

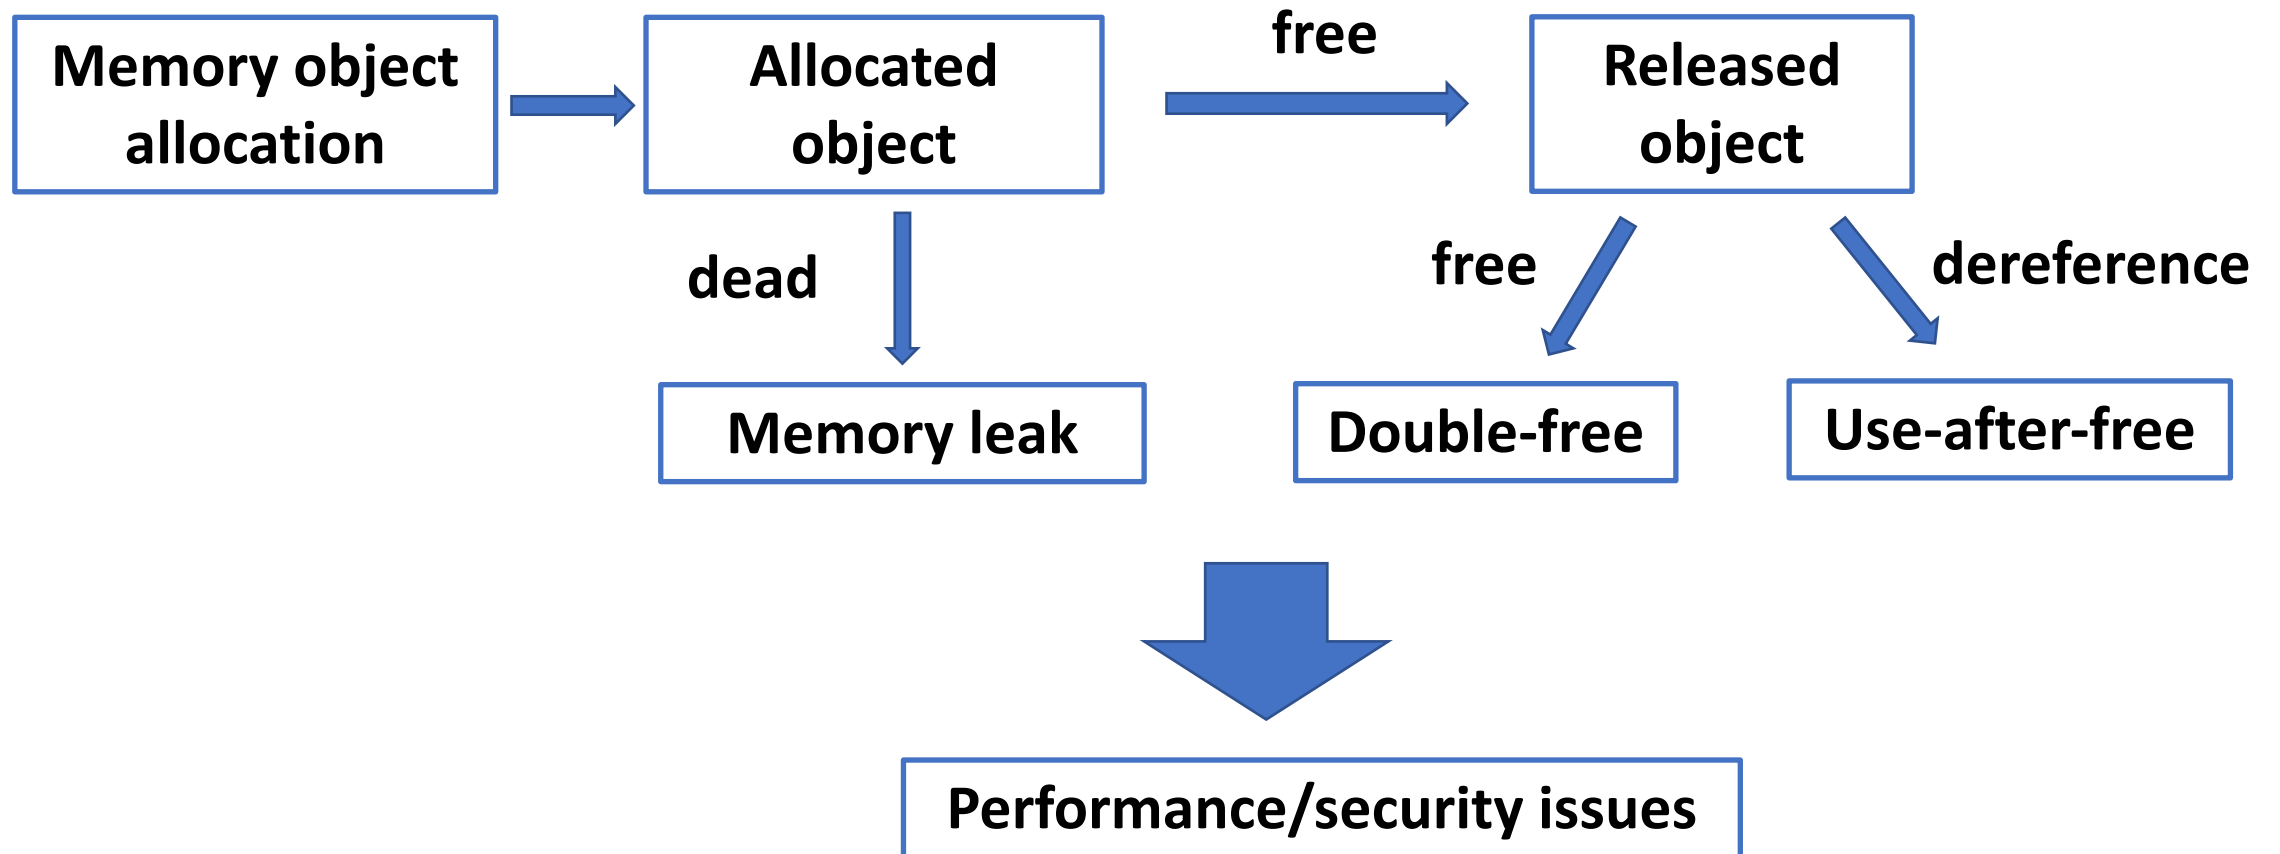
Detecting Kernel Memory Bugs through Inconsistent Memory Management Intention Inferences

Dinghao Liu Zhipeng Lu Shouling Ji Kangjie Lu
Jianhai Chen Zhenguang Liu Dexin Liu Renyi Cai Qinming He

USENIX Security 2024

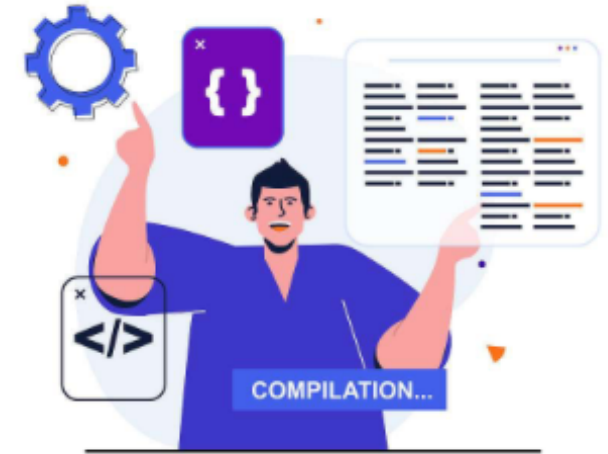
Background

➤ Kernel memory bugs



Kernel Memory Bug Detection

- **Data-flow analysis**
 - Keep tracking the lifecycle of a memory object.
 - **Path explosion for kernel analysis**
- **Function pairing**
 - The memory management function pairs are intended to operate in concert.
- **Similarity analysis**
 - Memory operations should be consistent within similar paths.
 - **Both rely on CORRECT code snippets**



Insight

- **Memory bugs frequently manifest within error handling paths, where programs attempt to release allocated memory resources.**
- **A multitude of memory bugs originate from a fundamental issue: the unclarity in ownership and lifecycle management of memory objects when premature release is necessitated.**

**There exists inconsistent intentions
regarding the management of memory objects.**

Memory Management Strategies

➤ Callee-based management

- Functions who allocate heap memories do the cleanup on failure
- The callers need not to take care of the cleanup
- The most widely-used strategy of memory management

```
1  /* net/smc/smc_er.c */
2  int smc_wr_alloc_link_mem(struct smc_link *link)
3  {
4      link->wr_tx_bufs = kcalloc(...);
5      if (!link->wr_tx_bufs)
6          goto no_mem;
7      link->wr_rx_bufs = kcalloc(...);
8      if (!link->wr_rx_bufs)
9          goto no_mem_wr_tx_bufs;
10     link->wr_tx_ibs = kcalloc(...);
11     if (!link->wr_tx_ibs)
12         goto no_mem_wr_rx_bufs;
13     ...
14     return 0;
15     ...
16 no_mem_wr_rx_bufs:
17     kfree(link->wr_rx_bufs);
18 no_mem_wr_tx_bufs:
19     kfree(link->wr_tx_bufs);
20 no_mem:
21     return -ENOMEM;
22 }
```

Memory Management Strategies

➤ Caller-based management

- The caller functions do the cleanup on failure of callees
- The error handling of callee functions would be simple and clear
- This strategy is also popular

```
1  /* drivers/media/usb/tm6000/tm6000-video.c */
2  static int tm6000_alloc_urb_buffers(struct tm6000_core *dev)
3  {
4      ...
5      dev->urb_buffer = kmalloc_array(...);
6      if (!dev->urb_buffer)
7          return -ENOMEM;
8
9      dev->urb_dma = kmalloc_array(...);
10     if (!dev->urb_dma)
11         return -ENOMEM;
12     ...
13     if (!dev->urb_buffer[i]) {
14         tm6000_err("unable to allocate ... buffer %i\n",
15                 dev->urb_size, i);
16         return -ENOMEM;
17     }
18     ...
19     return 0;
20 }
21
22 static int tm6000_prepare_isoc(struct tm6000_core *dev)
23 {
24     ...
25     if (tm6000_alloc_urb_buffers(dev) < 0) {
26         tm6000_err("cannot allocate memory for urb buffers\n");
27
28         /* call free, as some buffers might have been allocated */
29         tm6000_free_urb_buffers(dev);
30         ...
31         return -ENOMEM;
32     }
33     ...
34 }
```

Inconsistent MM Intentions

➤ Impact analysis

- **Memory corruption:** the callee function adopts callee-based management, while the caller function adopts caller-based management.
- **Memory leak:** the callee function adopts caller-based management, while the caller function adopts callee-based management.

➤ Causes of inconsistent MM intentions

- Complex error handling logic.
- Imperfect code updating.
- Implicit OS MM mechanisms.

Challenges & Solutions

➤ **Challenge 1: How to infer MM intentions?**

- Using static program analysis to infer explicit MM intentions.
- Using large language model to better understating implicit MM intentions.

➤ **Challenge 2: How to balance precision and efficiency?**

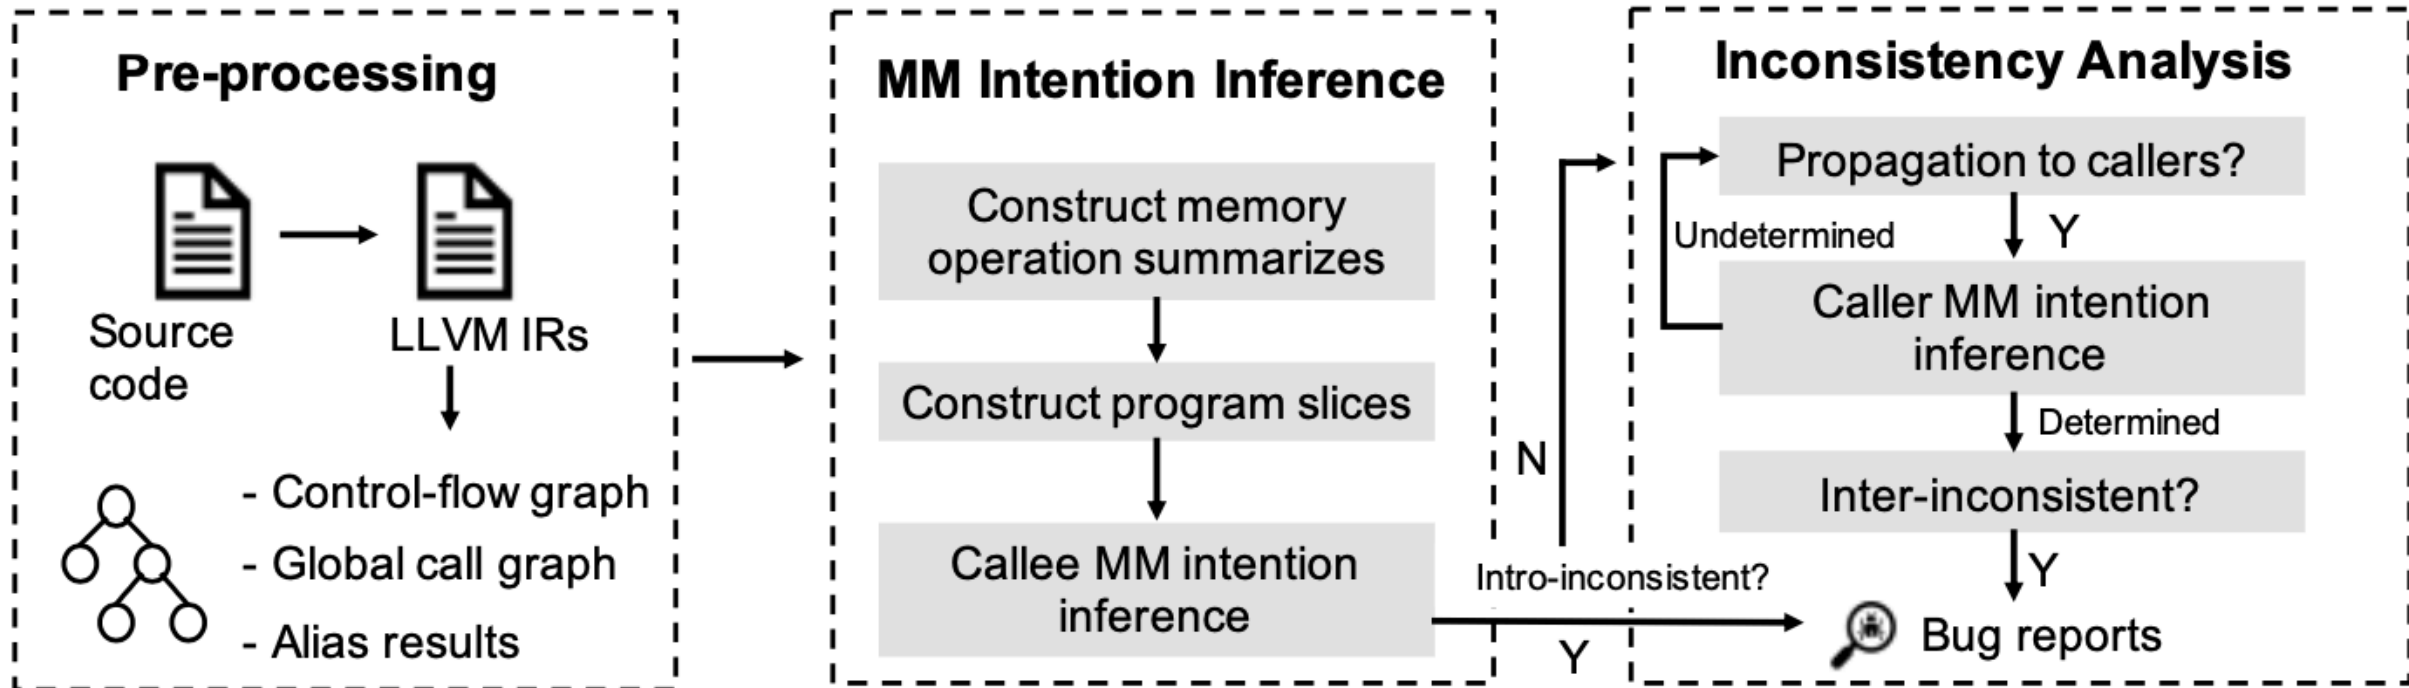
- Demand-driven memory operation summarization
- Object-based code slicing



System Design

Overview

IMMI (Inconsistent Memory Management Intentions)



Memory Operation Summarization

➤ Target: Simplify the inter-procedural analysis

Kernel developers typically implement memory management in a layered fashion

```
debug_close(struct inode *inode, struct file *file)
file_private_info_t *p_info;
p_info = (file_private_info_t *) file->private_data;
...
debug_info_free(p_info->debug_info_snap);
...
kfree(file->private_data);
...
```

```
debug_info_free(debug_info_t *db_info)
...
debug_areas_free(db_info);
kfree(db_info->active_pages);
kfree(db_info);
...
```

```
debug_areas_free(debug_info_t *db_info)
...
kfree(db_info->areas);
...
```

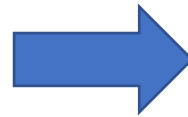
Memory Operation Summarization

➤ Techniques: Demand-driven summarization

First generate function level summaries

Find callers

```
debug_info_free(debug_info_t *db_info)
...
kfree(db_info->active_pages);
kfree(db_info);
...
```



```
F: debug_info_free, Arg:db_info, L0:
[ (Release-kfree, 0), db_info->active_pages;
  (Release-kfree, 0), db_info; ]
```

```
debug_close(struct inode *inode, struct file *file)
file_private_info_t *p_info;
p_info = (file_private_info_t *) file->private_data;
...
debug_info_free(p_info->debug_info_snap);
...
kfree(file->private_data);
...
```



```
F: debug_close, Arg:file, L0:
[ (Release-debug_info_free, 0),
  file->private_data->debug_info_snap;
  (Release-kfree, 0), file->private_data; ]
```

Memory Operation Summarization

➤ Techniques: Demand-driven summarization

Reconstruct complete summaries on-demand

```
F: debug_close, Arg:file, L0:
```

```
[ (Release-debug_info_free, 0),  
  file->private_data->debug_info_snap;  
  (Release-kfree, 0), file->private_data; ]
```

Find a customized release API, check its summary

```
F: debug_info_free, Arg:db_info, L0:
```

```
[ (Release-kfree, 0), db_info->active_pages;  
  (Release-kfree, 0), db_info; ]
```

```
F: debug_close, Arg:file, L0:
```

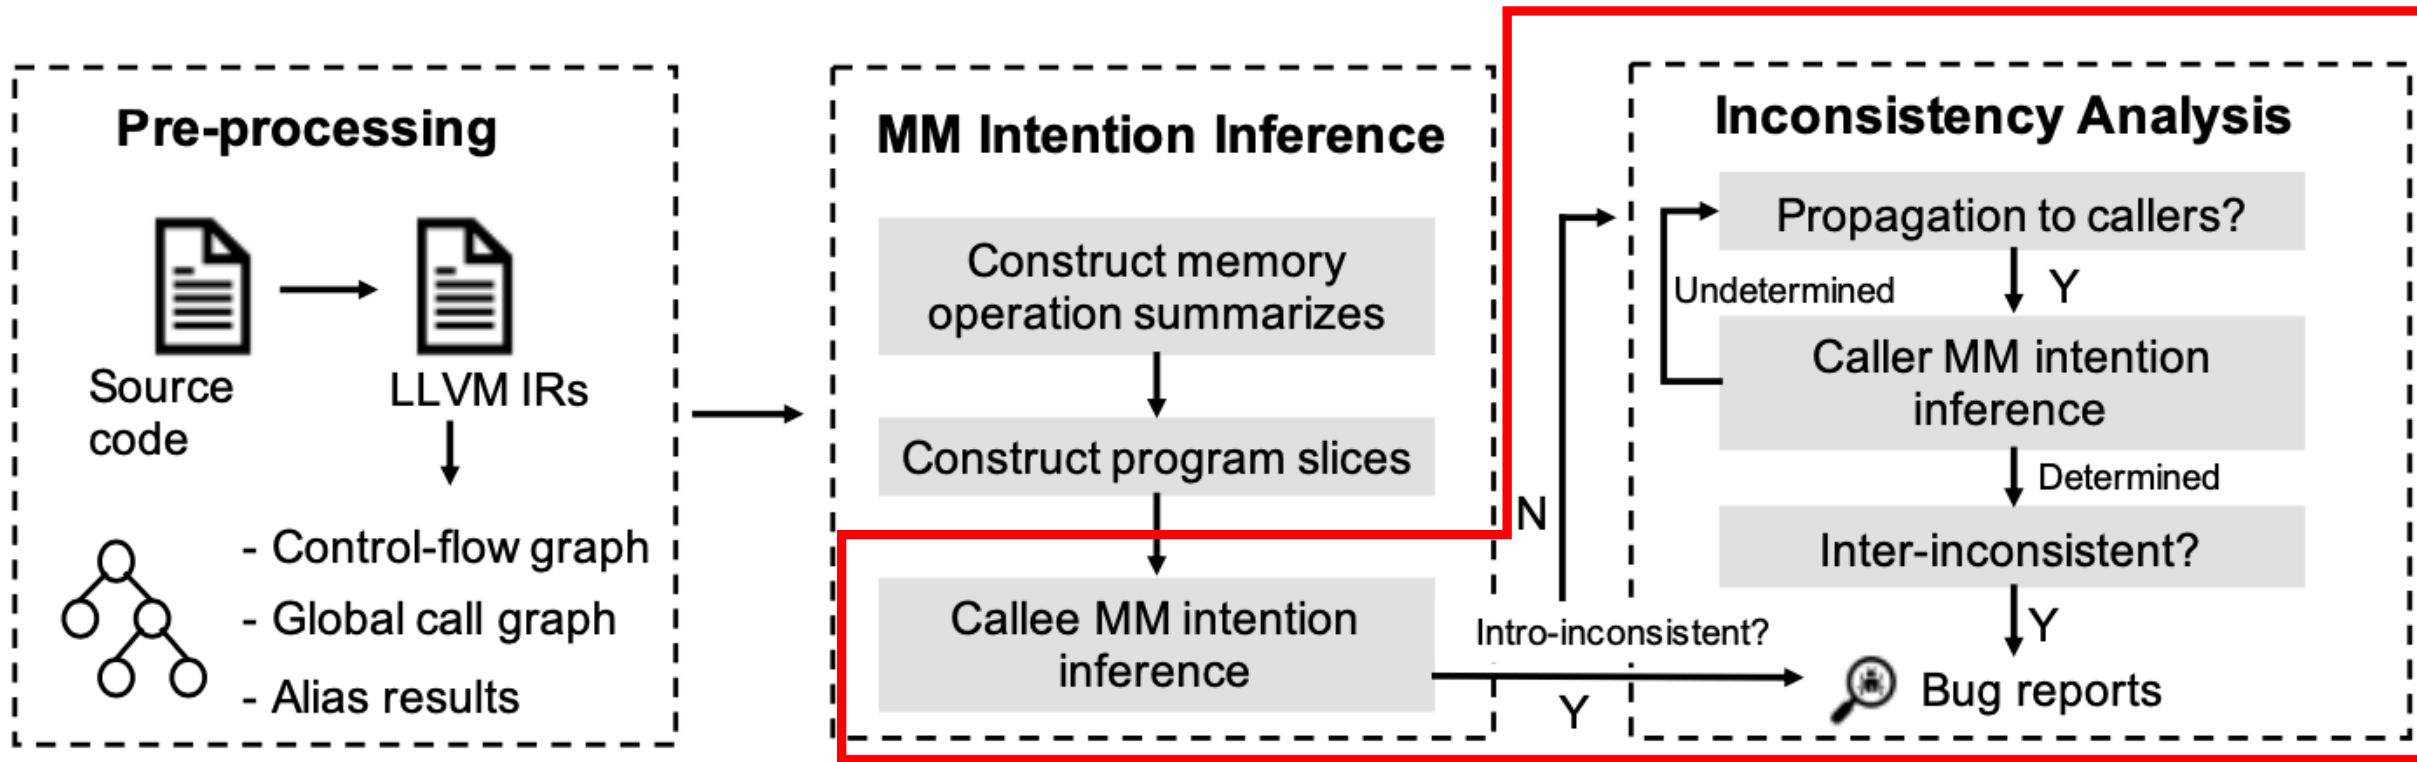
```
[ (Release-debug_info_free, 0),  
  file->private_data->debug_info_snap-  
  >active_pages;  
  (Release-debug_info_free, 0),  
  file->private_data->debug_info_snap;  
  (Release-kfree, 0), file->private_data; ]
```

Object-based Code Slicing

- **Target: Confine the analysis scope and enhance efficiency**
 - **Extracting memory object**
 - Extract memory object through official memory allocation APIs.
 - **Identifying error handling paths**
 - Combining the forward and backward data-flow analysis.
 - **Code slicing**
 - Tracking the usages of the memory object.
 - Finding the instruction that permits access to the object for other functions
 - Eliminating paths that are not reachable by this instruction.
 - Eliminating paths that reset the memory object after its release.

Overview

IMMI (Inconsistent Memory Management Intentions)



MM Intention Inference

➤ Callee MM intention inference


Upon a memory object is allocated, IMMI analyzes the host function:

- **Callee-based management**
 - All error paths within the code slices ensure the release of the memory objects.
- **Caller-based management**
 - None of the error paths within the code slices release the memory object.
- **Undetermined management**
 - Some error paths release the memory object, while others do not.

MM Intention Inference

➤ Callee MM intention inference

Upon a memory object is allocated, IMMI analyzes the host function:

- **Callee-based management**
 - All error paths within the code slices ensure the release of the memory objects.
- **Caller-based management**
 - None of the error paths within the code slices release the memory object.
- **Undetermined management**  **Intro-inconsistency, stop further analysis**
 - Some error paths release the memory object, while others do not.

MM Intention Inference

➤ **Caller MM intention inference**

When the memory object is propagated to caller functions, IMMI analyzes the callers:

- **Caller-based management**
 - The memory object is released on failure of the callee.
- **Callee-based management**
 - The memory object is not released upon failure of the callee, but is released in the subsequent error paths.
- **Undetermined management**
 - No release operation is detected, track the transfer to further callers.

MM Intention Inference

➤ Large language model integration

Target 1: Improving the understanding of implicit MM mechanisms.

Target 2: Improving the understanding of code comments.

Prompt message

- You re now a program static analysis expert. Your following analysis is based on the following function: [code given by IMMI]
- We define "error path" in a function as follows: A sequence of basic blocks that finally returns a non- zero number or null pointer. Note that if a call returns error but the path finally does not return a negative number or null pointer, this path is not an error path.
- A heap memory ``[object given by IMMI]“ is allocated through “[instruction given by IMMI]“ Please identify all of the LATER error paths after the allocation. Then, please analyze whether all of these paths have freed the heap memory. Pay attention to the implicit kernel memory release operations. Your final conclusion should be a separate line like: [Conclusion: Answer], "Answer" should only be "yes" or "no".



Case Study

A Memory Leak Bug Found by IMMI

```
1 /* drivers/net/ethernet/qlogic/qla3xxx.c */
2 static int ql_alloc_buffer_queues(struct ql3_adapter *qdev)
3 {
4     ...
5     qdev->lrg_buf = kmalloc_array(...);
6     if (qdev->lrg_buf == NULL)
7         return -ENOMEM;
8     ...
9     if (qdev->lrg_buf_q_alloc_virt_addr == NULL) {
10        netdev_err(qdev->ndev, "lBufQ failed\n");
11        return -ENOMEM;
12    }
13
14    return 0;
15 }
16
17 static int ql_alloc_mem_resources(struct ql3_adapter *qdev)
18 {
19     ...
20     if (ql_alloc_net_req_rsp_queues(qdev) != 0) {
21        netdev_err(qdev->ndev, ...);
22        goto err_req_rsp;
23    }
24
25     if (ql_alloc_buffer_queues(qdev) != 0) {
26        netdev_err(qdev->ndev, ...);
27        goto err_buffer_queues;
28    }
29
30     if (ql_alloc_small_buffers(qdev) != 0) {
31        netdev_err(qdev->ndev, ...);
32        goto err_small_buffers;
33    }
34     ...
35    return 0;
36
37 err_small_buffers:
38    ql_free_buffer_queues(qdev);
39 err_buffer_queues:
40    ql_free_net_req_rsp_queues(qdev);
41 err_req_rsp:
42    ...
43    return -ENOMEM;
44 }
```

Function `ql_alloc_buffer_queues` employs **caller-based** management for `qdev->lrg_buf`.

When `ql_alloc_buffer_queues` fails, `qdev->lrg_buf` is not freed.

When the following calls fail, `qdev->lrg_buf` is freed through `ql_free_buffer_queues`.

IMMI determines that this function employs **callee-based** management.



Evaluation

Evaluation Settings

Environment

- Use a Linux server with 126 GB RAM and an Intel Xeon Silver 4316 CPU
- Use **Clang-15** to implement IMMI

Target

- The Linux kernel of **v5.18**

LLM settings

- Model: **GPT-4 (gpt-4-1106-preview)**
- Query each prompt 4 times, at least **3 responses should maintain consistent**
- Let LLM present intermediate results, subsequently synthesizing them into a definitive outcome
- Act as a post-filter for static analyzer

Evaluation - IMMI

Performance

Bitcode Loading	ICall Analysis	Summa- rization	Bug Analysis	LLM I/O	Total
1m 11s	3m 58s	2m 15s	1m 59s	25m 49s	35m 12s

Bug findings

Category	Reported Bugs		Real Bugs	Precision
	SA	SA+LLM		
Intro-inconsistency	63	52	35	67.3%
Inter-inconsistency	95	71	45	63.4%
Total	158	123	80	65.0%

80 new bugs: 57 memleak bugs, 16 double-free bugs, 6 UAF bugs, 1 null-pointer-dereference bug

Evaluation - IMMI

Comparison with existing tools

Tool	Methodology	Supported Bug Types	Analysis Time	Precision	Overlap with IMMI
IPPO	Similarity analysis	Memleak, memory corruption, ref-count leak, deadlock, missing check	2h	37%	28
HERO	Function pairing	Memleak, memory corruption, ref-count leak, deadlock	11h	52%	24
MLEE	Rules checking	Memleak	30 min	82%	26
Goshawk	Memory operation synopsis	Memory corruption	7h	63%	0
IMMI	MM intention analysis	Memleak, memory corruption	35 min	65%	80

IMMI effectively balances performance, precision, and the ability to uncover bugs.

Evaluation - LLM

Performance of LLM (GPT-4)

- Total bug reports: 158 -> 123
- Among the eliminated 35 reports, 32 are valid (false positives of IMMI)
- Overall false discovery rate of IMMI: 47.5% -> 35.0%

Comparison with different LLMs

LLM	Valid Reports	TP	FP	TN	FN
GPT-4	128 (81.0%)	67	26	32	3
GPT-3.5	140 (88.6%)	41	33	36	30
Llama-3	144 (91.1%)	58	47	19	20

Evaluation - LLM

Effectiveness of LLM (GPT-4)

```
1 /* drivers/dax/bus.c */
2 static int devm_register_dax_mapping(struct dev_dax *dev_dax, ...)
3 {
4     ...
5     mapping = kzalloc(sizeof(*mapping), GFP_KERNEL);
6     if (!mapping)
7         return -ENOMEM;
8     ...
9     mapping->id = ida_alloc(&dev_dax->ida, GFP_KERNEL);
10    if (mapping->id < 0) {
11        kfree(mapping);
12        return -ENOMEM;
13    }
14    ...
15    rc = device_add(dev);
16    if (rc) {
17        put_device(dev);
18        return rc;
19    }
20    ...
21    rc = devm_add_action_or_reset(..., unregister_dax_mapping,
22                                dev);
23    if (rc)
24        return rc;
25    return 0;
26 }
```

➤ Reference counting analysis

· `device_add(dev)` is called, and if it fails, `put_device(dev)` is called, **which will decrement the reference count and should trigger the release of the device and its associated memory if the count reaches zero.**

➤ Indirect release call analysis

· `devm_add_action_or_reset` is called to add a cleanup action (`unregister_dax_mapping`) to the device-managed resource list of `dax_region->dev`. If this fails, the function returns the error code, **and the device-managed resource list will take care of cleaning up the resources.**

Evaluation - LLM

False negative of LLM (GPT-4)

```
1 static int qedf_alloc_global_queues(struct qedf_ctx *qedf)
2 {
3     ...
4     qedf->global_queues = kzalloc(...);
5     if (!qedf->global_queues) {
6         QEDF_ERR(&(qedf->dbg_ctx), "Unable to allocate global "
7             "queues array ptr memory\n");
8         return -ENOMEM;
9     }
10    ...
11    status = qedf_alloc_bdq(qedf);
12    if (status) {
13        QEDF_ERR(&qedf->dbg_ctx, "Unable to allocate bdq.\n");
14        goto mem_alloc_failure;
15    }
16    ...
17    return 0;
18
19 mem_alloc_failure:
20     qedf_free_global_queues(qedf);
21     return status;
22 }
```

• The function `qedf_free_global_queues` is not defined in the provided code snippet, but **based on the naming convention and typical practices in C, we can infer that this function is responsible for freeing the resources allocated for the global queues.** Since all error paths after the allocation of `qedf->global_queues` lead to this function call, we can conclude that the function is designed to free the allocated heap memory.

qedf_free_global_queues() does not actually release *qedf->global_queues*

Conclusion

- **Inconsistent memory management intentions could lead to many kernel memory bugs like memleaks and memory corruptions.**
- **We presented IMMI to detect kernel memory bugs.**
 - Demand-driven summarization.
 - Object-based code slicing.
 - Combining program analysis and LLM in MM intention inference.
- **We evaluated IMMI on the Linux kernel**
 - Find 80 new memory bugs.
 - IMMI could effectively detect bugs that missed by existing tools



Dinghao Liu: dinghao.liu@zju.edu.cn

Zhipeng Lu: alexious@zju.edu.cn