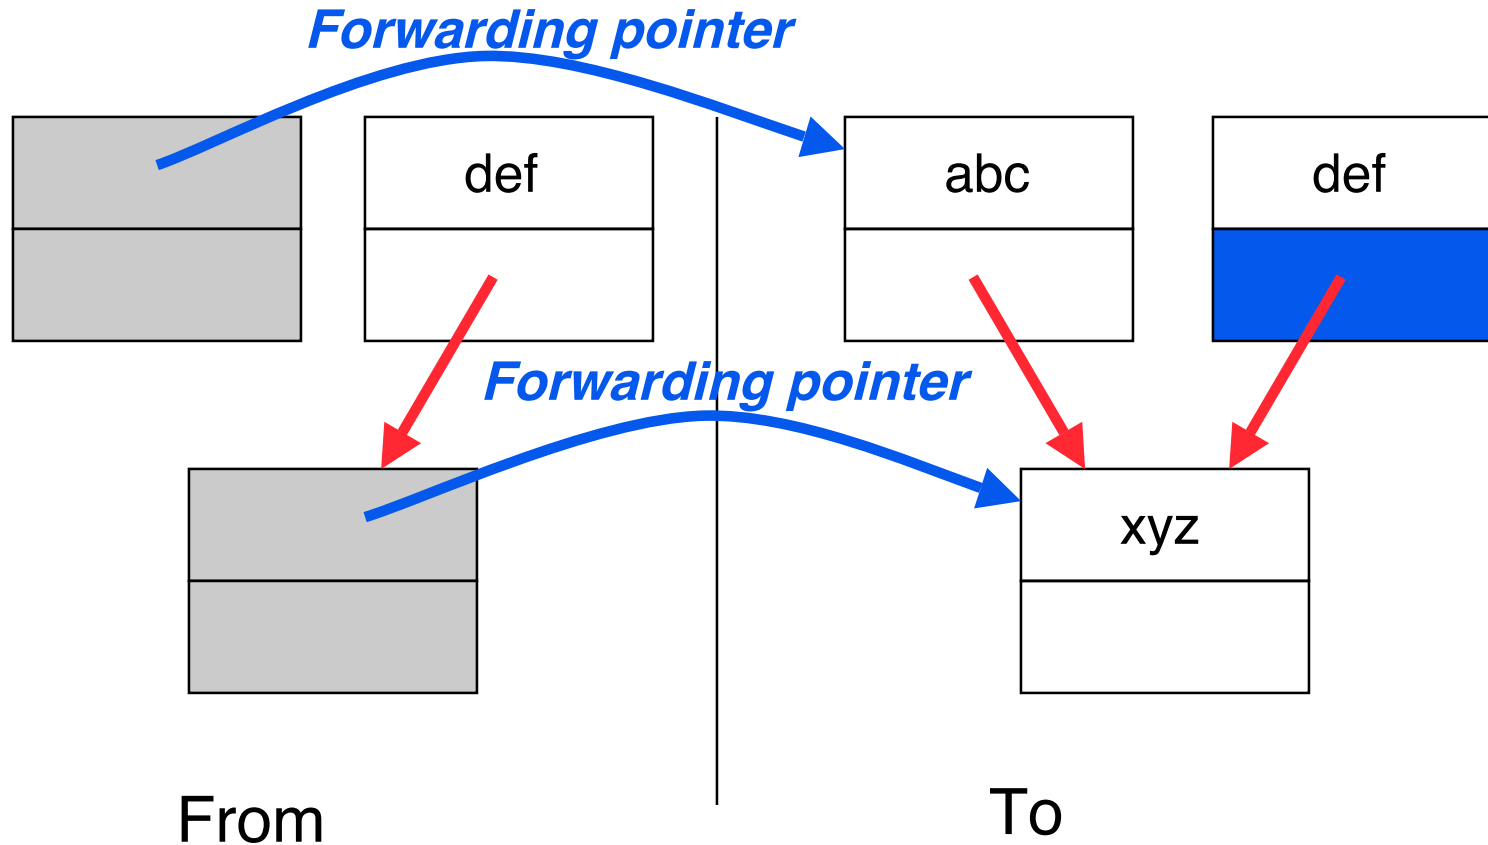
Forwarding Pointers: now copy “def”



*Since xyz was already copied,
def uses xyz's forwarding pointer
to find its new location*



Forwarding Pointers



*Since xyz was already copied,
def uses xyz's forwarding pointer
to find its new location*




Assembly Language

- **Basic job of a CPU: execute lots of *instructions*.**
- **Instructions are the primitive operations that the CPU may execute.**
- **Different CPUs implement different sets of instructions. The set of instructions a particular CPU implements is an *Instruction Set Architecture (ISA)*.**
 - **Examples: Intel 80x86 (Pentium 4), IBM/Motorola PowerPC (Macintosh), MIPS, Intel IA64, ...**



Book: *Programming From the Ground Up*

*“A new book was just released which is based on a new concept - teaching computer science through assembly language (Linux x86 assembly language, to be exact). This book teaches how the machine itself operates, rather than just the language. I've found that the key difference between mediocre and excellent programmers is whether or not they know assembly language. **Those that do tend to understand computers themselves at a much deeper level.** Although [almost!] unheard of today, this concept isn't really all that new -- there used to not be much choice in years past. Apple computers came with only BASIC and assembly language, and there were books available on assembly language for kids. This is why the old-timers are often viewed as 'wizards': they **had** to know assembly language programming.”*



-- slashdot.org comment, 2004-02-05



Instruction Set Architectures

- **Early trend was to add more and more instructions to new CPUs to do elaborate operations**
 - **VAX architecture had an instruction to multiply polynomials!**
- **RISC philosophy (Cocke IBM, Patterson, Hennessy, 1980s) – Reduced Instruction Set Computing**
 - **Keep the instruction set small and simple, makes it easier to build fast hardware.**
 - **Let software do complicated operations by composing simpler ones.**



MIPS Architecture

- MIPS – semiconductor company that built one of the first commercial RISC architectures
- We will study the MIPS architecture in some detail in this class (also used in upper division courses CS 152, 162, 164)
- Why MIPS instead of Intel 80x86?
 - MIPS is simple, elegant. Don't want to get bogged down in gritty details.
 - MIPS widely used in embedded apps, x86 little used in embedded, and more embedded computers than PCs



Most HP LaserJet workgroup printers are driven by MIPS-based™ 64-bit processors.



Assembly Variables: Registers (1/4)

- **Unlike HLL like C or Java, assembly cannot use variables**
 - Why not? Keep Hardware Simple
- **Assembly Operands are registers**
 - limited number of special locations built directly into the hardware
 - operations can only be performed on these!
- **Benefit: Since registers are directly in hardware, they are very fast (faster than 1 billionth of a second)**



Assembly Variables: Registers (2/4)

- **Drawback:** Since registers are in hardware, there are a predetermined number of them
 - **Solution:** MIPS code must be very carefully put together to efficiently use registers
- **32 registers in MIPS**
 - **Why 32? Smaller is faster**
- **Each MIPS register is 32 bits wide**
 - **Groups of 32 bits called a word in MIPS**



Assembly Variables: Registers (3/4)

- Registers are numbered from 0 to 31
- Each register can be referred to by number or name
- Number references:
\$0, \$1, \$2, ... \$30, \$31



Assembly Variables: Registers (4/4)

- By convention, each register also has a name to make it easier to code
- For now:
 - \$16 - \$23 → \$s0 - \$s7
(correspond to C variables)
 - \$8 - \$15 → \$t0 - \$t7
(correspond to temporary variables)
- Later will explain other 16 register names
- In general, use names to make your code more readable



C, Java variables vs. registers

- In C (and most High Level Languages) variables declared first and given a type
 - Example:

```
int fahr, celsius;  
char a, b, c, d, e;
```
- Each variable can **ONLY** represent a value of the type it was declared as (cannot mix and match `int` and `char` variables).
- In Assembly Language, the registers have no type; operation determines how register contents are treated



“And in Conclusion...”

- **In MIPS Assembly Language:**
 - Registers replace C variables
 - One Instruction (simple operation) per line
 - Simpler is Better
 - Smaller is Faster
- **New Registers:**
 - C Variables: $\$s0 - \$s7$
 - Temporary Variables: $\$t0 - \$t7$

