

# Reference Counting

Kevin Köster

Uni Hamburg

31. März 2013



## Frage

Wann wissen wir, ob ein Objekt freigegeben werden kann?

```
1 void foo() {  
2     char* data = malloc(1024);  
3  
4     calc_unblocked(data);  
5  
6     calc_blocked(data);  
7  
8     // free?!?  
9 }
```

## Möglichkeiten

- Garbage Collection
- Reference Counting

# Garbage Collection

```
1 public static void foo(){  
3     ArrayList<String> liste = new ArrayList<String>(); //  
        heap  
5     bar(liste);  
}
```

Listing 1: GC

# Garbage Collection

## Eigenschaften

Überwacht alle Objekte und löscht diese, wenn nötig.

# Garbage Collection

## Vorteile

- Einfach (Benutzer)
- Sicher

## Nachteile

- Aufwändige Implementierung
- Kostet Rechenleistung
- Zufällig



# Reference Counting

## Eigenschaften

Zählung der Referenzen auf ein bestimmtes Objekt.

# Reference Counting

## Vorteile

- Leichte Implementierung
- Schnell

## Nachteile

- Aufwändig für Benutzer
- Fehleranfällig
- Zyklen


# Zyklen

## Problem


Zwei Objekte referenzieren sich gegenseitig

# Zyklen

Refcount:  
1




Refcount:  
1



# Zyklen


Refcount:  
2

Refcount:  
2



# Zyklen

Refcount:  
1



Refcount:  
1



# Inode

## Inhalt eines Inode

- Zugriffsrechte
- Eigentümer und Gruppe (UID und GID)
- Dateityp
- Größe
- Änderungsdaten
- Referenzzähler

# Inode

Inode 15728



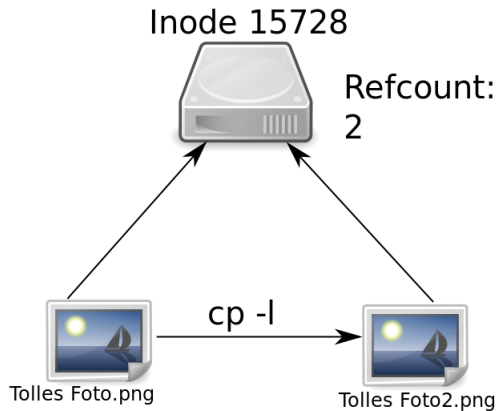
RefCount:  
1



Tolles Foto.png



# Inode



# Inode

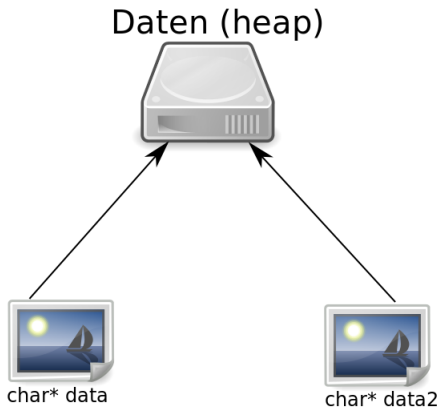
## Referenzzähler

- Erhöht sich für jede Referenz um 1
- Jeder Hardlink eine Referenz (Nicht unterscheidbar)
- Kein direktes Löschen nur unlink

# stat

```
root@kevin-ubuntu: ~  
root@kevin-ubuntu ~ # stat /  
Datei: »/“  
Größe: 4096          Blöcke: 8          EA Block: 4096  Verzeichnis  
Gerät: 808h/2056d   Inode: 2          Verknüpfungen: 25  
Zugriff: (0755/drwxr-xr-x) Uid: (  0/  root)  Gid: (  0/  root)  
Zugriff   : 2013-01-16 20:33:49.282818286 +0100  
Modifiziert: 2012-12-21 14:39:30.495831681 +0100  
Geändert  : 2012-12-21 14:39:30.495831681 +0100  
Geburt    : -  
root@kevin-ubuntu ~ # █
```

# Inode



# Python

```
#!/usr/bin/env python
2 import sys
4 a = 42
  print(sys.getrefcount(a))
6 b = a
  print(sys.getrefcount(a))
```

Listing 2: Python

```
1 9
  10
```

Listing 3: Output

# PHP

```
1 <?php  
2 $a = 42;  
3 xdebug_debug_zval( 'a' );  
4 $b = $a;  
5 xdebug_debug_zval( 'a' );  
6 ?>
```

Listing 4: PHP

```
1 a :  
2 ( refcount=1, is_ref=0), int 42  
3  
4 a :  
5  
6 ( refcount=2, is_ref=0), int 42
```

Listing 5: Output

```
16 struct memoryObject{  
    void* data;  
18     unsigned int refCount;  
    };  
20  
    struct memoryObject* alloc(size_t);  
22 struct memoryObject* getRef(struct memoryObject*);  
    void release(struct memoryObject*);
```

Listing 6: Header File

```
12 struct memoryObject* alloc(size_t size){  
    struct memoryObject* object = 0;  
  
14    object = (struct memoryObject*) malloc(sizeof(struct  
        memoryObject));  
  
16    object->refCount = 1;  
    object->data = malloc(size);  
  
18    return object;  
20 }
```

Listing 7: Allocate

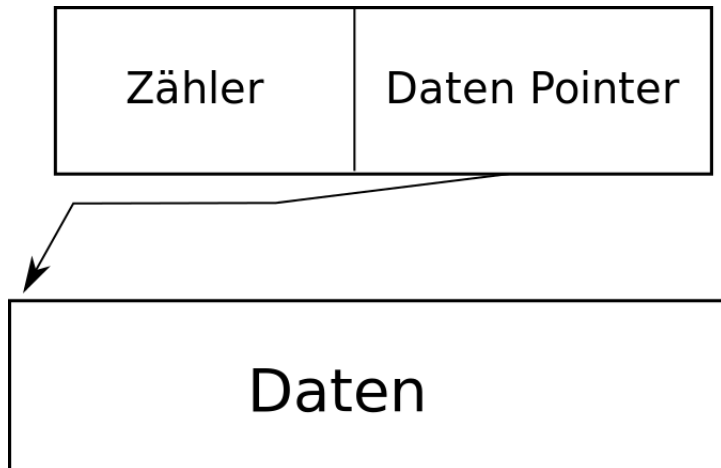


```
22 struct memoryObject* getRef(struct memoryObject* object
    ){
    object->refCount++;
24 return object;
}
```

Listing 8: Get new reference

```
28 void release(struct memoryObject* object){  
30     object->refCount--;  
32     if( object->refCount == 0){  
34         free(object->data);  
36         free(object);  
    }  
}
```

Listing 9: Release




```
26 void usage(){  
    struct memoryObject* object = alloc(sizeof(int));  
  
28    *((int*)object->data) = 5;  
  
30    struct memoryObject* object2 = getRef(object);  
  
32    release(object);  
    release(object2);  
34 }
```

Listing 10: Benutzung

```
16 typedef struct {  
    unsigned int refCount;  
    //void* data;  
18 }memoryObject;  
  
20 void* alloc(size_t);  
void* getRef(void*);  
22 void release(void*);
```

Listing 11: Header File

# Speicher



# Speicher



# Speicher





```
12 void* alloc(size_t size){  
14     size_t overheadSize = sizeof(memoryObject);  
16     memoryObject* o = 0;  
18     char* ptr = 0;  
20     o = (memoryObject*) malloc(overheadSize + size);  
22     ptr = (char*) o;  
24     ptr += sizeof(memoryObject);  
    o->refCount = 1;  
    return (void*) ptr;  
}
```

Listing 12: Allocate

```
28 void* getRef(void* pointer){  
30     memoryObject* o;  
32     char* cptr;  
34     cptr = (char*)pointer;  
36     cptr -= sizeof(memoryObject);  
38     o = (memoryObject*)cptr;  
39     o->refCount++;  
40     return pointer;  
41 }
```

Listing 13: Get new reference

```
42 void release(void* pointer){  
    memoryObject* o;  
    char* cptr;  
  
44     cptr = (char*)pointer;  
46     cptr -= sizeof(memoryObject);  
    o = (memoryObject*)cptr;  
  
48     o->refCount --;  
  
50     if( o->refCount == 0){  
52         free(o);  
    }  
54 }
```

Listing 14: Release

# Benutzung

```
28 void usage(){  
    int* intPtr = (int*)alloc(sizeof(int));  
  
30    *intPtr = 5;  
  
32    int* ptr2 = (int*)getRef(intPtr);  
  
34  
36    release(ptr2);  
    release(intPtr);  
}
```

Listing 15: Benutzung

# Zyklen

```
12 struct link{  
13     void* data;  
14     struct link* ptr;  
15 };  
16 void circular(){  
17     struct link* ptr1 = (struct link*)alloc(sizeof(struct  
18         link));  
19     struct link* ptr2 = (struct link*)alloc(sizeof(struct  
20         link));  
21     ptr1->ptr = (struct link*)getRef(ptr2);  
22     ptr2->ptr = (struct link*)getRef(ptr1);  
23     release(ptr1);  
24     release(ptr2);  
25 }
```

Listing 16: Zyklus

# Zyklen

```
1 $valgrind --tool=memcheck ./ReferenceCounting
  ==12551== Memcheck, a memory error detector
3 ==12551== Copyright (C) 2002-2011, and GNU GPL'd, by Julian Seward et al.
  ==12551== Using Valgrind-3.7.0 and LibVEX; rerun with -h for copyright info
5 ==12551== Command: ./ReferenceCounting
  ==12551==
7 ==12551==
  ==12551== HEAP SUMMARY:
9 ==12551==     in use at exit: 36 bytes in 2 blocks
  ==12551==   total heap usage: 2 allocs, 0 frees, 36 bytes allocated
11 ==12551==
  ==12551== LEAK SUMMARY:
13 ==12551==    definitely lost: 36 bytes in 2 blocks
  ==12551==   indirectly lost: 0 bytes in 0 blocks
15 ==12551==    possibly lost: 0 bytes in 0 blocks
  ==12551==   still reachable: 0 bytes in 0 blocks
17 ==12551==    suppressed: 0 bytes in 0 blocks
  ==12551== Rerun with --leak-check=full to see details of leaked memory
19 ==12551==
  ==12551== For counts of detected and suppressed errors, rerun with: -v
21 ==12551== ERROR SUMMARY: 0 errors from 0 contexts (suppressed: 2 from 2)
```

Listing 17: Valgrind

# C++

## Vorteile

- Operatoren können überladen werden
- Kein manuelles zerstören

## Nachteile

- Langsamer
- Nicht in C

# C++

## Shared Pointer

- Überladene Operatoren
- Reference Counting
- Zyklen



# Benutzung

```
18 void usage(){  
    shared_ptr<int> ptr1(new int);  
20  
    *ptr1= 10;  
22  
    shared_ptr<int> ptr2 = ptr1;  
24 }
```

Listing 18: Shared Pointer

## Weak Pointer

- Erhöht die Referenz NICHT
- Kann temporär zum Strong Pointer werden
- Kann Zyklen verhindern
- Zielobjekt kann schon gelöscht sein

# C++

```
void foo() {  
2  
    if (auto ptr = shared_ptr.lock()) {  
4        bar();  
    }  
6    else {  
    load();  
8    bar();  
    }  
10 }  
}
```

Listing 19: Weak Pointer

# System

- Intel Core2 Quad CPU Q6600 @ 2.40GHz
- 4GB RAM DDR2 800MHz
- 64Bit
- Linux 3.5.0-21-generic
- Ubuntu 12.10
- Compiler
  - gcc (Ubuntu/Linaro 4.7.2-2ubuntu1) 4.7.2
  - OpenJDK 1.7.0\_09

# Programm

## Optionen

- 10 000 000 Iterationen
- -O0 -lrt

# Normal

```
78 for(i = 0; i < ITERATIONS; ++i){  
    pointers[i] = (int *)malloc(sizeof(int));  
}
```

Listing 20: Alloc

```
82 for(i = 0; i < ITERATIONS; ++i){  
84   linkPtr = pointers[i];  
}
```

Listing 21: Get Reference

```
88 for(i = 0; i < ITERATIONS; ++i){  
    free(pointers[i]);  
}
```

Listing 22: Release

# Normal

Typ	Zeit	Anteil
Alloc time	0.583308s	0.71
getRef time	0.032474s	0.04
Release time	0.211202s	0.26
Total time	0.826984s	

# Naiv

```
46 for (i = 0; i < ITERATIONS; ++i){  
48     pointers[i] = (int*) alloc(sizeof(int));  
}
```

Listing 23: Alloc

```
52 for (i = 0; i < ITERATIONS; ++i){  
    linkPtr = (int*) getRef(pointers[i]);  
54     release(linkPtr);  
}
```

Listing 24: Get Reference

```
58 for (i = 0; i < ITERATIONS; ++i){  
    release(pointers[i]);  
}
```

Listing 25: Release



# Naiv

Typ	Zeit	Anteil
Alloc time	1.139242s	0.6
getRef time	0.296284s	0.16
Release time	0.468051s	0.25
Total time	1.903577s	

# Transparent

```
46 for (i = 0; i < ITERATIONS; ++i){  
48     pointers[i] = (int*) alloc(sizeof(int));  
}
```

Listing 26: Alloc

```
52 for (i = 0; i < ITERATIONS; ++i){  
    linkPtr = (int*) getRef(pointers[i]);  
54     release(linkPtr);  
}
```

Listing 27: Get Reference

```
58 for (i = 0; i < ITERATIONS; ++i){  
    release(pointers[i]);  
}
```

Listing 28: Release

# Transparent

Typ	Zeit	Anteil
Alloc time	0.648253s	0.48
getRef time	0.313264s	0.23
Release time	0.381382s	0.28
Total time	1.342898s	

## Shared Pointer

```
34  for(i = 0; i < ITERATIONS; ++i){  
    pointers[i].reset(new int);  
}
```

Listing 29: Alloc

```
38  for(i = 0; i < ITERATIONS; ++i){  
40  linkPtr = pointers[i];  
}
```

Listing 30: Get Reference

# Shared Pointer

Typ	Zeit	Anteil
Alloc time	1.844750s	0.35
getRef time	2.017531s	0.39
Release time	1.356881s	0.26
Total time	5.219162s	

# Java

```
12 public static void bench(){
13     int i;
14     Integer linkPtr;
15     Integer [] pointers = new Integer [ITERATIONS];
16
17     start_time = System.nanoTime();
18     for(i = 0; i < ITERATIONS; ++i){
19         pointers[i] = new Integer(0);
20     }
21     alloc_time = System.nanoTime();
22
23     for(i = 0; i < ITERATIONS; ++i){
24         linkPtr = pointers[i];
25     }
```

Listing 31: Java

# Java

Typ	Zeit	Anteil
Alloc time	6.129283253s	0.94
getRef time	0.00263494s	0.0004
Release time	0.419077866s	0.06
Total time	6.548361119s	

Normal	Naiv	Transparent	C++	Java
0.826984s	1.903577s	1.342898s	5.219162s	6.548361119s
+0%	+130%	+62%	+531%	+692%
4Byte	20Byte	8Byte	20Byte	?



## Was haben wir gelernt?

- Reference Counting kann Probleme lösen...
- ...jedoch nicht alle!
- Birgt auch Probleme
- Benötigt oft Hilfe vom Benutzer
- Kein Wunderheilmittel!
- Wird in manchen Bereichen aber zwingend benötigt
- C ist am schnellsten ;)

- <http://e2fsprogs.sourceforge.net/ext2intro.html>
- <http://www.eosgarden.com/en/articles/c-refcount/>
- Boost/STL/Java/Python Dokumentation
- [http://en.wikipedia.org/wiki/Reference\\_counting](http://en.wikipedia.org/wiki/Reference_counting)