

# *Memory Management*

- ❖ Memory management - allocation of memory chunks to programs and processes
  - 1) Operating system requires 50Mb to store a program to be executed
  - 2) Program requests block of memory for an instance of a structure
  
- ❖ Heap - portion of memory to be managed -  $M[0..N-1]$ .
  - blocks are allocated from the heap
  - when a block has been allocated it is reserved or in use
  - blocks can later be freed or deallocated



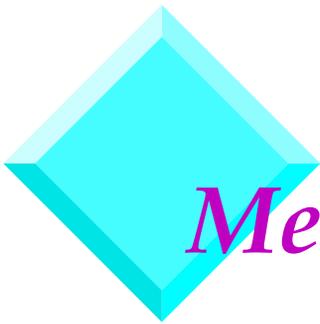
# *Factors determining choice of memory management scheme*

- ❖ Blocks of fixed size versus blocks of various sizes.
  - Some languages (LISP) usually request blocks of fixed size to build lists
  - Operating systems request blocks of various sizes to store programs and their data
  - C++ programs require blocks of different sizes for different structures
- ❖ Linked blocks versus unlinked blocks.
  - If A and B are blocks of allocated memory, and A contains a pointer to B, then B cannot be moved without updating references to B from within A.
    - ◆ Might need to move B to create a block of a large size
  - Leads to dangling pointers



# *Memory management factors*

- ❖ **Small blocks versus large blocks.**
  - Small blocks can be
    - ◆ moved easily
    - ◆ initialized efficiently since time spent is proportional to size of the block
- ❖ **Time versus memory.**
  - Trade-offs between utilization of heap and complexity of managing heap.
- ❖ **Explicit versus implicit release.**
  - Operating system can explicitly release blocks of programs no longer running
  - Many programs return storage implicitly through “garbage collection.”



# *Memory management factors*

- ❖ **Scheduled versus non-scheduled release.**
  - Is information available about the relative order of allocation and deallocation requests?
  - Example: Management of run time stack - blocks are released in a last-in-first-out fashion.
- ❖ **Initialized versus un-initialized blocks.**
  - Initialize large structures to value specified by language semantics or programmer specification
  - Blank out memory so a program cannot read the results of a previous occupant of memory



## *Reference count method - propagating effects of block de- allocation*

- ❖ Reference count - count of the number of pointers that point to a block
  - when a new copy of the address of a block is created, we increment the cell's reference count
  - when a copy is destroyed (say by a pointer assignment) we decrement the reference count
  - when the reference count goes to 0, we de-allocate
  - when block, X, is deactivated, the reference count of any block pointed at by X is decremented.



# *Disadvantages of reference counts*

- ❖ Requires storage in each active block for the reference count
  - ◆ Unless the maximum number of references is known beforehand, can overflow reference count field
- ❖ Doesn't work for circular lists
  - If X and Y point to one another, then even if there are no outside references to either, their reference counts will never go to 0.



## *Mark and sweep garbage collection*

- ❖ Wait for the free list to become empty, and then scavenge memory for unused blocks.
- ❖ Associate a mark field (one bit) with each block. Initially all mark fields are **off**.
- ❖ A marking process will set the mark fields **on** for all blocks that can be reached by some chain of pointers beginning outside of T
- ❖ Memory is then scanned, and all blocks with their mark fields off are collected into the free list. At the same time, the mark fields of used blocks are turned off to prepare for the next round of garbage collection.
- ❖ Big question: How do we do the marking?



# Marking

- Marking algorithms are like tree traversals except
  - 1) blocks can have many pointers and not just 2 - although we will assume they all have the same number
  - 2) blocks may be linked into graphs. But if a block can be reached by more than one path, it is marked the first time it is encountered so that it is not explored when encountered along other paths.
- Let  $C_0, C_1, \dots, C_{k-1}$  be the  $k$  pointers stored in any block
- $\text{Atom}(P)$  is true if  $P$  points outside the heap (probably back to a named program variable).



# *Recursive, depth-first marking*

- ❖ Preorder versions of recursive traversals
  - preorder is important so we mark a block before marking its children - this prevent re-entering circular structures

FullyRecursiveMark(pointer P);

if not Atom(P) and not marked(P) then

    Mark(P)

    for j from 0 to k-1 do

        FullyRecursiveMark( $C_j(P)$ )



# *Recursive, depth-first marking*

- ❖ Preorder versions of recursive traversals
  - preorder is important so we mark a block before marking its children - this prevent re-entering circular structures

FullyRecursiveMark(pointer P);

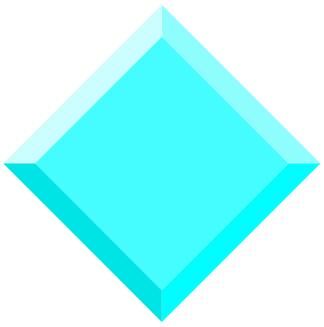
if not Atom(P) and not marked(P) then

Mark(P)

for j from 0 to k-1 do

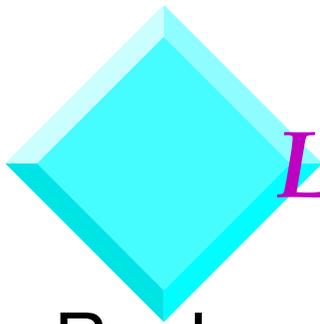
FullyRecursiveMark( $C_j(P)$ )

- ❖ Problem - stack memory needed can be as large as number of links traversed before encountering an atom. But when garbage collection is called, memory is gone!



# *Traversing binary trees*

- ❖ Lindstrom scanning - scan the tree so that every node is visited, although not necessarily in any special order.
- ❖ Importance of traversal algorithms:
  - 1) garbage collection based on ideas in traversal
  - 2) traversal central to many other problems - copying, testing for equality of structures, printing



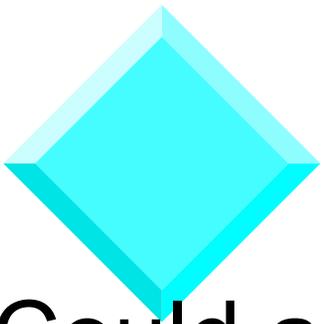
## *Link Inversion Traversals*

- ❖ Replace downward pointers with upward pointers during descent of the tree.
- ❖ Restore pointers to their original values during ascent.
- ❖ Each node has a TAG field that indicates which link field has been reset to point towards its parent. TAG = 1 if the right link has been reset, and is 0 otherwise.
- ❖ Algorithm can be applied to binary trees with shared subtrees (subtrees pointed to by more than one ancestral node in the whole tree).



## *Link inversion for simple lists*

- ❖ Consider a dictionary stored in a singly linked list. Given a word,  $w$ , suppose we want to find the last word in the list,  $L$  that lexicographically precedes it and ends in the same last letter.
  - $L = (\text{cannery}, \text{cat}, \text{chickadee}, \text{coelacanth}, \text{collie}, \text{corn}, \text{cup})$
  - $w = \text{crabapple}$
  - answer = collie
- ❖ Can solve this trivially by traversing the list, remembering the last word we saw that ends in "e" and halting when we find the entry following "crabapple"
- ❖ But we can also solve it if we could walk backwards in the list once "cup" was encountered



## *Link inversion*

- ❖ Could also solve the problem with a stack of pointers
  - as we traverse the list push the pointers onto the stack
  - when we want to back up, start popping the pointers until we reach the record of interest
- ❖ But we can avoid the stack by turning around the pointers in the list itself in a way that lets us restore the list at the end of the algorithm

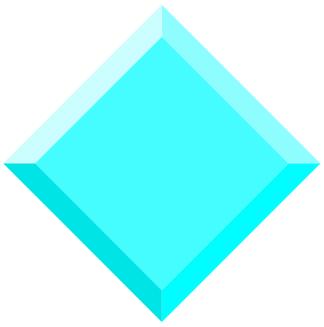


## *Link inversion*

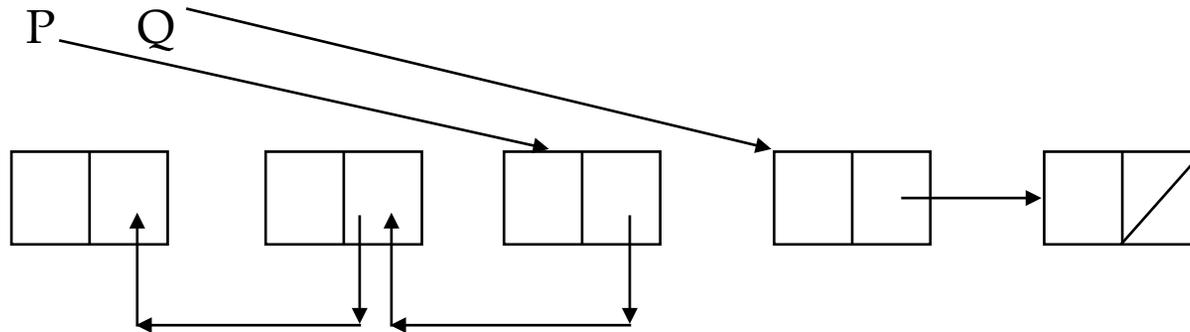
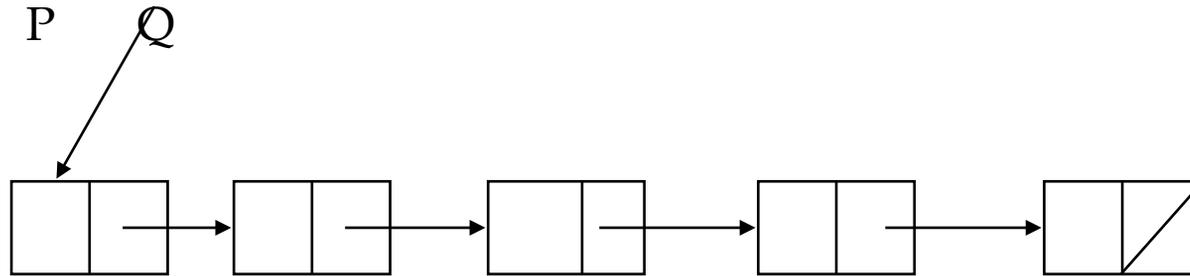
❖ Algorithm uses two pointers

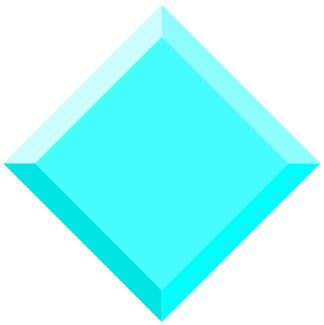
1) Q - points to the first unvisited cell in the remainder of the list. Following link fields from Q will take us in the forward direction to the tail of L from some point on:  $x_i, x_{i+1}, \dots, x_n$ .

2) P - points to the cell containing  $x_{i-1}$ . This cell's link field has been changed to point to  $x_{i-2}$  and so on. Following links from P gets us back to the front of the list.



# *Link inversion*





## *Link inversion*

❖ Following operations can be used to move around the list using link inversion:

1) Start a traversal

$P \leftarrow \text{nil}$

$Q \leftarrow L$  (first node on list)

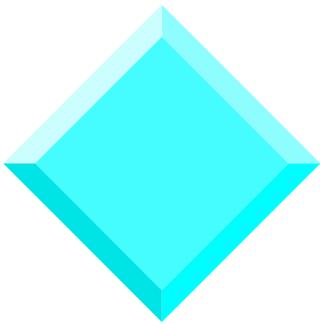
2) Move forward

$P \leftarrow Q$

$Q \leftarrow \text{link}(Q)$

$\text{link}(Q) \leftarrow P$

NOTE: On the left hand side are the new values - all values on the right are old values - simultaneous assignment.



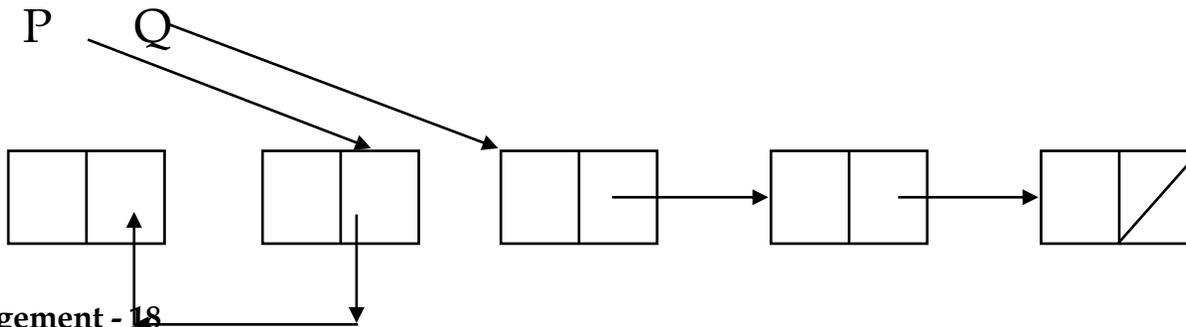
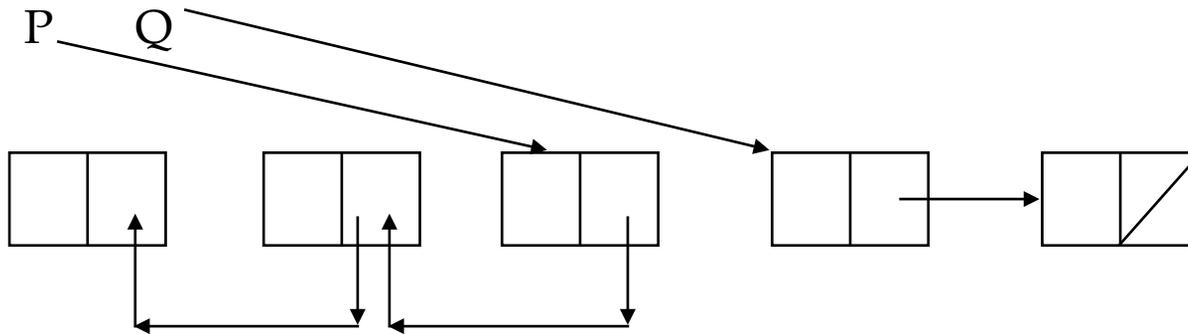
# Link inversion

3) Move backwards in the list; Simultaneous assignment:

$P \leftarrow \text{link}(P)$

$\text{link}(P) \leftarrow Q$

$Q \leftarrow P$





## *Link inversion*

Note:

- 1) restoring the list to its original form involves completely backing out of the list
  - 2) no other use of the list can be accomplished while this traversal is going on - can't be used if several processes require concurrent access to the same data structure
- We'll now generalize this to (almost) stackless traversal of trees



## *Tree traversal with link inversion*

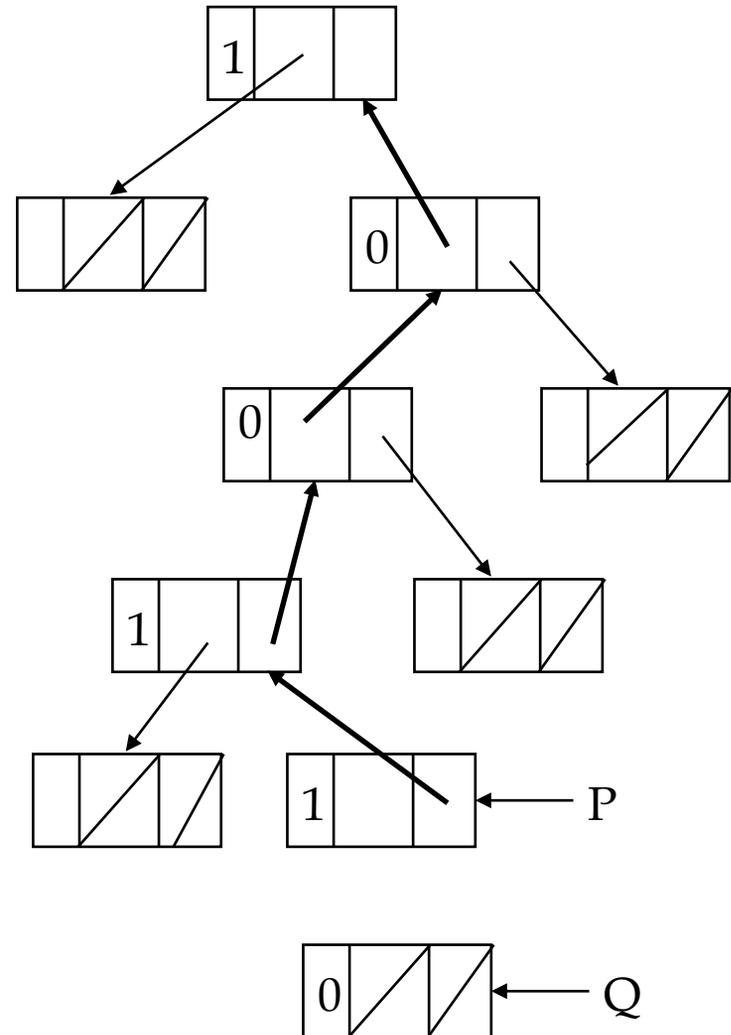
- ❖ Algorithm will descend through the tree by following L and R links
- ❖ The pointer followed will be changed to point to its parent, so the stack used by the recursive algorithm is (mostly) stored in the tree itself
- ❖ When we ascend the tree, we will restore the pointers back to their original values. But how do we know which pointer to restore?
- ❖ Answer: Use a tag bit which will tell us if we changed the L or the R link, and change accordingly.
- ❖ So stack is, effectively, reduced to one bit per stack element.

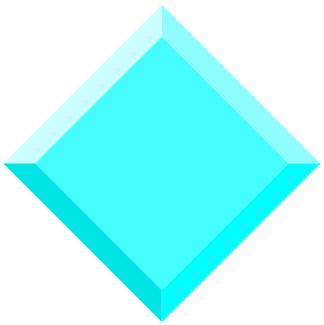


# Link inversion tree traversal

Algorithm uses two variables

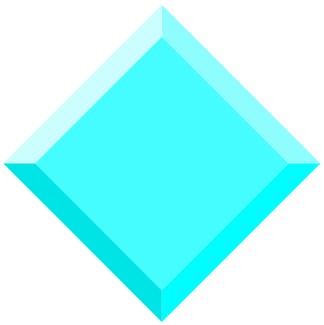
- 1) Q - points to the next **untraversed** tree node. Everything that can be reached from Q looks like the original binary tree
- 2) P - points to the node that is the parent of the one that Q points to. P is the top node on the embedded stack. If we follow pointers back from P (using the tag fields) we get back to the root.





## *Tree link inversion*

- ❖ Algorithm uses four code fragments - descend left or right, and ascend left or right. These cause the P-Q pair to move one edge up or down the tree
  - 1) on descent, the direction indicates which child pointer to follow
  - 2) on ascent, the direction indicates which field of the node to which P points contains the address of that node's parents (use the tag).
- ❖ A sequence descend left-ascend left (or descend right-ascend right) leaves the tree as it was originally.
- ❖ Tag bit is 0 during the visit to the node's left subtree and is 1 during the visit to the nodes right subtree. Tag tells us which link to ascend through.



# Link tree traversal

❖ Code fragments use simultaneous assignment:

1) descend left

$P \leftarrow Q$

$Q \leftarrow \text{Left}(Q)$

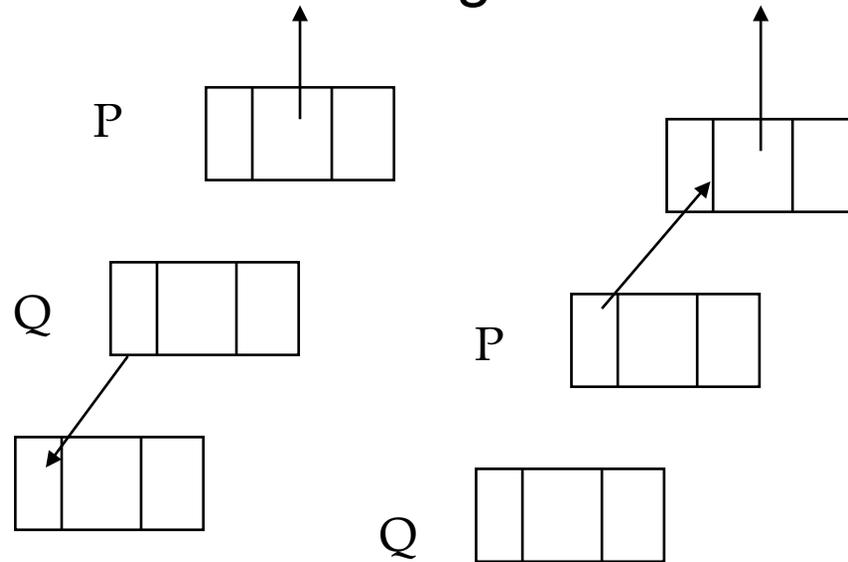
$\text{Left}(Q) \leftarrow P$

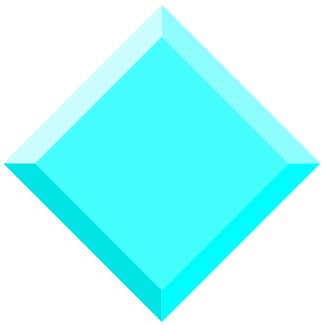
2) Descend right

$P \leftarrow Q$

$Q \leftarrow \text{Right}(Q)$

$\text{Rlink}(Q) \leftarrow P$





## *Link tree traversal*

❖ 3) ascend from left

$Q \leftarrow P$

$P \leftarrow \text{Left}(P)$

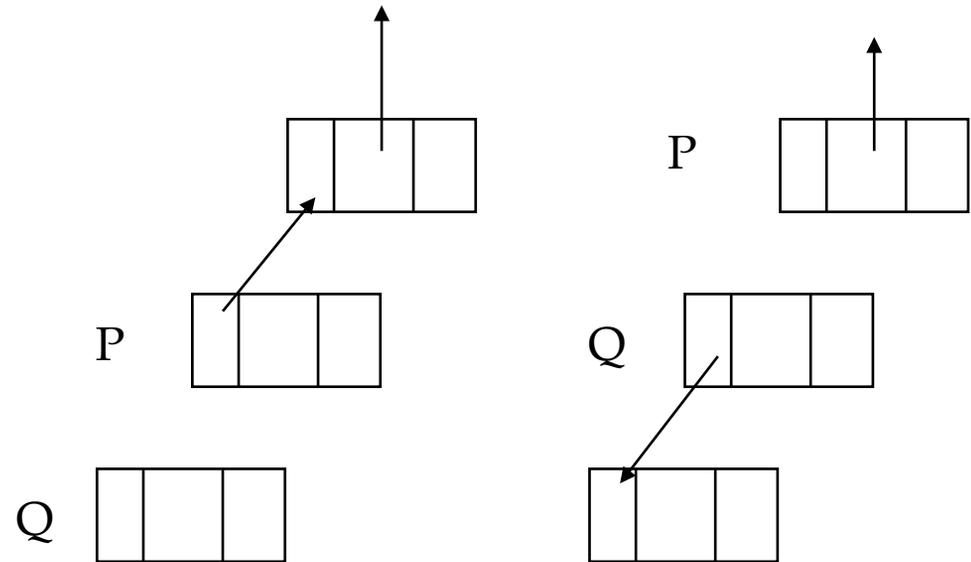
$\text{Left}(P) \leftarrow Q$

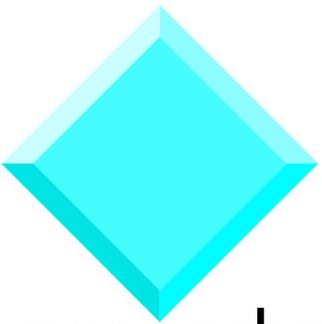
4) ascend from right

$Q \leftarrow P$

$P \leftarrow \text{Right}(P)$

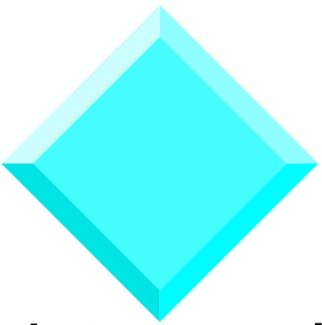
$\text{Right}(P) \leftarrow Q$



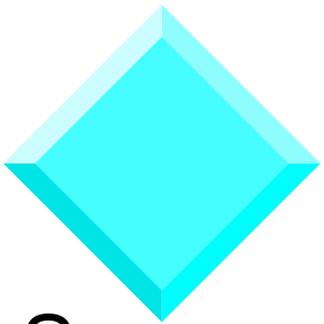


## *Link tree traversal*

```
procedure linkInversionPreOrderTraversal
(pointer Q):
// Initially Q points to the root of the tree}
P <- Nil
// Descend as far as possible to the left
while Q <> Nil do
    visit(Q) // For preorder traversal
    tag(Q) <- 0
    descend_to_left()
```



```
// Ascend as far as possible from the right  
(maybe not at all)  
  while P <> nil and Tag(P) = 1 do  
    ascend_from_right()  
if P = Nil then return  
  else  
    ascend_from_left()  
    Tag(Q) <- 1  
    descend_to_right()
```



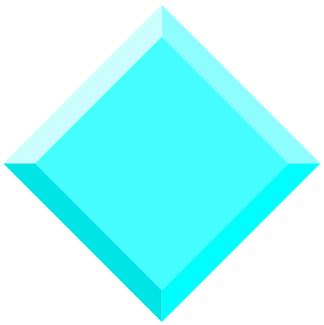
## *Storage requirements*

- ❖ Suppose we have a perfectly balanced binary tree containing  $2^n$  nodes
  - Classic tree traversal algorithm will require a stack of depth  $n$ , with each stack element being  $n$  bits long to represent the address of a node in the tree -  $n^2$  bits in all
  - The in-place traversal requires 1 bit per node as a tag field, or  $2^n$  bits of auxiliary space! Not good.
- ❖ The in-place algorithm doesn't really use a tag field, but a tag bit **stack**, which can never be bigger than  $n$  elements, each of which requires 1 bit.
  - So, auxiliary storage is decreased from  $n^2$  to  $n$ .



# *Robson traversal*

- ❖ No tag bits, no auxiliary stack! (but only for trees)
- ❖ Use the empty link fields of leaf nodes to create the traversal stack
  - since there are plenty of these empty fields, we'll have enough space even for the worst structured trees
  - Left links will be used to chain together the leaf nodes that form the stack
  - Right link fields will point to the interior tree nodes that compose the stack.
  - These will be interior nodes that
    - ◆ have nonempty left subtrees that have been traversed
    - ◆ have nonempty right subtrees that are in the process of being traversed



# *Link-Inversion*

## *Schorr-Deutsch-Waite Marking*

- ❖ Store the stack as a pointer chain in the structure to be marked - algorithm will use  $\text{roof}[\log k] \cdot h$  bits for its "stack", where  $h$  is the length of the longest pointer chain from the root to any cell.
- ❖ Generalization of binary tree traversal with link inversion. Differences:
  - 1) Instead of a tag bit to remember whether we descended through left or right pointer, need a field large enough to store the index  $(0..k-1)$  of the link field traversed. Can be either kept as a field in the cells (bad idea!) or as an auxiliary stack.
  - 2) Ascend and descend code now indexed by pointer



# Ascend and descend code

- descend via  $C_j$   
(simultaneous assignment)

$P \leftarrow Q$

$Q \leftarrow C_j(Q)$

$C_j(Q) \leftarrow P$

- ascend via  $C_j$  (sim. assignment)

$Q \leftarrow P$

$P \leftarrow C_j(P)$

$C_j(P) \leftarrow Q$

procedure

LinkinversionMark( pointer Q):

{Mark all cells reachable from Q}

$P \leftarrow \text{Nil}$ ;

$S \leftarrow \text{MakeEmptyStack}()$

repeat forever

if  $Q \neq \text{Nil}$  and not Atom(Q) and not Marked(Q) then

Mark(Q)

Push(0,S)

descend via  $C_0$

else if  $P = \text{Nil}$  then return

else

$j \leftarrow \text{Pop}(S)$

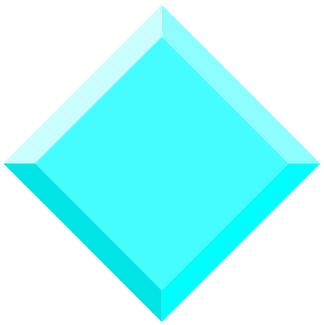
ascend via  $C_j$

$j \leftarrow j+1$

If  $j < k$  then

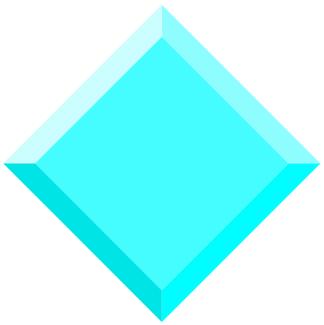
Push(j,S)

descend via  $C_j$



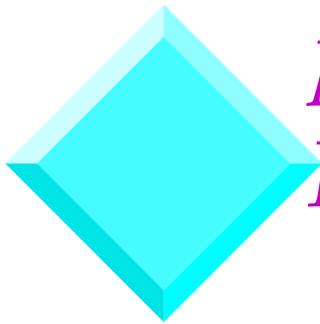
## *Other garbage collection algorithms*

- ❖ Link inversion algorithms that take constant space, independent of the length of the longest path from the root to any cell.
  - still use a mark bit in each cell
  - require much more computation than stack-based link inversion
- ❖ Copying algorithms - break memory into two parts, and allocate storage from one. When that segment of memory is fully utilized, copy and compact non-garbage cells from one area to the other.
  - algorithms do not touch garbage
  - but must have a lot of memory



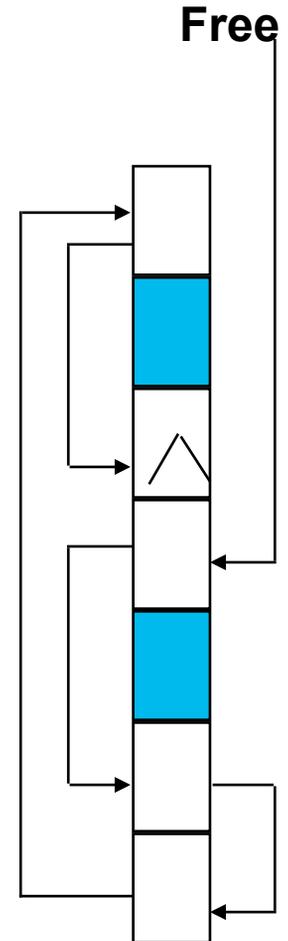
## *Cautions on garbage collection*

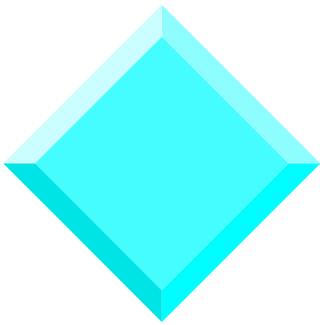
- ❖ System performance usually degrades quickly once garbage collection is required
  - not much available storage is really available, so system spends a lot of time marking or copying to reclaim only a small amount of storage
  - almost immediately have to garbage collect again, as more memory is consumed
- ❖ Incremental garbage collection
  - spend a fixed amount of time collecting garbage
  - therefore, collect only a little garbage at a time.



# Managing the Heap - Records of a single size

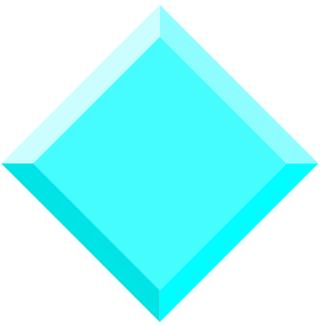
- ❖ Suppose all requests are for blocks of size  $k$
- ❖ Partition heap once into  $\text{floor}[n/k]$  blocks of size  $k$ . Let  $T[0..m-1]$  be those blocks, where  $m = \text{floor}[n/k]$
- ❖ At any time, some blocks in  $T$  are free and some are in use.
- ❖ Storage management
  - Maintain a free list of blocks in  $T$  - arrange free blocks in a singly linked list
  - List is accessed as a stack - allocation is a POP and deallocation is a PUSH. So, allocation takes  $O(1)$  time.
  - How can we determine when to deallocate implicitly?





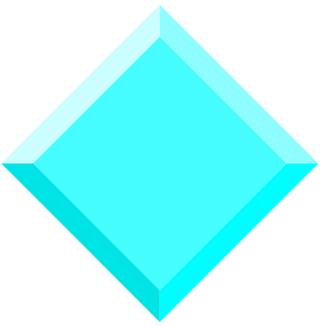
## *Records of various sizes*

- ❖ If blocks can be of different sizes, then we need a more complex memory management mechanism.
- ❖ **Method 1:** Allocate in address order from a single block of large memory.
  - when heap is exhausted, some blocks may have been de-allocated
  - compress them to get one big block, and continue



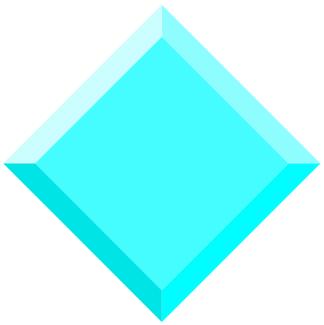
## *Records of various sizes*

- ❖ **Method 2:** Keep a single list of blocks of all sizes.
  - If a request for a block of size  $k$  is encountered, find a block of size  $h > k$ , allocate  $k$  words and return a block of size  $h-k$  to free storage. The set of blocks is called the storage pool.
- ❖ **Two main problems:**
  - 1) the leftovers ( $h-k$ ) might be too small for anything
  - 2) when a block is deallocated, it may need to be merged with an adjacent free block



## *Records of various sizes*

- ❖ **Method 3:** Keep several pools of free blocks, according to size.
  - allocate by choosing a block from the correct pool
  - need a subdivision strategy that allows easy construction of large(st) blocks as memory is de-allocated.
  - "Buddy system"



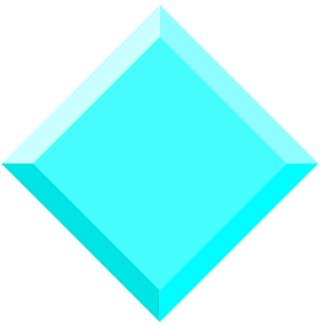
## *Method 1: Compression*

### ❖ Major problems:

- 1) time it takes to move blocks within memory
- 2) adjustment of pointers to blocks that have moved.

### ❖ Algorithm assumptions:

- 1) The heap  $M[0..N-1]$  is partitioned into blocks, some of which are free and others in use.
- 2)  $\text{Marked}(P)$  is true if  $P$  points to a block that is in use, and false otherwise
- 3)  $\text{First}$  is the address of  $M[0]$  and  $\text{Last}$  is address for  $M[N-1]$ .
- 4) Each block has a size field; the block after the one beginning at address  $P$  begins at address  $P + \text{Size}(P)$



# Compression

## ❖ Assumptions

5) Pointer fields of a block in use are known in advance (for updating references)

6) Each block in use contains a pointer field forwardingaddress that is reserved for use by the compaction algorithm. The value of this field will be computed before any blocks are moved.

## ❖ Algorithm will use the following abbreviation:

for each block P, in address order, do for

P <-- First

while P <> Last do

\*\*\*\*\*

P <-- P + size(P)



# Compaction algorithm

```
procedure Compact;  
{Compact blocks into low memory  
addresses, adjusting pointers; marking of  
blocks in use has already been done}  
{Compute forwarding addresses}  
Dest <-- First  
for each block, P, in a-order, do  
  if Marked(P) then  
    ForwardingAddress(P) <-- Dest  
    Dest <-- Dest + Size(P)  
{Adjust internal links}  
for each block, P, in a-order, do  
  if Marked(P) then  
    for each pointer field Link do  
      Link(P) <--  
        ForwardingAddress(Link(P))
```

```
{move the cells}
```

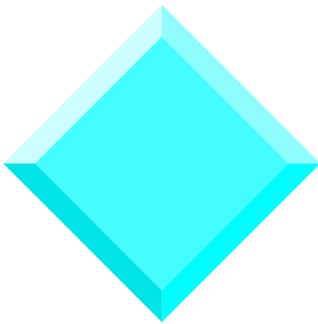
```
for each block, P, in a-order, do
```

```
  if Marked(P) then
```

```
    Copy Size(P) bytes beginning at  
    P to ForwardingAddress(P)
```

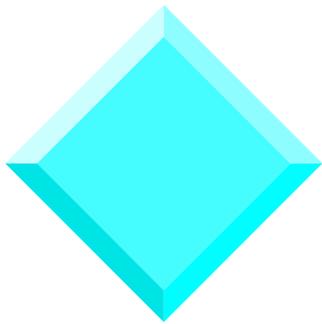
## •Observations:

- 1) algorithm is linear in size of memory compacted
- 2) but, practically, it takes three scans over the memory to complete (four counting the initial marking)



## *Method 2: Blocks of various sizes*

- ❖ If blocks are of different sizes, and some blocks are very large, compaction is not a good solution because it might get called too often.
- ❖ Assume explicit deallocation
- ❖ Strategy for selecting blocks, subdividing them and recombining them determined by:
  - Memory utilization - do not fail to satisfy a request if the aggregate amount of memory in use is only a small percentage of the heap size
  - Memory overhead - the memory manager data structures should be small compared to the heap
  - Time efficiency - allocation and deallocation should be fast



## *It's not a perfect world*

- ❖ Not all sequences of requests for memory will be satisfied by any memory management algorithm, even if it is possible to satisfy them by some variation on the algorithm.

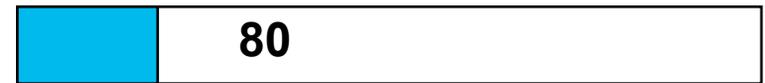
- ❖ Example: Heap of size 100

- $B_1$   $\leftarrow$  allocate (20)

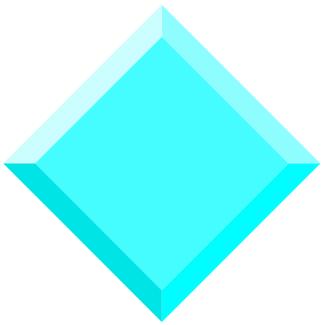
- $B_2$   $\leftarrow$  Allocate (40)

- Free( $B_1$ )

- $B_3$   $\leftarrow$  Allocate (50)

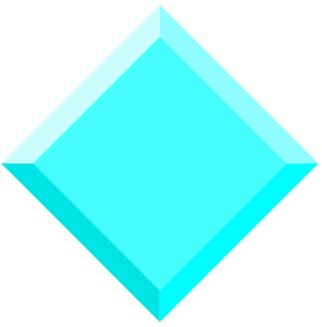


- ❖ If we had allocated  $B_2$  from the end, then we could have satisfied the last request
  - But, generally, deciding if a (known) sequence of requests is satisfiable is NP-complete



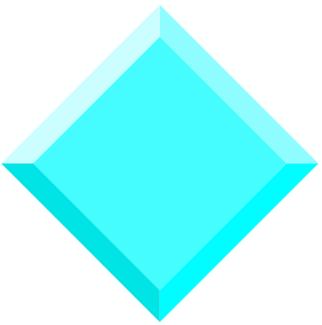
# *Fragmentation*

- ❖ Two blocks of size  $s_1$  and  $s_2$  are never more useful than one larger block of size  $s_1 + s_2$ .
- ❖ Allocation strategy should try to minimize the number of free blocks, and maximize their sizes
- ❖ External fragmentation - splitting of free storage into a relatively large number of relatively small free blocks.
  - External fragmentation can be combatted by not subdividing a free block if the leftover space is too small to be useful - just allocate the larger block
  - smaller block would clutter free list.



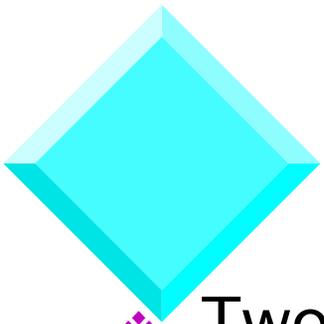
# *Fragmentation*

- ❖ Internal fragmentation - distribution of significant amount of free storage in allocated blocks
  - always some "wasted" space because blocks require fields for use by memory manager - size, pointers to other blocks, etc.



## *Summary*

- ❖ Maintain a pool of free blocks
- ❖ Develop an allocation strategy for assigning a block to a request
  - retain unused portion of block for use by other requests
  - must decide on a data structure for the pool of free blocks. Can be organized either by size or by location or by both.
- ❖ Develop a freeing strategy for returning a freed block to the pool, and coalescing it with other free blocks to form largest possible free blocks



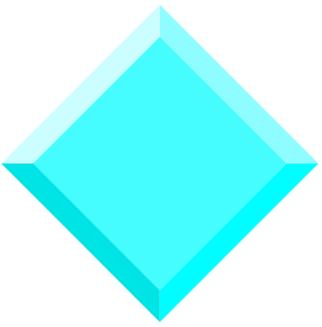
# *Allocation strategies*

## ❖ Two basic strategies:

- 1) **best fit** allocates a request for a block of size  $n$  by finding the smallest block larger than  $n$ .
- 2) **first fit** allocates using the first block found of size at least  $n$ .

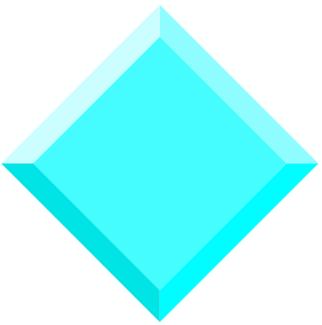
## ❖ Best fit

- implement pool of free blocks as a linked list ordered by increasing size
- first block on this list whose size is  $\geq n$  is the best fit
- on average, half the list is scanned to find the best fit.
- must scan list again to put leftover block back into free list.
- if list is maintained in address order, then entire list would have to be scanned to find the best fit, but reinsertion of leftover fragment is avoided.



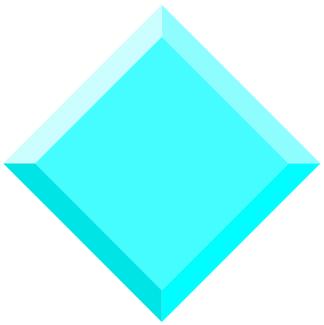
## *Allocation strategies*

- ❖ Use a linked list representation of the pool of free blocks
- ❖ Problem - blocks near the beginning of the list get subdivided more often than blocks near the end
- ❖ **Solution:** The **roving pointer** - begin search for a block at the point where the previous search ended.
  - Small blocks will not accumulate at any one area of the list.
  - A block that has just been subdivided, or was too small to fulfill the previous request, will not be considered again until after all blocks have been examined. This gives it time to coalesce with freed blocks



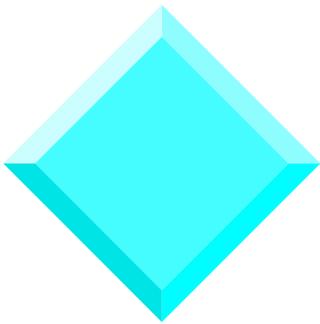
## *Data structures for freeing*

- ❖ What fields of information would we need to implement block deallocation and coalescing efficiently?
  - 1) Mark field = 0 for a free block and 1 for a block in use
  - 2) Size field
  - 3) Next and Prev fields to maintain a doubly linked list of free blocks.
- ❖ Only needed when the block is free - process can use the fields when the block is in use
- ❖ Makes subdivision easy - allocate from the end and set  $\text{Size} \leftarrow \text{Size} - n$



## *Data structures for freeing*

- ❖ Do not depend on assumptions about order of blocks in free list to support the Free operation
  - Either size order or address order will entail search of the free list, and we want to avoid this search
- ❖ Suppose we free the block B.
  - To determine if the next higher (in address) block is free, we look at  $\text{Mark}(B + \text{Size}(B))$
  - Determining if the previous block is free is harder
  - Replicate the mark field at the end of each block
  - So, can examine a field at a fixed displacement before the beginning of block B to determine if the previous block is free
  - This is called  $\text{LowMark}(B)$ .



# Data structures for freeing

- ❖ If B's lower neighbor is free, then its Size field is also replicated near its end
  - Call this field LowSize(B).
- ❖ These fields are stored in the block prior to B, and not in B itself

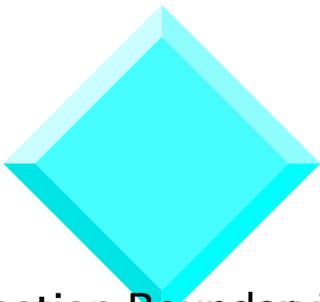
Low Mark	Low Size
Mark =1	Size
LowM = 1	Low Size
Mark	Size

**Block in use**

Low Mark	Low Size
Mark =0	Size
Next	
Prev	
LowM = 0	Low Size
Mark	Size

**Free block**

- This is called the Boundary Tag Method



# Allocation algorithm

```
function BoundaryTagAllocate (integer:n):pointer
{return address of a free block of size n, or nil if no large enough block
is available}
```

```
  SaveRover <-- Rover
```

```
  repeat
```

```
    m <-- Size(Rover)
```

```
    if m < n then
```

```
      Rover <-- Next(Rover)
```

```
    else
```

```
      if m-n < delta then {almost exact match, no
```

```
      leftover}
```

```
        P <-- Rover
```

```
        DoublyLinked Delete (P)
```

```
      else {inexact match - divide the block}
```

```
        P <-- Rover + m - n
```

```
        Size(Rover) <-- LowSize (P) <-- m-n
```

```
        Size (P) <-- LowSize (P+n) <--n
```

```
        LowMark(P) <-- 0
```

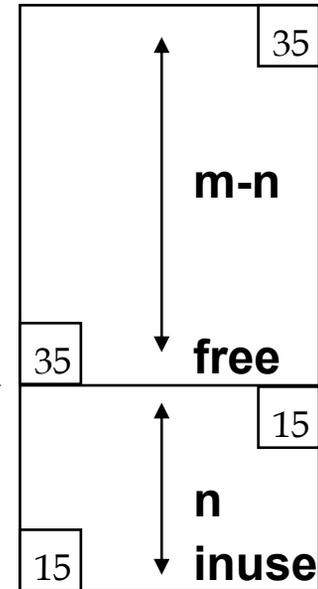
**Rover**

10

**P** →

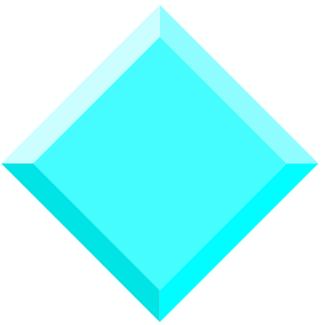
45

60



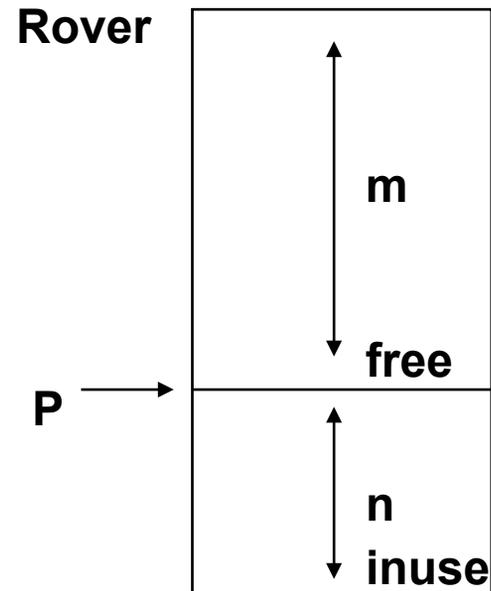
m = 50

n = 15



## *Allocation algorithm*

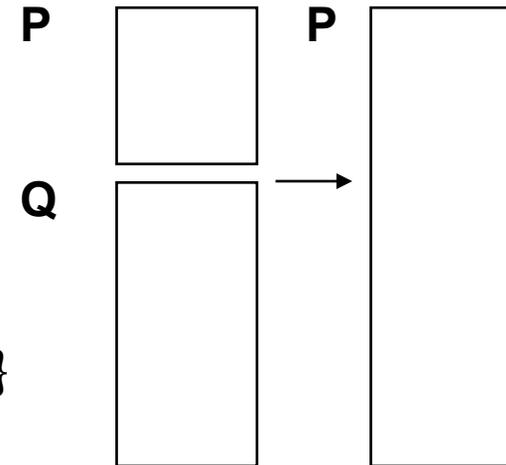
```
Mark (P) <-- 1  
LowMark (P + Size(P)) <-- 1  
Rover <-- Next (Rover)  
Return (P)  
until Rover = SaveRover  
return nil
```





# Freeing algorithm

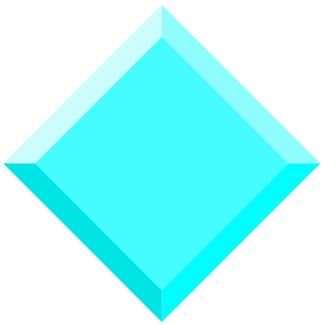
```
procedure BoundaryTagFree (pointer P)
{deallocate the block pointed to by P}
if LowMark (P) = 0 then {Preceding block is free, merge}
    Q <-- P
    P <-- P - LowSize (P)
    Size (P) <-- Size (P) + Size (Q)
else { Preceding block in use; mark this block as free}
    Mark (P) <-- 0
Q <-- P + Size (P) { Q <-- address of subsequent block}
if Mark(Q) = 0 then {merge with following block}
    Size (P) <-- Size (P) + Size (Q)
    if Rover = Q then Rover <-- P
    DoublyLinkedDelete (Q)
Q <-- P + Size (P) {P points to block to be
placed in free; Q points just after the end of this block}
```



```
LowSize(Q) <--
Size(P)
```

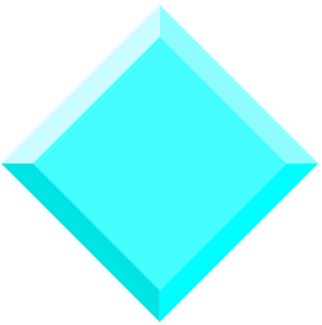
```
LowMark(Q) <-- 0
```

```
DoublyLinkedInsert(P,
Free)
```



## *Method 3: Buddy systems*

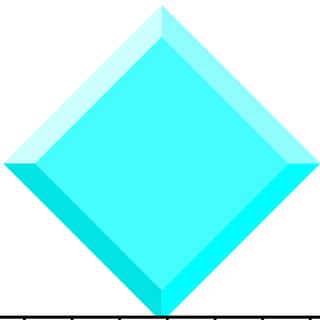
- ❖ Trade external fragmentation for internal fragmentation
  - only blocks of small number of fixed sizes are allocated. Additionally, the starting addresses for the blocks are constrained.
  - a request for a block of size  $m$  is rounded up to the smallest size  $n$  that is supported
- ❖ When a block is freed, it may not be merged with an adjacent free block
  - every block has a "buddy block" of the same size
  - when a block is freed, if its buddy is free, it is merged with its buddy to create a larger block
  - this larger block may be merged with its buddy ...



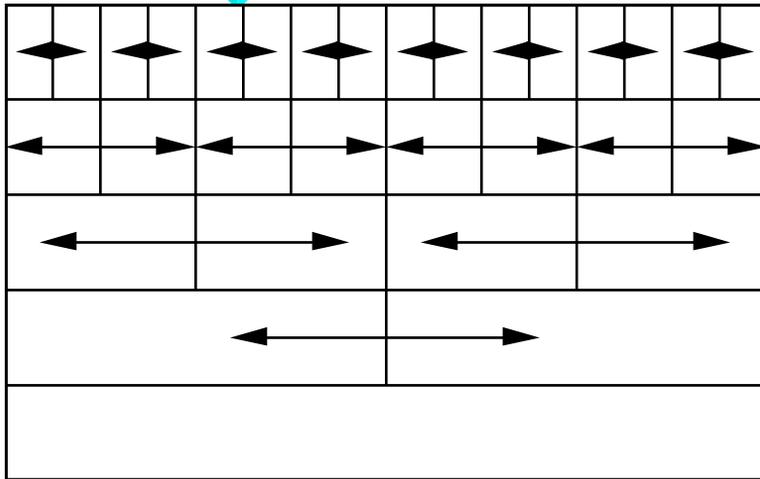
## *Binary buddies*

- ❖ All allocatable blocks have sizes that are powers of 2.
- ❖ Heap size is also a power of 2, say  $2^m$ .
- ❖ The allowable block sizes are  $2^m, 2^{m-1}, \dots, 2^1, 2^0$ 
  - this means that there are  $m+1$  Free lists, with list headers  $\text{Free}[m], \text{Free}[m-1], \dots, \text{Free}[0]$ .
- ❖ Each block of size  $2^k$  begins at an address which is a multiple of  $2^k$  - begins at  $p \cdot 2^k$ 
  - ends at  $(p+1) \cdot 2^k - 1$
  - $0 \leq p < 2^{m-k}$





# Binary buddies



❖ If block N of size  $2^k$  begins at address P, then B's buddy begins at either

–  $P - 2^k$  or  $P + 2^k$

❖ Easy to tell by looking at k'th bit of P. If k'th bit is 1, then buddy is at  $P - 2^k$ . If k'th bit is

• Block of size 4 beginning at  $P=4 = 0100$ , buddy is at  $P+2^k$

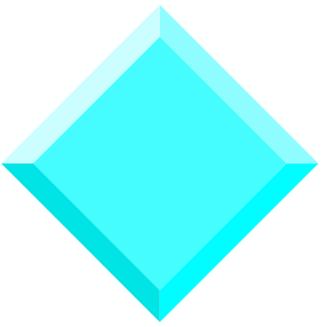
0100. k'th bit = 1 (start from right at 0, so buddy is at 0000)

• Block of size 2 starting at  $P=10 = 1010$  has buddy at  $1000 = 8$ .

• Buddy of a block of size  $2^k$  at address P is the block of size  $2^k$  beginning at  $P \oplus 2^k$

– So just **complement** bit k in the binary representation of B's address

– But this is just the **exclusive or** of the block's position and its size!



## *Binary buddies*

- ❖ Free lists are kept as doubly linked to facilitate deletions
- ❖ Fields are same as boundary tag method, except you don't need the extra copy of the mark and size fields

mark = 1	size

mark=0	size
next	
prev	



## *Binary buddies - allocation*

**function** BinaryBuddyAllocate (integer n): pointer

{Return pointer to a block whose size is the next power of 2 > n}

{ Return nil if no sufficiently large block exists}

j <-- k <-- roof(log n)

**while** j <= m and IsEmptyList (Free[j]) **do** j <-- j+1

**if** j > m **then return** nil

P <-- Free[j]

DoublyLinkedDelete (P)

Mark(P) <-- 1

**while** j > k **do**

Size(P) <-- k

    j <-- j-1

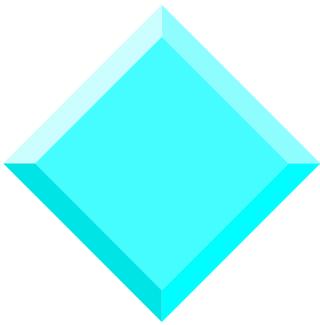
**return** P.

    Q <-- P + 2<sup>j</sup>

    Mark(Q) <-- 0

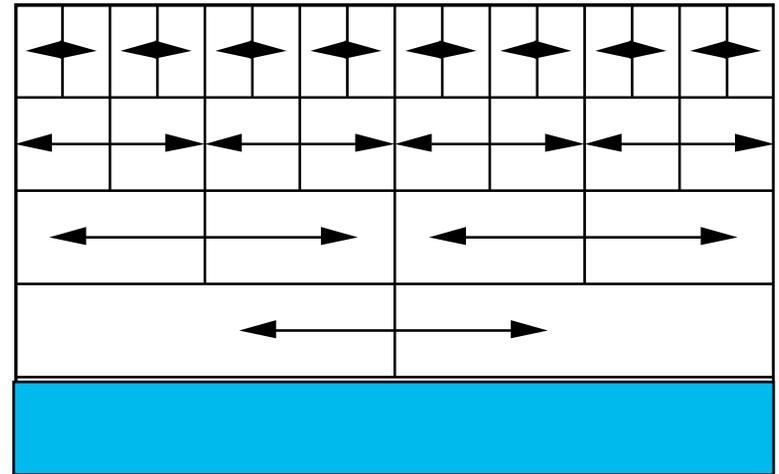
    Size(Q) <-- j

    DoublyLinkedInsert (Q, Free[[j]]) {Inserts Q onto an empty list Free[j]}

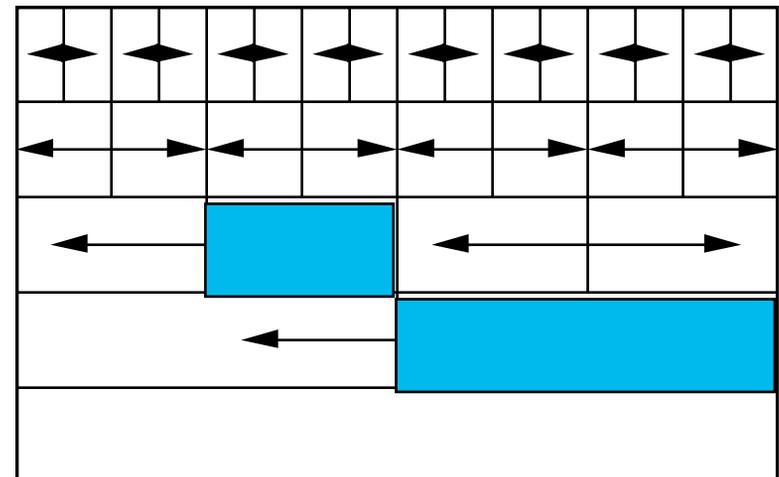


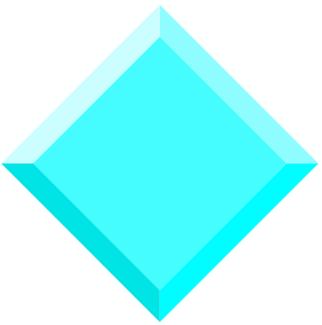
## Example

- ❖  $\text{Free}[4] = 0$
- ❖  $\text{Free}[3] = \text{Free}[2] = \text{Free}[1] = \text{Free}[0] = \text{NIL}$
- ❖ Allocate a block of size 3
- ❖ Fill with a block of size 4
  - $j \leftarrow k \leftarrow 2$
  - $j$  gets bumped up to 4 and the block of size 16 starting at  $P = 0$  is used



- Place the block of size 8 starting at position 8 ( $0 + 2^3$ ) into  $\text{Free}[3]$
- Place the block of size 4 starting at location 4 ( $0 + 2^2$ ) into  $\text{Free}[2]$





## *Block freeing*

- ❖ May reverse the splits that occurred during allocation
  - for example, if we allocate a block of size  $k$  and then immediately free it, we restore free lists to their original state.
- ❖ If the freed block's buddy is free, it is removed from free list and merged to create a larger block.
- ❖ This is repeated until either the buddy is in use, or the block becomes the entire heap.
- ❖ Can have at most  $\log(n)$  frees, so freeing is  $O(\log n)$  in the worst case.



# Block freeing

```
procedure BinaryBuddyFree (pointer P):  
{Free the block of size  $2^k$  beginning at P}  
k  $\leftarrow$  Size(P)  
while k < m and Mark(P EOR  $2^k$ ) = 0 and  
  Size(P EOR  $2^k$ ) = k do  
  {P's buddy is free, so merge}  
  Q  $\leftarrow$  P EOR  $2^k$   
  DoublyLinked Delete (Q)  
  if Q < P then P  $\leftarrow$  Q  
  k  $\leftarrow$  k + 1  
Mark(P)  $\leftarrow$  0  
Size (P)  $\leftarrow$  k  
DoublyLinkInsert (P, Free[k])
```

P

