

# Feature Consistency in Compile-Time-Configurable System Software

## Facing the Linux 10000 Feature Problem

**Reinhard Tartler**, Daniel Lohmann, Julio Sincero,  
Wolfgang Schröder-Preikschat

System Software Group

Friedrich-Alexander University  
Erlangen-Nuremberg

September 7, 2011



supported by **DFG**

# Configuration Complexity

- Linux has become incredibly configurable



# Configuration Complexity

- Linux has become incredibly configurable
- Complexity increases considerably




# Configuration Complexity

- Linux has become incredibly configurable
- Complexity increases considerably

→ Source of **bugs!**





Linux v3.0 contains:

**7.702** Features

**893** Kconfig files


**31.281** Source files

**88.897** #ifdef blocks



# The Problem

---



Configuration

**Source of  
Inconsistencies!**

Implementation



- Bugs in **declaration** and **implementation**
- Excellent tool support for **static analysis**:
  - Coccinelle: Faults in Linux: Ten Years Later (ASPLOS'11)
  - Dingo: Taming Device Drivers (EuroSys'09)
  - KLEE: Automatic generation of high-coverage tests (EuroSys'08)
  - RWset: Attacking path explosion (TACAS'08)
  - EXE: Automatically generating inputs of death (CCS'06)
  - ...





# Finding Bugs with Tools for Static Analysis

- Bugs in **declaration** and **implementation**
- Excellent tool support for **static analysis**:
  - Coccinelle: Faults in Linux: Ten Years Later (ASPLOS'11)
  - Dingo: Taming Device Drivers (EuroSys'09)
  - KLEE: Automatic generation of high-coverage tests (EuroSys'08)
  - RWset: Attacking path explosion (TACAS'08)
  - EXE: Automatically generating inputs of death (CCS'06)
  - ...
- Each of them check a **single** configuration:



# Symbolic Inconsistency

```
config HOTPLUG_CPU
  bool "Support for hot-pluggable CPUs"
  depends on SMP && HOTPLUG
  ---help---
```



# Symbolic Inconsistency

```
config HOTPLUG_CPU
    bool "Support for hot-pluggable CPUs"
    depends on SMP && HOTPLUG
    ---help---
```

```
static int
hotplug_cfd(struct notifier_block *nfb, unsigned long action, void *hcpu)
{
    // [...]
    switch (action) {
        case CPU_UP_PREPARE:
        case CPU_UP_PREPARE_FROZEN:
            // [...]
    #ifdef CONFIG_CPU_HOTPLUG
        case CPU_UP_CANCELED:
        case CPU_UP_CANCELED_FROZEN:

        case CPU_DEAD:
        case CPU_DEAD_FROZEN:
            free_cpumask_var(cfd->cpumask);
            break;
    #endif
    };
    return NOTIFY_OK;
```



# Symbolic Inconsistency


```
config HOTPLUG_CPU
    bool "Support for hot-pluggable CPUs"
    depends on SMP && HOTPLUG
    ---help---
```

```
static int
hotplug_cfd(struct notifier_block *nb, unsigned long action, void *hcpu)
{
    // [...]
    switch (action) {
        case CPU_UP_PREPARE:
        case CPU_UP_PREPARE_FROZEN:
            // [...]
#ifdef CONFIG_CPU_HOTPLUG
        case CPU_UP_CANCELED:
        case CPU_UP_CANCELED_FROZEN:


        case CPU_DEAD:
        case CPU_DEAD_FROZEN:
            free_cpumask_var(cfd->cpumask);
            break;
#endif
    };
    return NOTIFY_OK;
}
```

Symbolic ↯

■ Result: Fix for a **critical bug**




# Logic Inconsistencies



```
#ifdef CONFIG_DISCONTIGMEM
    // Block1
    static ... int pfn_to_mid(...)
#endif
#ifdef CONFIG_NUMA
    // Block2
#else
    // Block3
#endif
#endif
```




# Logic Inconsistencies



- Feature DISCONTIGMEM requires NUMA
- Inner block is not configuration dependent anymore




# Logic Inconsistencies



- Feature DISCONTIGMEM requires NUMA
- Inner block is not configuration dependent anymore
- Result: **code cleanup**



# General Approach




```
#ifdef CONFIG_DISCONTIGMEM
    // Block1
    static ... int pfn_to_mid(...)
#endif
#ifdef CONFIG_NUMA
    // Block2
#else
    // Block3
#endif
#endif
```






# General Approach



```
#ifdef CONFIG_DISCONTIGMEM
    // Block1
    static ... int pfn_to_mid(...)
#endif
#ifdef CONFIG_NUMA
    // Block2
#else
    // Block3
#endif
#endif
```




# General Approach



```
#ifdef CONFIG_DISCONTIGMEM
// Block1
static ... int pfn_to_mid(...)
#endif
#ifdef CONFIG_NUMA
// Block2
#else
// Block3
#endif
#endif
```


$$\begin{aligned} C = & (\text{FLATMEM} \rightarrow \text{MEMORY\_MODEL}) \\ & \wedge (\text{DISCONTIGMEM} \rightarrow \text{MEMORY\_MODEL}) \\ & \wedge (\text{SPARSEMEM} \rightarrow \text{MEMORY\_MODEL}) \\ & \wedge (\text{NUMA} \rightarrow \text{MEMORY\_MODEL}) \\ & \wedge (\text{DISCONTIGMEM} \rightarrow \text{NUMA}) \end{aligned}$$


# General Approach



```
#ifdef CONFIG_DISCONTIGMEM
    // Block1
    static ... int pfn_to_mid(...)
#endif
#ifdef CONFIG_NUMA
    // Block2
#else
    // Block3
#endif
#endif
```


extract

$$\begin{aligned} C = & (\text{FLATMEM} \rightarrow \text{MEMORY\_MODEL}) \\ & \wedge (\text{DISCONTIGMEM} \rightarrow \text{MEMORY\_MODEL}) \\ & \wedge (\text{SPARSEMEM} \rightarrow \text{MEMORY\_MODEL}) \\ & \wedge (\text{NUMA} \rightarrow \text{MEMORY\_MODEL}) \\ & \wedge (\text{DISCONTIGMEM} \rightarrow \text{NUMA}) \end{aligned}$$

extract

$$\begin{aligned} I = & (\text{Block}_1 \leftrightarrow \text{DISCONTIGMEM}) \\ & \wedge (\text{Block}_2 \leftrightarrow \text{Block}_1 \wedge (\text{NUMA})) \\ & \wedge (\text{Block}_3 \leftrightarrow \text{Block}_1 \wedge \neg \text{Block}_2) \end{aligned}$$


# General Approach



Crosscheck both formulas with a SAT solver:

$$\mathbf{dead?} = \text{sat}(C \wedge \mathcal{I} \wedge \text{Block}_N)$$

$$\mathbf{undead?} = \text{sat}(C \wedge \mathcal{I} \wedge \neg \text{Block}_N \wedge \text{parent}(\text{Block}_N))$$

$$\wedge (\text{DISCONTIGMEM} \rightarrow \text{MEMORY\_MODEL})$$

$$\wedge (\text{SPARSEMEM} \rightarrow \text{MEMORY\_MODEL})$$

$$\wedge (\text{NUMA} \rightarrow \text{MEMORY\_MODEL})$$

$$\wedge (\text{DISCONTIGMEM} \rightarrow \text{NUMA})$$

$$\wedge (\text{Block}_2 \leftrightarrow \text{Block}_1 \wedge (\text{NUMA}))$$

$$\wedge (\text{Block}_3 \leftrightarrow \text{Block}_1 \wedge \neg \text{Block}_2)$$



## ■ Accuracy

- Conceptually **no false positives**
- **Exact** identification of variation points



## ■ Accuracy

- Conceptually **no false positives**
- **Exact** identification of variation points

## ■ Coverage

- Extract configuration model for all **22 architectures**
- Defect  $\rightsquigarrow$  detected on **each** architecture



## ■ Accuracy

- Conceptually **no false positives**
- **Exact** identification of variation points

## ■ Coverage

- Extract configuration model for all **22 architectures**
- Defect  $\rightsquigarrow$  detected on **each** architecture

## ■ Performance

- Easy and fast to use during **incremental builds**
- Possible by **problem slicing**
- **Complete run** on Linux in less than **10** minutes



# Results

subsystem	#ifdefs	logic	symbolic	total
arch/	33757	345	581	926
drivers/	32695	88	648	736
fs/	3000	4	13	17
include/	7241	6	11	17
kernel/	1412	7	2	9
mm/	555	0	1	1
net/	2731	1	49	50
sound/	3246	5	10	15
virt/	53	0	0	0
other subsystems	601	4	1	5
$\Sigma$	85291	460	1316	1776
fix proposed		150 (1)	214 (22)	364 (23)
confirmed defect		38 (1)	116 (20)	154 (21)
confirmed rule-violation		88 (0)	21 (2)	109 (2)
pending		24 (0)	77 (0)	101 (0)





## Results

We have found **1776** configurability issues

Submitted **123** patches for **364** defects

**20** are confirmed **new bugs** (affecting binary code)

Cleaned up **5129** lines of cruft code

pending

24 (0)

77 (0)

101 (0)



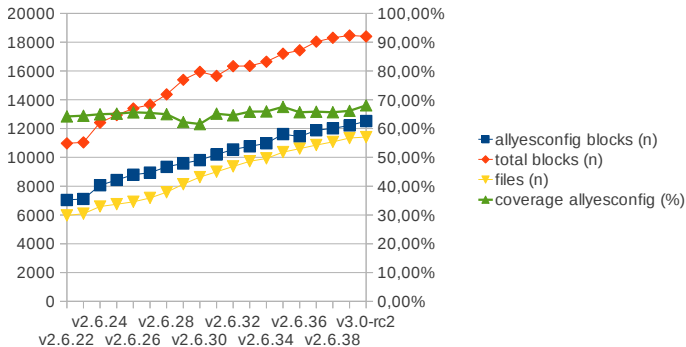
## Further Application: Configuration Coverage

---

- Current ongoing work, accepted at PLOS'11
- Configuration Coverage is defined as:  
fraction of *selected* configuration-conditional blocks  
divided by the number of *available* configuration-conditional blocks.
- How to catch bugs that apply only on specific kernel configurations?  
⇒ Test them on as many configurations as possible
- Static analyzers (sparse, smatch, ...) scan a particular kernel configuration  
⇒ How to *effeciently* exand their coverage?



# Historical analysis of allyes coverage



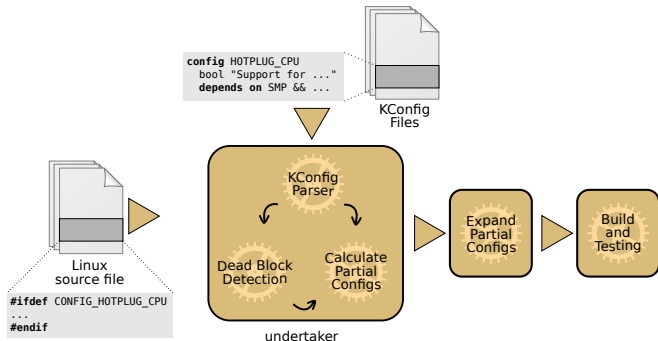
# Concrete Example

```
#ifdef CONFIG_DISCONTIGMEM Block 1
static inline int pfn_to_nid(unsigned long pfn)
{
#ifdef CONFIG_NUMA Block 2
    return((int) physnode_map[(pfn) / PAGES_PER_ELEMENT]);
#else Block 3
    return 0;
#endif
}
#endif Block 1
```

- Possible Configurations:
  - Neither, DISCONTIGMEM, DISCONTIGMEM  $\wedge$  NUMA
- Additionally testing the configuration NUMA does not increase the Configuration Coverage.



# Realization



# Implementation Challenges

- Proper extraction of Configurations constraints
  - Kconfig (implemented in undertaker)
  - Kbuild constraints (largely unhandled)
- Expansion of Partial Configurations
  - Naïve approach has some surprising effects (i.e., fails sometimes)
  - Kconfig-sat seems promising, but unfortunately discontinued



## Evaluation

---

Analyzed files	10,365
Files with variability	3,163
Rate of files with variability	30.52%
Sum of all (partial) configurations	4,435
<hr/>	
Sum of <i>configuration controlled</i> conditional blocks	16,444
Sum of blocks selected by allyesconfig	11,511
Sum of all blocks selected by undertaker-coverage	13,844
Coverage allyesconfig (non-dead-corrected)	70.00%
Coverage undertaker (non-dead-corrected)	84.19%
<hr/>	
Dead blocks	1,778
Selectable blocks (excluding dead blocks)	14,666
Selected by allyesconfig	11,511
Covered by undertaker	13,844
allyesconfig coverage	78.49%
undertaker coverage	94.40%
undertaker coverage / allyesconfig coverage	1.20

---



## Evaluation

---

Analyzed files	10,365
<b>Files with variability</b>	<b>3,163</b>
Rate of files with variability	30.52%
<b>Sum of all (partial) configurations</b>	<b>4,435</b>
<hr/>	
Sum of <i>configuration controlled</i> conditional blocks	16,444
Sum of blocks selected by allyesconfig	11,511
Sum of all blocks selected by undertaker-coverage	13,844
Coverage allyesconfig (non-dead-corrected)	70.00%
Coverage undertaker (non-dead-corrected)	84.19%
<hr/>	
Dead blocks	1,778
Selectable blocks (excluding dead blocks)	14,666
Selected by allyesconfig	11,511
Covered by undertaker	13,844
allyesconfig coverage	78.49%
undertaker coverage	94.40%
undertaker coverage / allyesconfig coverage	1.20

---





## Evaluation

---

Analyzed files	10,365
<b>Files with variability</b>	<b>3,163</b>
Rate of files with variability	30.52%
<b>Sum of all (partial) configurations</b>	<b>4,435</b>
<hr/>	
Sum of <i>configuration controlled</i> conditional blocks	16,444
Sum of blocks selected by allyesconfig	11,511
Sum of all blocks selected by undertaker-coverage	13,844
<b>Coverage allyesconfig (non-dead-corrected)</b>	<b>70.00%</b>
<b>Coverage undertaker (non-dead-corrected)</b>	<b>84.19%</b>
<hr/>	
Dead blocks	1,778
Selectable blocks (excluding dead blocks)	14,666
Selected by allyesconfig	11,511
Covered by undertaker	13,844
allyesconfig coverage	78.49%
undertaker coverage	94.40%
undertaker coverage / allyesconfig coverage	1.20

---



## Evaluation

---

Analyzed files	10,365
<b>Files with variability</b>	<b>3,163</b>
Rate of files with variability	30.52%
<b>Sum of all (partial) configurations</b>	<b>4,435</b>
<hr/>	
Sum of <i>configuration controlled</i> conditional blocks	16,444
Sum of blocks selected by allyesconfig	11,511
Sum of all blocks selected by undertaker-coverage	13,844
<b>Coverage allyesconfig (non-dead-corrected)</b>	<b>70.00%</b>
<b>Coverage undertaker (non-dead-corrected)</b>	<b>84.19%</b>
<hr/>	
Dead blocks	1,778
Selectable blocks (excluding dead blocks)	14,666
Selected by allyesconfig	11,511
Covered by undertaker	13,844
allyesconfig <b>coverage</b>	<b>78.49%</b>
undertaker <b>coverage</b>	<b>94.40%</b>
undertaker coverage / allyesconfig coverage	1.20

---



## Evaluation

Analyzed files	10,365
<b>Files with variability</b>	<b>3,163</b>
Rate of files with variability	30.52%
<b>Sum of all (partial) configurations</b>	<b>4,435</b>

With **30 percent** more compiler calls  
(static analysis runs)

We get **15 percent** more Configuration Coverage

Covered by undertaker	13,844
allyesconfig <b>coverage</b>	<b>78.49%</b>
undertaker <b>coverage</b>	<b>94.40%</b>
undertaker coverage / allyesconfig coverage	1.20



# Conclusions

---

- Configurability has to be seen as a significant cause of software defects in its own respect
- **Configuration** and **implementation** need to be kept consistent
- Configuration Coverage increases the effectiveness of existing tools.



# Conclusions

---

- Configurability has to be seen as a significant cause of software defects in its own respect
- **Configuration** and **implementation** need to be kept consistent
- Configuration Coverage increases the effectiveness of existing tools.
- Vision:
  - **Explorative** tool for **visualizing** and **checking** Variability in Kconfig and realization
  - Linux Feature Explorer (**LIFE**)



# Conclusions

---

- Configurability has to be seen as a significant cause of software defects in its own respect
- **Configuration** and **implementation** need to be kept consistent
- Configuration Coverage increases the effectiveness of existing tools.
- Vision:
  - **Explorative** tool for **visualizing** and **checking** Variability in Kconfig and realization
  - Linux Feature Explorer (**LIFE**)

<http://vamos.informatik.uni-erlangen.de/trac/undertaker>

