The 18th Annual Network and Distributed System Security Symposium

SigGraph: Brute Force Scanning of Kernel Data Structure Instances Using Graph-based Signatures

Zhiqiang Lin[1]

Junghwan Rhee[1], Xiangyu Zhang[1], Dongyan Xu[1], Xuxian Jiang[2]

[1]Purdue University
[2]North Carolina State University

February 7th, 2011

Problem Statement

- Given a kernel data structure definition
- Identifying *instances* of this data structure in a kernel memory image at arbitrary location

```
struct task {
  [0] struct thread *thread;
  [4] struct memory *mm;
  [8] struct signal *signal;
 [12] struct task *parent;
 [16] int magic_number;
}
```

A simplified Linux Kernel **`task_struct`**

Security Applications: Memory Forensics

Password

```
struct user_account {
00: short int u_type;
04: pid_t u_pid;
08: char u_line[32];
40: char uid[4];
44: char user[32];
76: char password[128];
204: char u_host[128];
332: short int e_termination;
334: short int e_exit;
336: long int u_session;
340: struct timeval u_tv;
348: int32_t u_addr_v6[4];
}
```

Data structure signatures **play a critical role in memory forensics**

Security Applications: Kernel Rootkit Defense

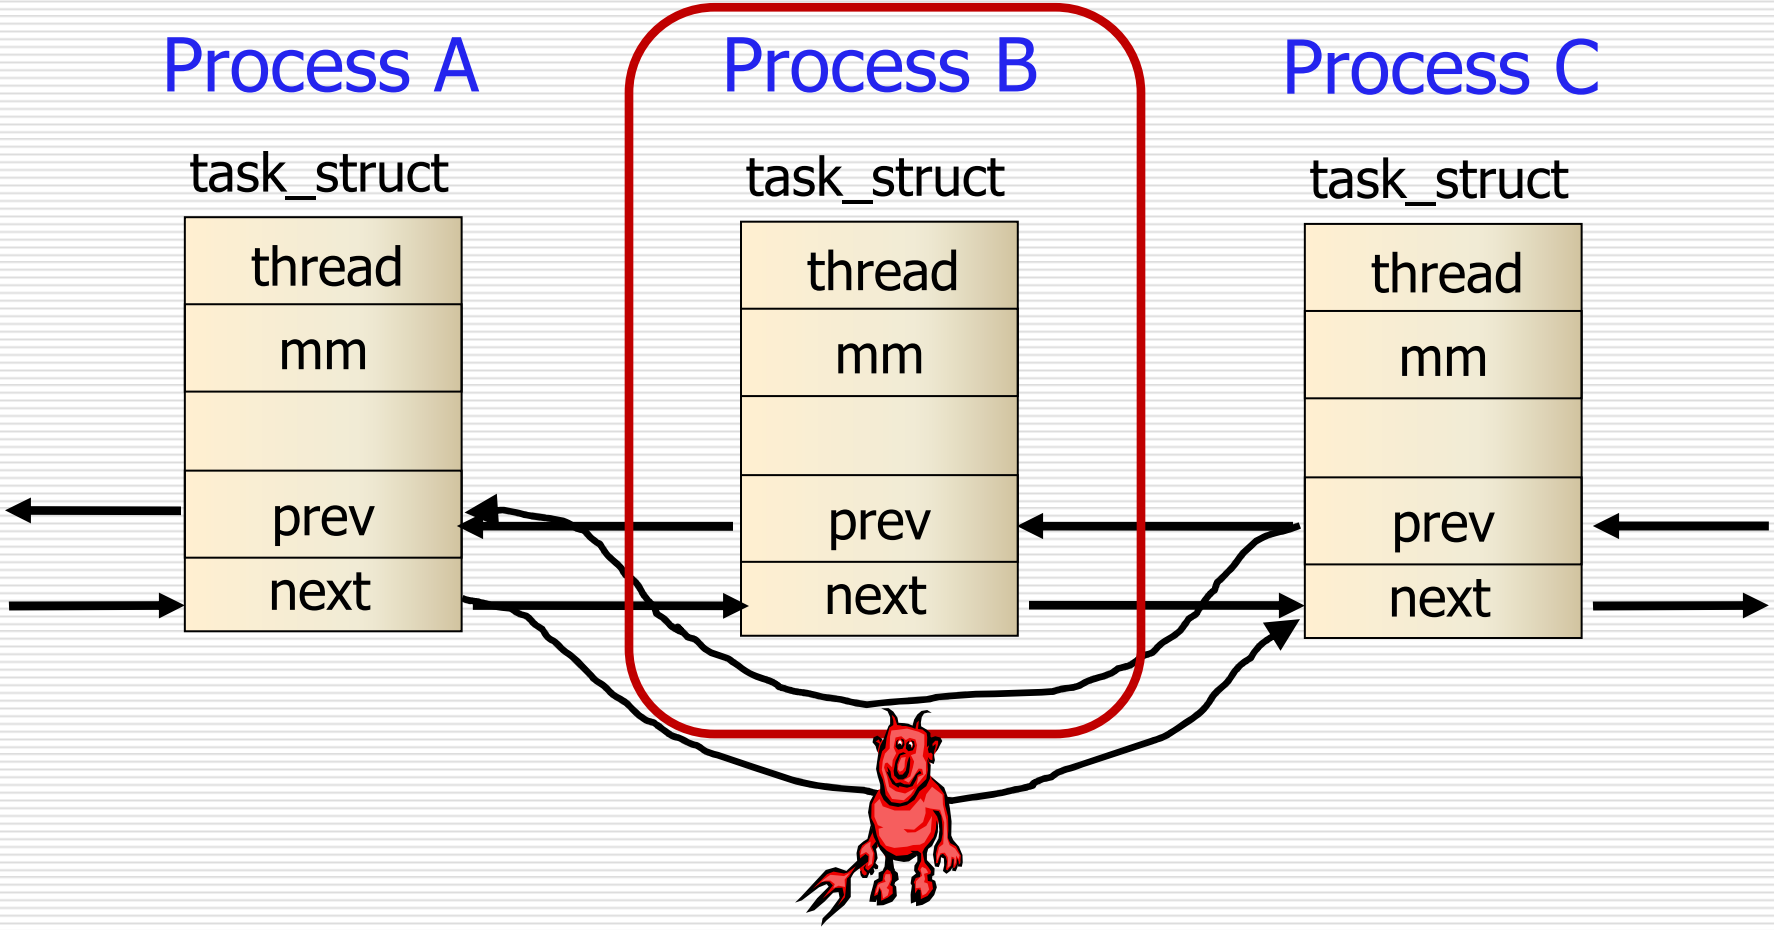

State-of-the-art

- Value-invariant signature schemes
  - Klist [Rutkowska,2003], GREPEXEC [bugcheck, 2006], Volatility [Walters, 2006], [Schuster, 2006], [Dolan-Gavitt et al., CCS'09]

Invariant value can be changed?

[Dolan-Gavitt et al., CCS'09]

```
struct task {
  [0] struct thread *thread;
  [4] struct memory *mm;
  [8] struct signal *signal;
 [12] struct task *parent;
 [16] int magic_number;
}
```

magic_number=0xabcdef0f

Field w/o value invariant?

Key Idea

```
struct task {
  [0] struct thread *thread;
  [4] struct memory *mm;
  [8] struct signal *signal;
 [12] struct task *parent;
}

struct thread {
  [0] struct task *task;
}

struct memory {
  [0] struct vma *mmap;
  [4] void (*map_area)
      (struct memory*
mmap);
}

struct signal {
  [0] struct task_status
*status;
}
```

task(x) ⇔ thread(*(x+0)) ∧ mm(*(x+4)) ∧ signal(*(x+8)) ∧ task(*(x+12))

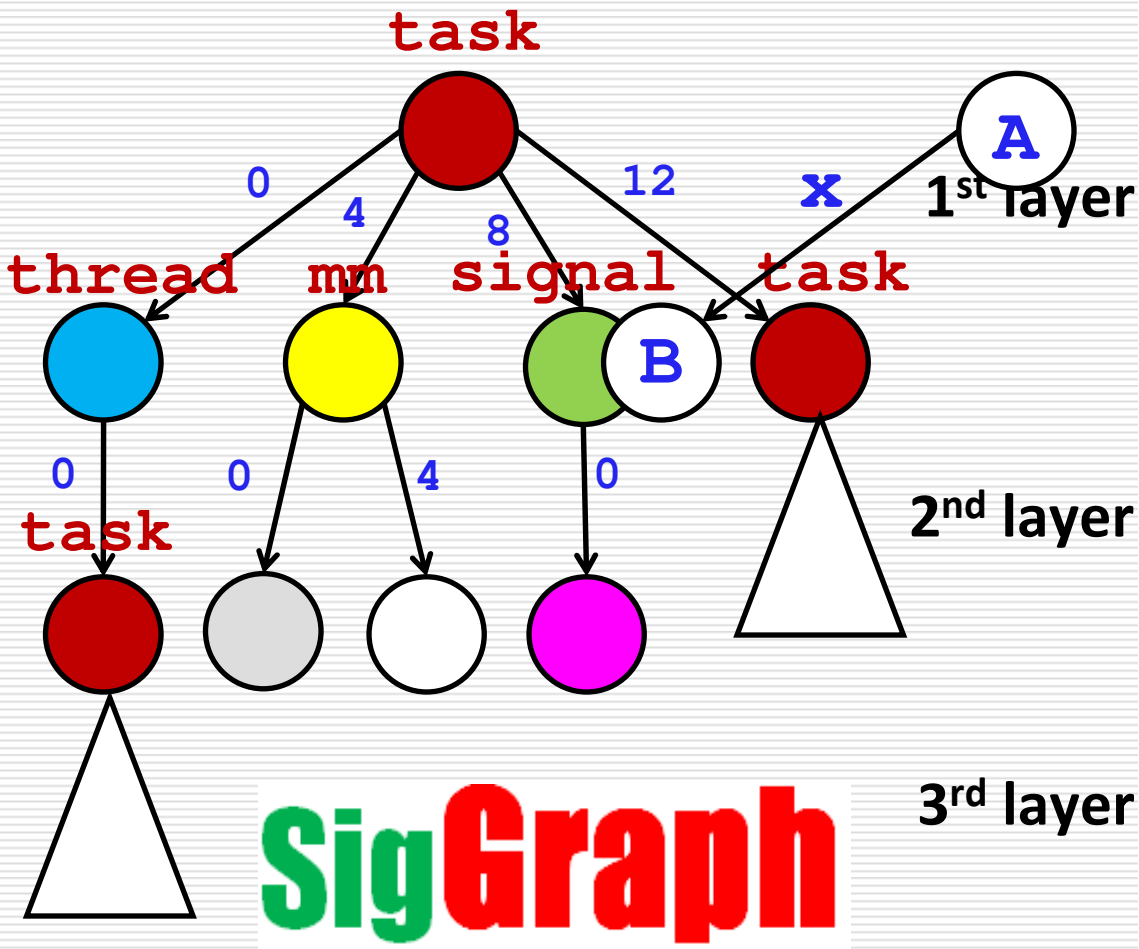

How to Use SigGraph

```
0xc001c0a8: 0xc002c0a8  0xc002bee0  0xc002caa0  0xc00ddbb0
...
0xc002c0a8: 0xc12a0e7c  0xc727faa8  0xbfbb9195  0x00000009
...
0xc002bee0: 0xc001c114  0xc001c16c  0xffb29122  0x00201001
...
0xc002caa0: 0xb002ca20  0xb021d00a  0xc05b9f5c  0x00000000
...
0xc00ddbb0: 0xc12a0e7c  0xc727faa8  0xc001c114  0xc001c16c
...
```

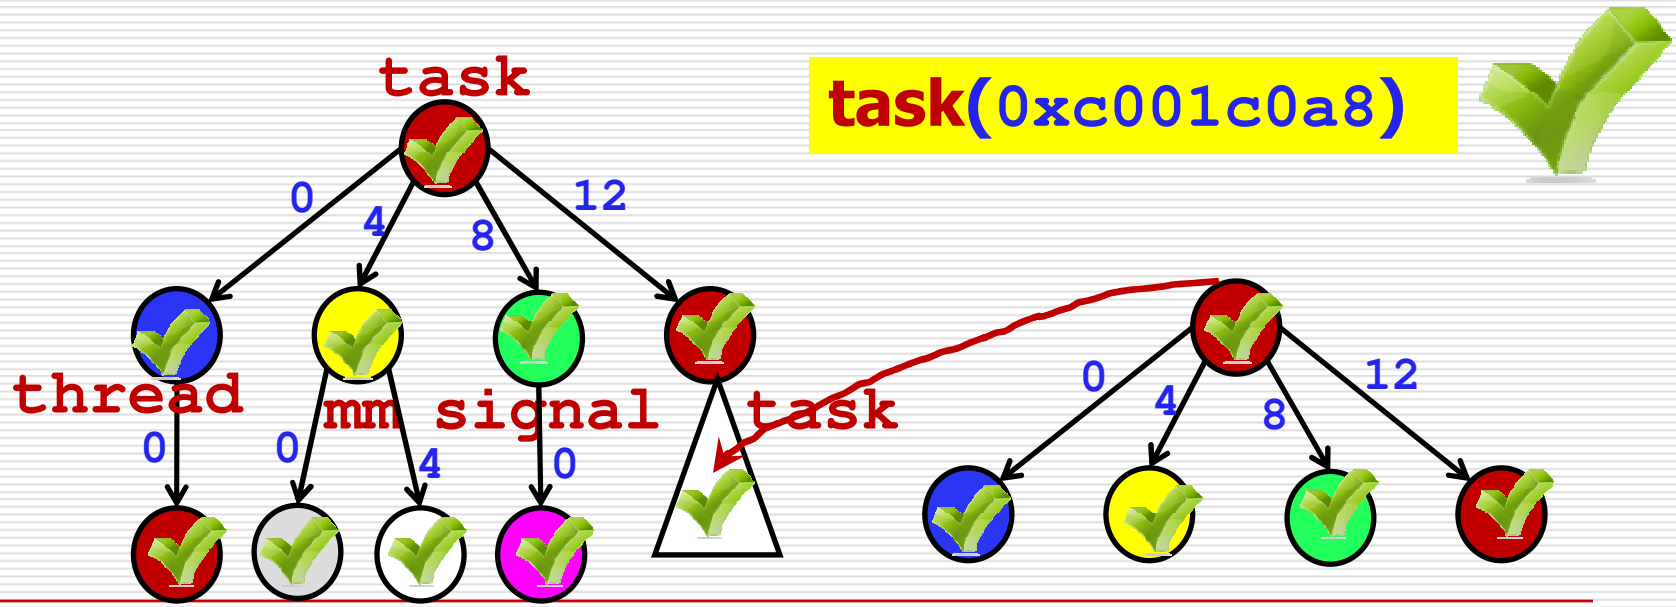

SigGraph Overview

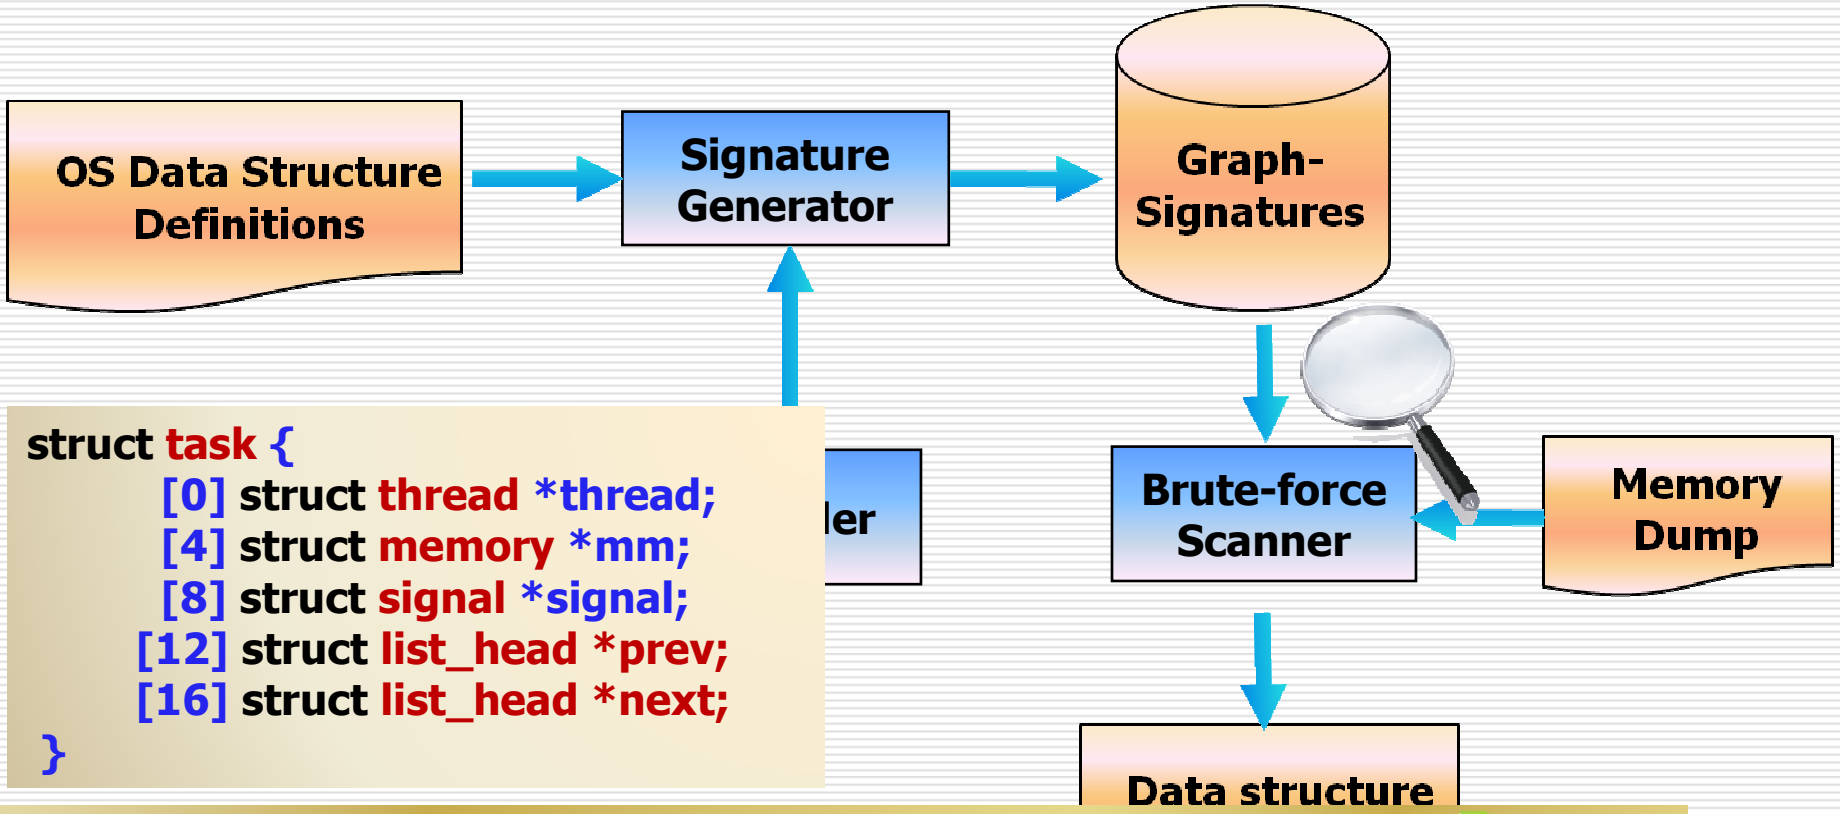

(1) Compiler approach
(2) Extracting from debug information
(3) Reverse engineering kernel

Signature Generator

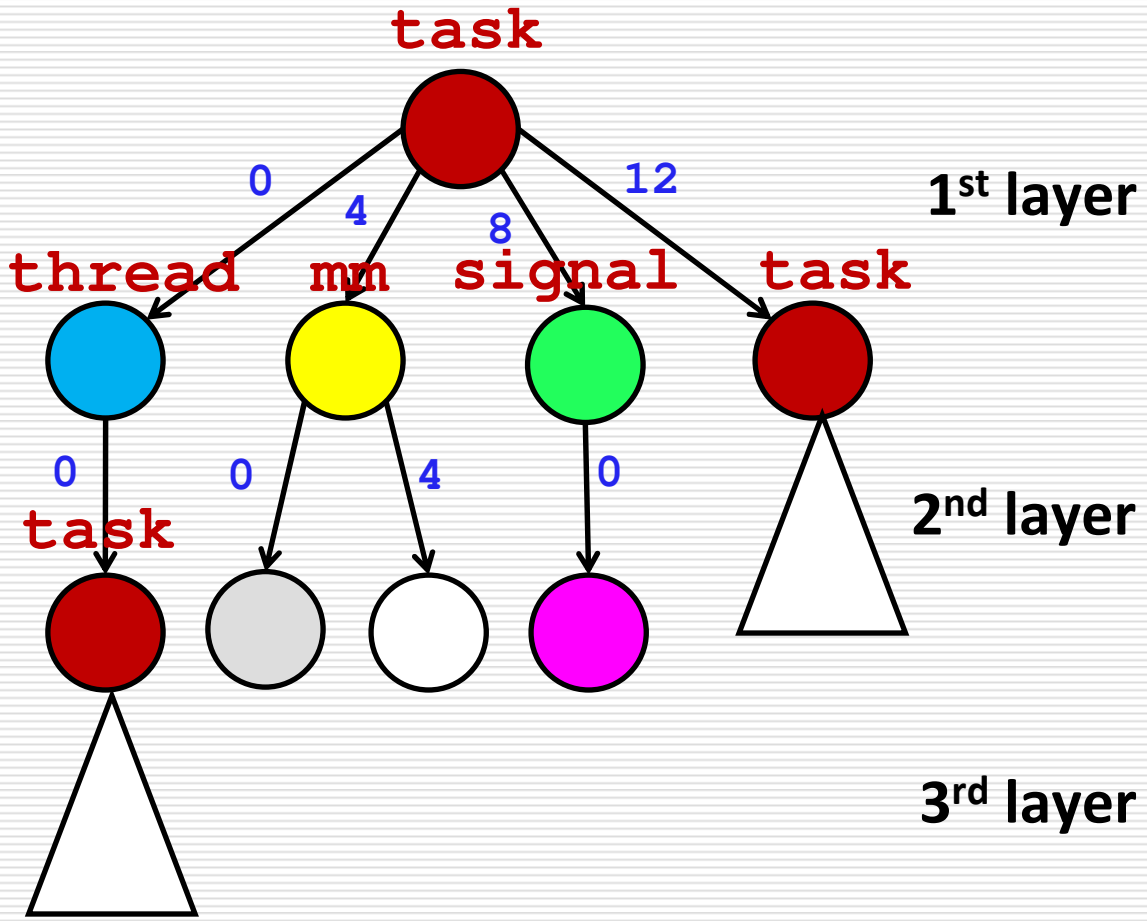

- Challenge: Signatures must be *unique*, *non-isomorphic* among each other.

Isomorphism

```
struct B {
   [0]  E * b1;
   [4]  B * b2;
}
struct E {
     ...
   [12] G * e1;
     ...
   [24] H * e3;
}
```

```
struct G {
     ...
   [10] int * g;
}
struct GG {
     ...
   [4] char * gg1;
   [8] char * gg2;
}
```

```
struct BB {
   [0]  EE * bb1;
   [4]  BB * bb2;
}
struct EE {
     ...
   [12] GG * ee1;
     ...
   [24] HH * ee3;
}
```

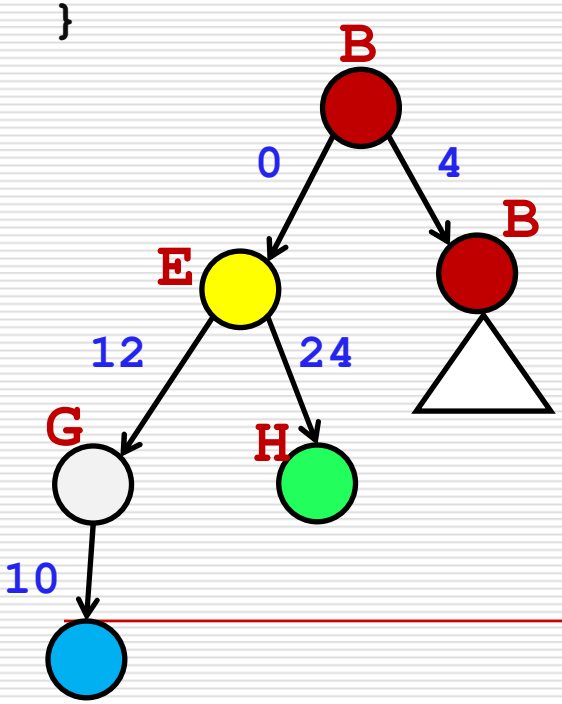

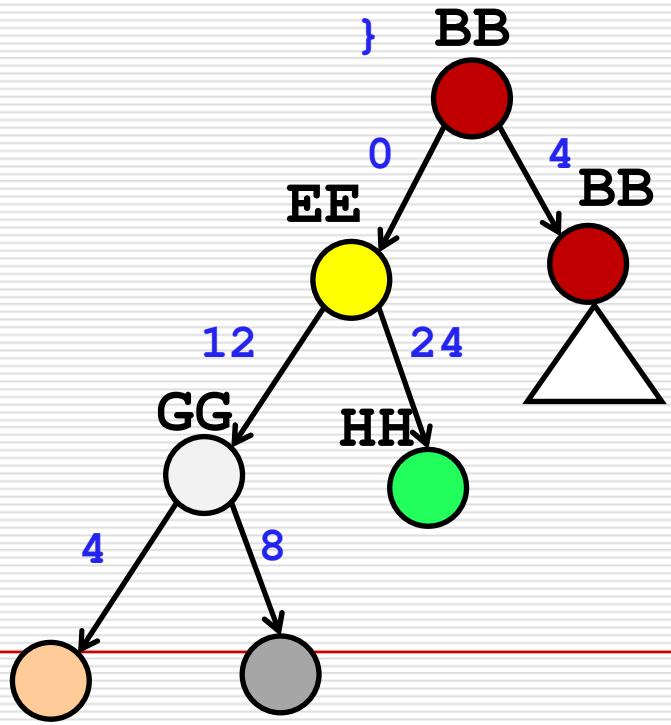

Isomorphism

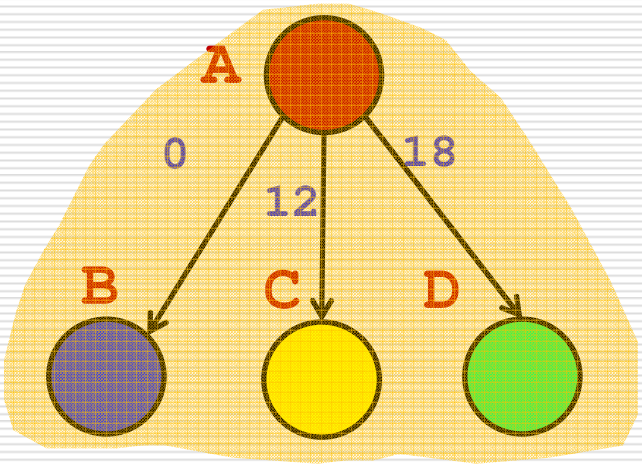

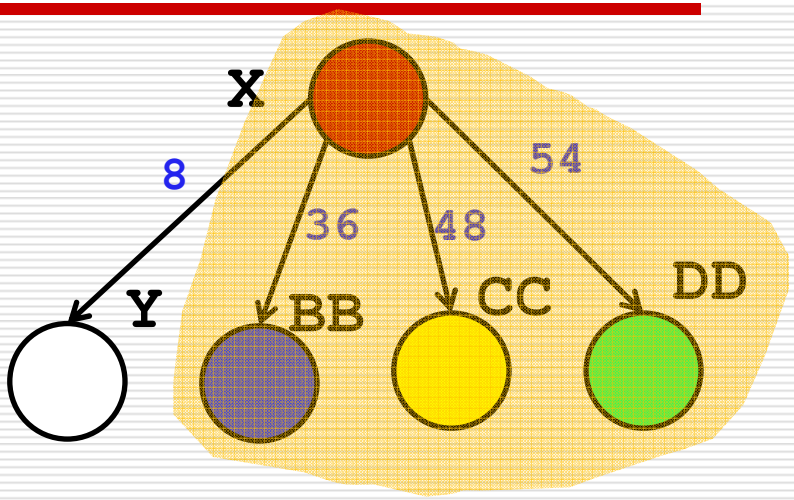


```
struct A {
  [0]  struct B * a1;
    ...
  [12] struct C * a2;
    ...
  [18] struct D * a3;
}
```

```
struct X {
    ...
  [8]  struct Y  * x1;
    ...
  [36] struct BB * x2;
    ...
  [48] struct CC * x3;
    ...
  [54] struct DD * x4;
}
```

Our Solution

- *Immediate pointer pattern (IPP):* one-layer pointer structure as a *string*

$$IPP(T) = f_1 \cdot t_1 \cdot (f_2 - f_1) \cdot t_2 \cdot \ldots \cdot (f_n - f_{n-1}) \cdot t_n$$

```
struct B {
   [0]   E * b1;
   [4]   B * b2;
}
```

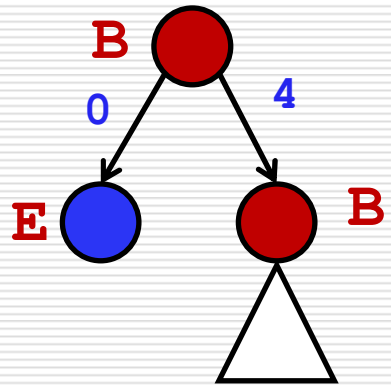

$$IPP(B) = 0 \cdot E \cdot 4 \cdot B$$

- Pointer expansion '$\xrightarrow{T}$'

$$IPP(B) \xrightarrow{B} 0 \cdot E \cdot 4 \cdot (0 \cdot E \cdot 4 \cdot B)$$

Problem Formulation

$IPP(T_1, T_2, \ldots, T_n) = t_1 \cdot t_2 \cdot \ldots \cdot t_n$

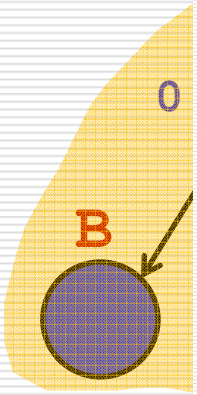

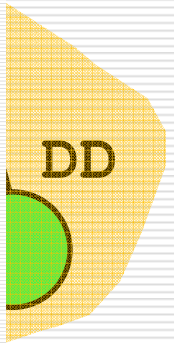

How to Use SigGraph

```
0xc001c0a8: 0xc002c0a8  0xc002bee0  0xc002caa0  0xc00ddbb0
...
0xc002c0a8: 0xc12a0e7c  0xc727faa8  0xbfbb9195  0x00000009
...
0xc002bee0: 0xc001c114  0xc001c16c  0xffb29122  0x00201001
```

Ignore the symbol type at specific layer

```
0xc00ddbb0: 0xc12a0e7c  0xc727faa8  0xc001c114  0xc001c16c
...
```

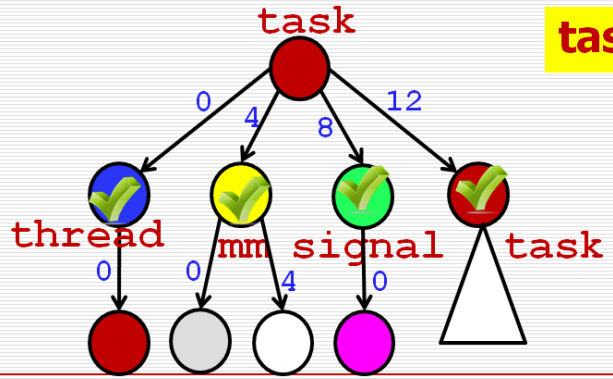

task(0xc001c0a8)

$IPP(A) = 2 \cdot CC \cdot 6 \cdot DD$

"If IPP(A) is a substring of IPP(X)"

Profiler

- Practical pointer issues
  - `null` Pointer
  - `void` Pointer
  - Special Pointer
    - LIST_POISON1 (0x00100100)
    - LIST_POISON2 (0x00200200)
    - SPINLOCK_MAGIC (0xdead4ead)

Pruning a few noisy pointer fields does not degenerate the uniqueness of the graph-based signatures

LiveDM [Rhee et al., RAID'10]

Evaluation

- Memory snapshot collection
  - QEMU

- Ground truth acquisition
  - RedHat crash utility
  - Symbolic information
    - system.map

crash-utility.redhat.com

- Profiling run
  - Long runs with typical workload

Evaluation on Memory Forensics

Data Structures of Interest	"True" Instance	SigGraph		Value-invariant	
		FP%	FN%	FP%	FN%
task_struct	88	0.00	0.00	0.00	0.00
thread_info	88	0.00	0.00	6.45	1.08
mm_struct	52	0.00	0.00	0.00	0.00
vm_area_struct	2174	0.40	0.00	7.52	0.00
files_struct	53	0.00	0.00	0.00	0.00
fs_struct	52	0.00	0.00	0.00	0.00
dentry	31816	0.01	0.00	0.01	0.00
sysfs_dirent	2106	0.52	0.00	97.63	0.00
socket	55	0.00	0.00	0.00	12.24
sock	55	0.00	0.00	0.00	27.90
user_struct	10	0.00	0.00	99.91	0.00

Application: Rootkit Detection

Rootkit Name	Target Object	Inside View	Crash Tool		SigGraph	
		#obj.s	#obj.s	detected	#obj.s	detected
adore-ng-2.6	module	23	23	✗	24	✓
adore-ng-2.6’	task_struct	62	63	✓	63	✓
cleaner-2.6	module	22	22	✗	23	✓
enyelkm 1.0	module	23	23	✗	24	✓
hp-2.6	task_struct	56	57	✓	57	✓
linuxfu-2.6	task_struct	59	60	✓	60	✓
modhide-2.6	module	22	22	✗	23	✓
override	task_struct	58	59	✓	59	✓
rmroots	task_struct	56	N/A	✗	55	✓
rmroots’	module	23	N/A	✗	24	✓

```
ps
lsmod
```

Related Work

- Kernel memory mapping and analysis
  - Copilot [Petroni et al., Security'04], [Petroni et al., CCS'07]
  - Gibraltar [Baliga et al., ACSAC'08]
  - KOP [Carbone et al.,CCS'09]
- Memory forensics
  - Memory graph-based: Redhat crash utility, KOP
  - Value-invariant Signature: Klist [Rutkowska,2003], GREPEXEC [bugcheck, 2006], Volatility [Walters, 2006], [Schuster, 2006], [Dolan-Gavitt et al., CCS'09]
- Dynamic heap type inference [Polishchuk et al., 2007]

Conclusion

- Points-to relations can be leveraged to generate graph-based signatures for brute force scanning
- SigGraph, a framework that generates *non-isomorphic structural-invariant signatures*
  - Complements *value-invariant* signatures
- Applications:
  - Kernel memory forensics
  - Kernel rootkit detection

Q&A

Thank you

For more information

{zlin,rhee,xyzhang,dxu}@cs.purdue.edu