

# Many-Core Compiler Fuzzing

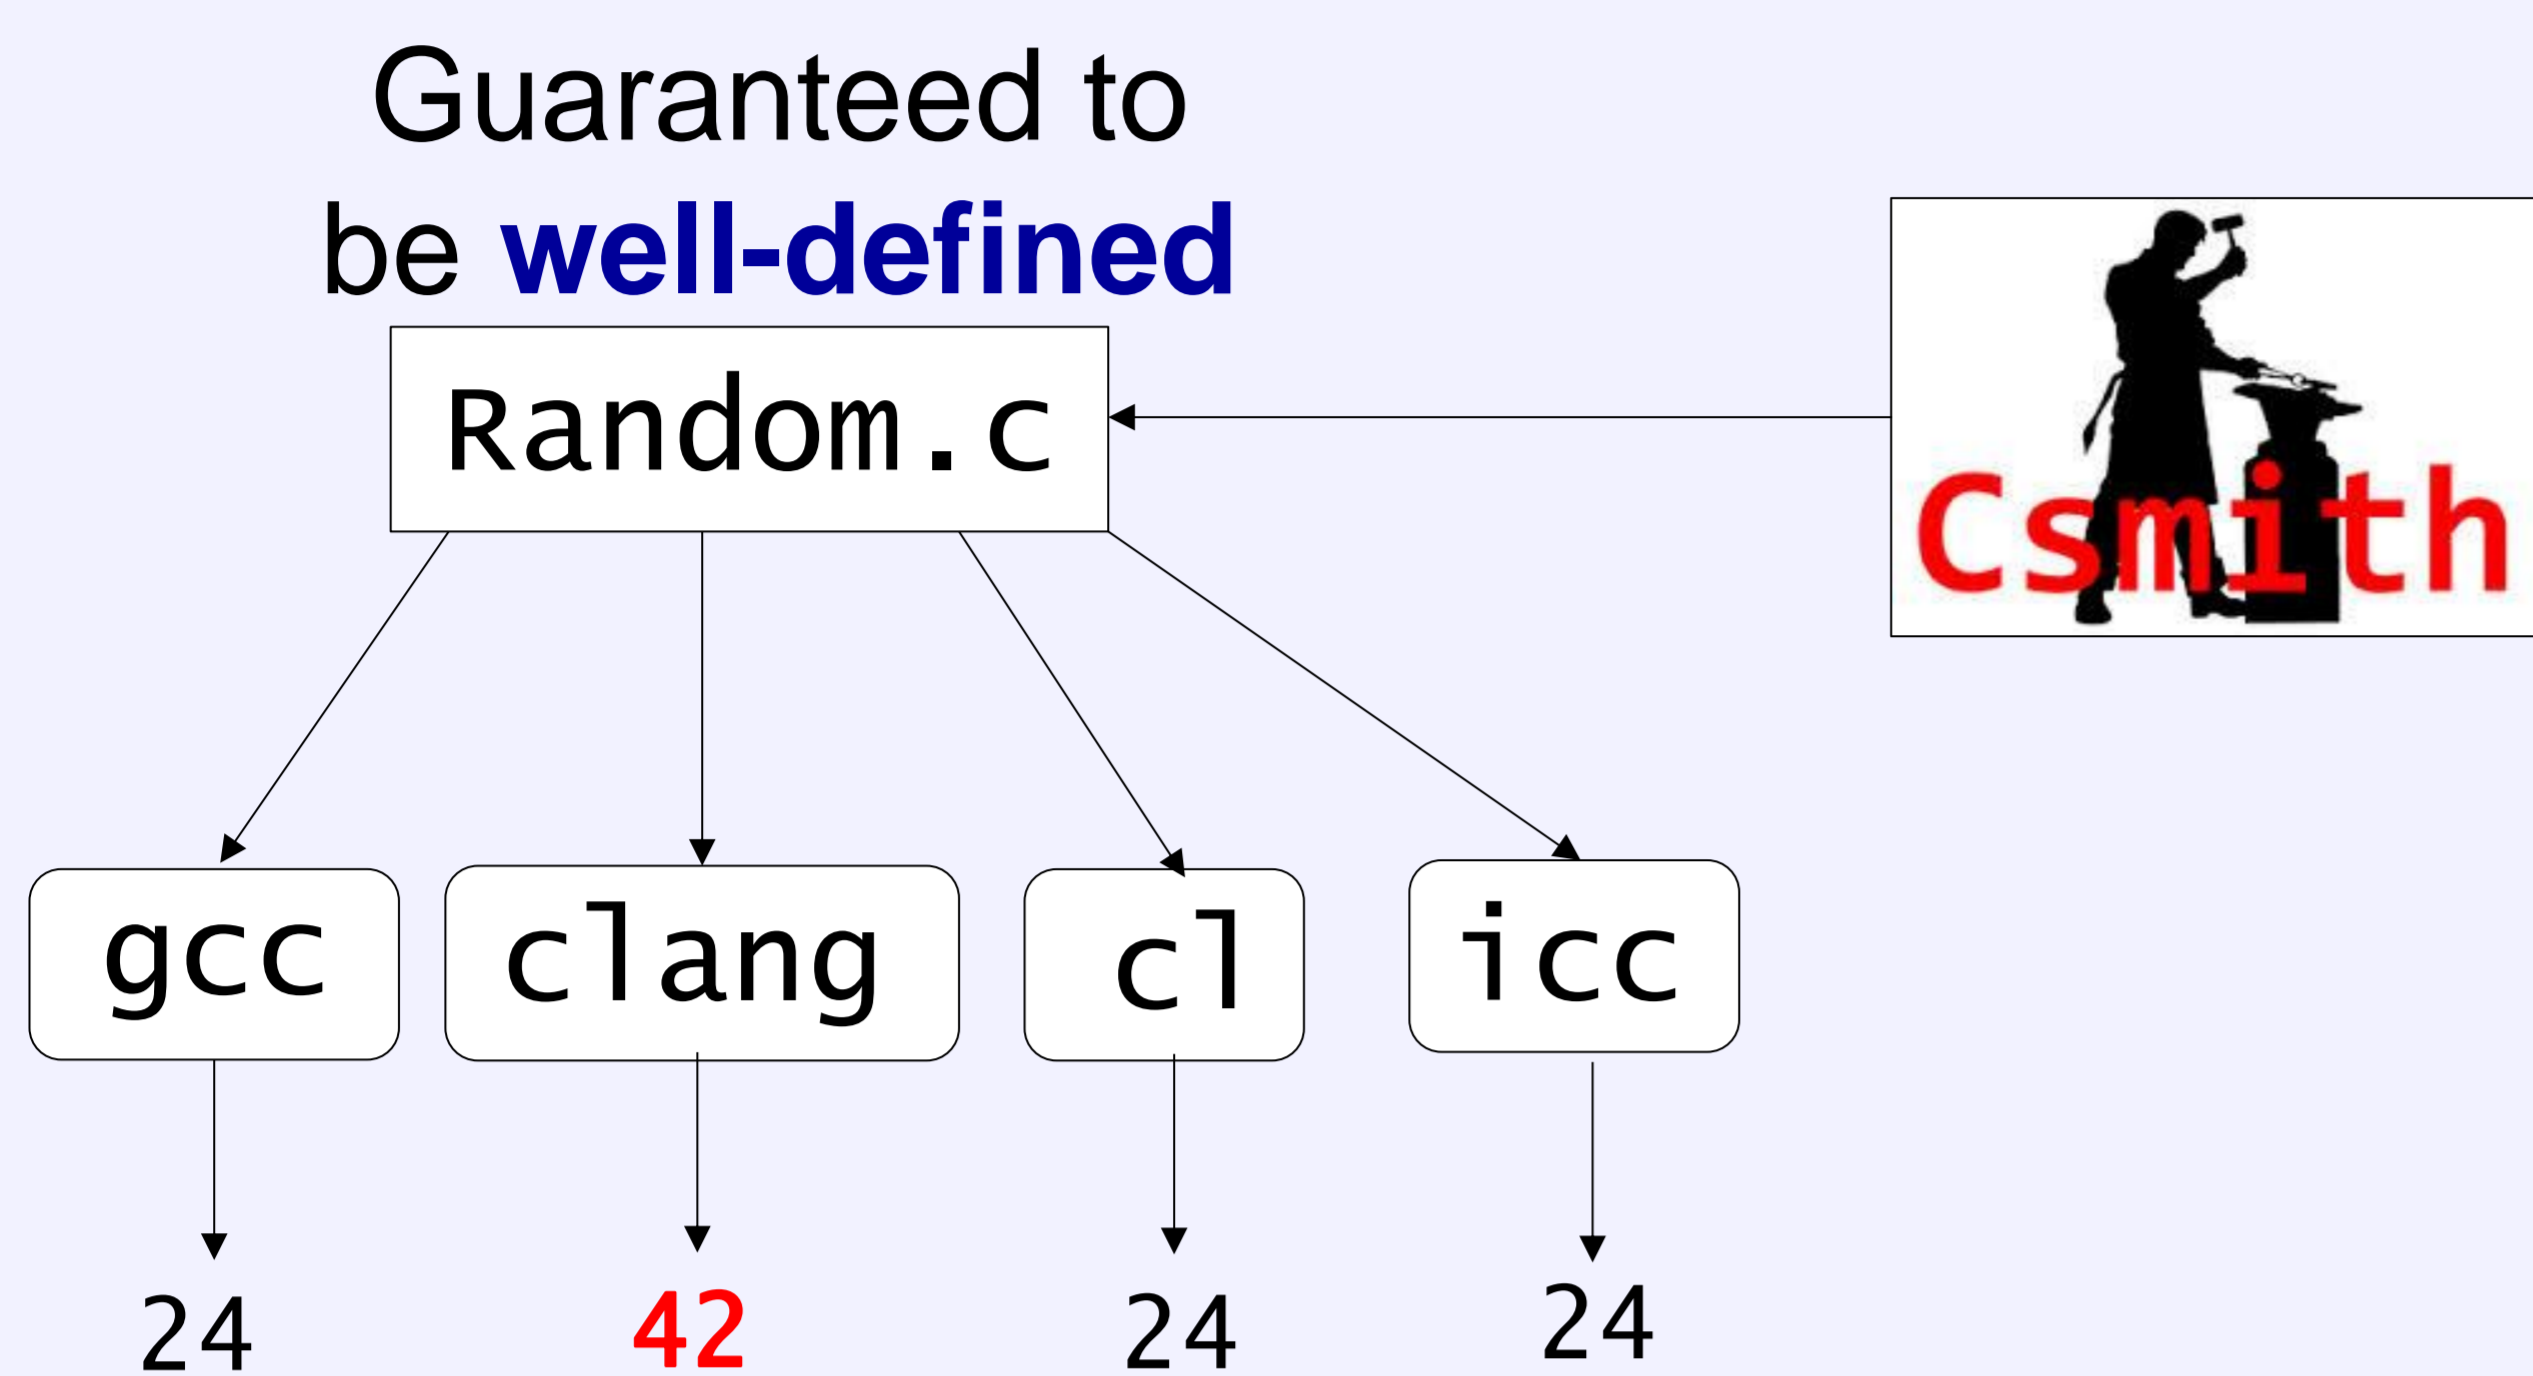


Christopher Lidbury,<sup>1</sup> Andrei Lascau,<sup>1</sup> Nathan Chong,<sup>2</sup> Alastair F. Donaldson<sup>1</sup>  
<sup>1</sup>Imperial College London, <sup>2</sup>University College London

**Research aim:** validate effectiveness of existing compiler testing methods...

## Random Differential Testing

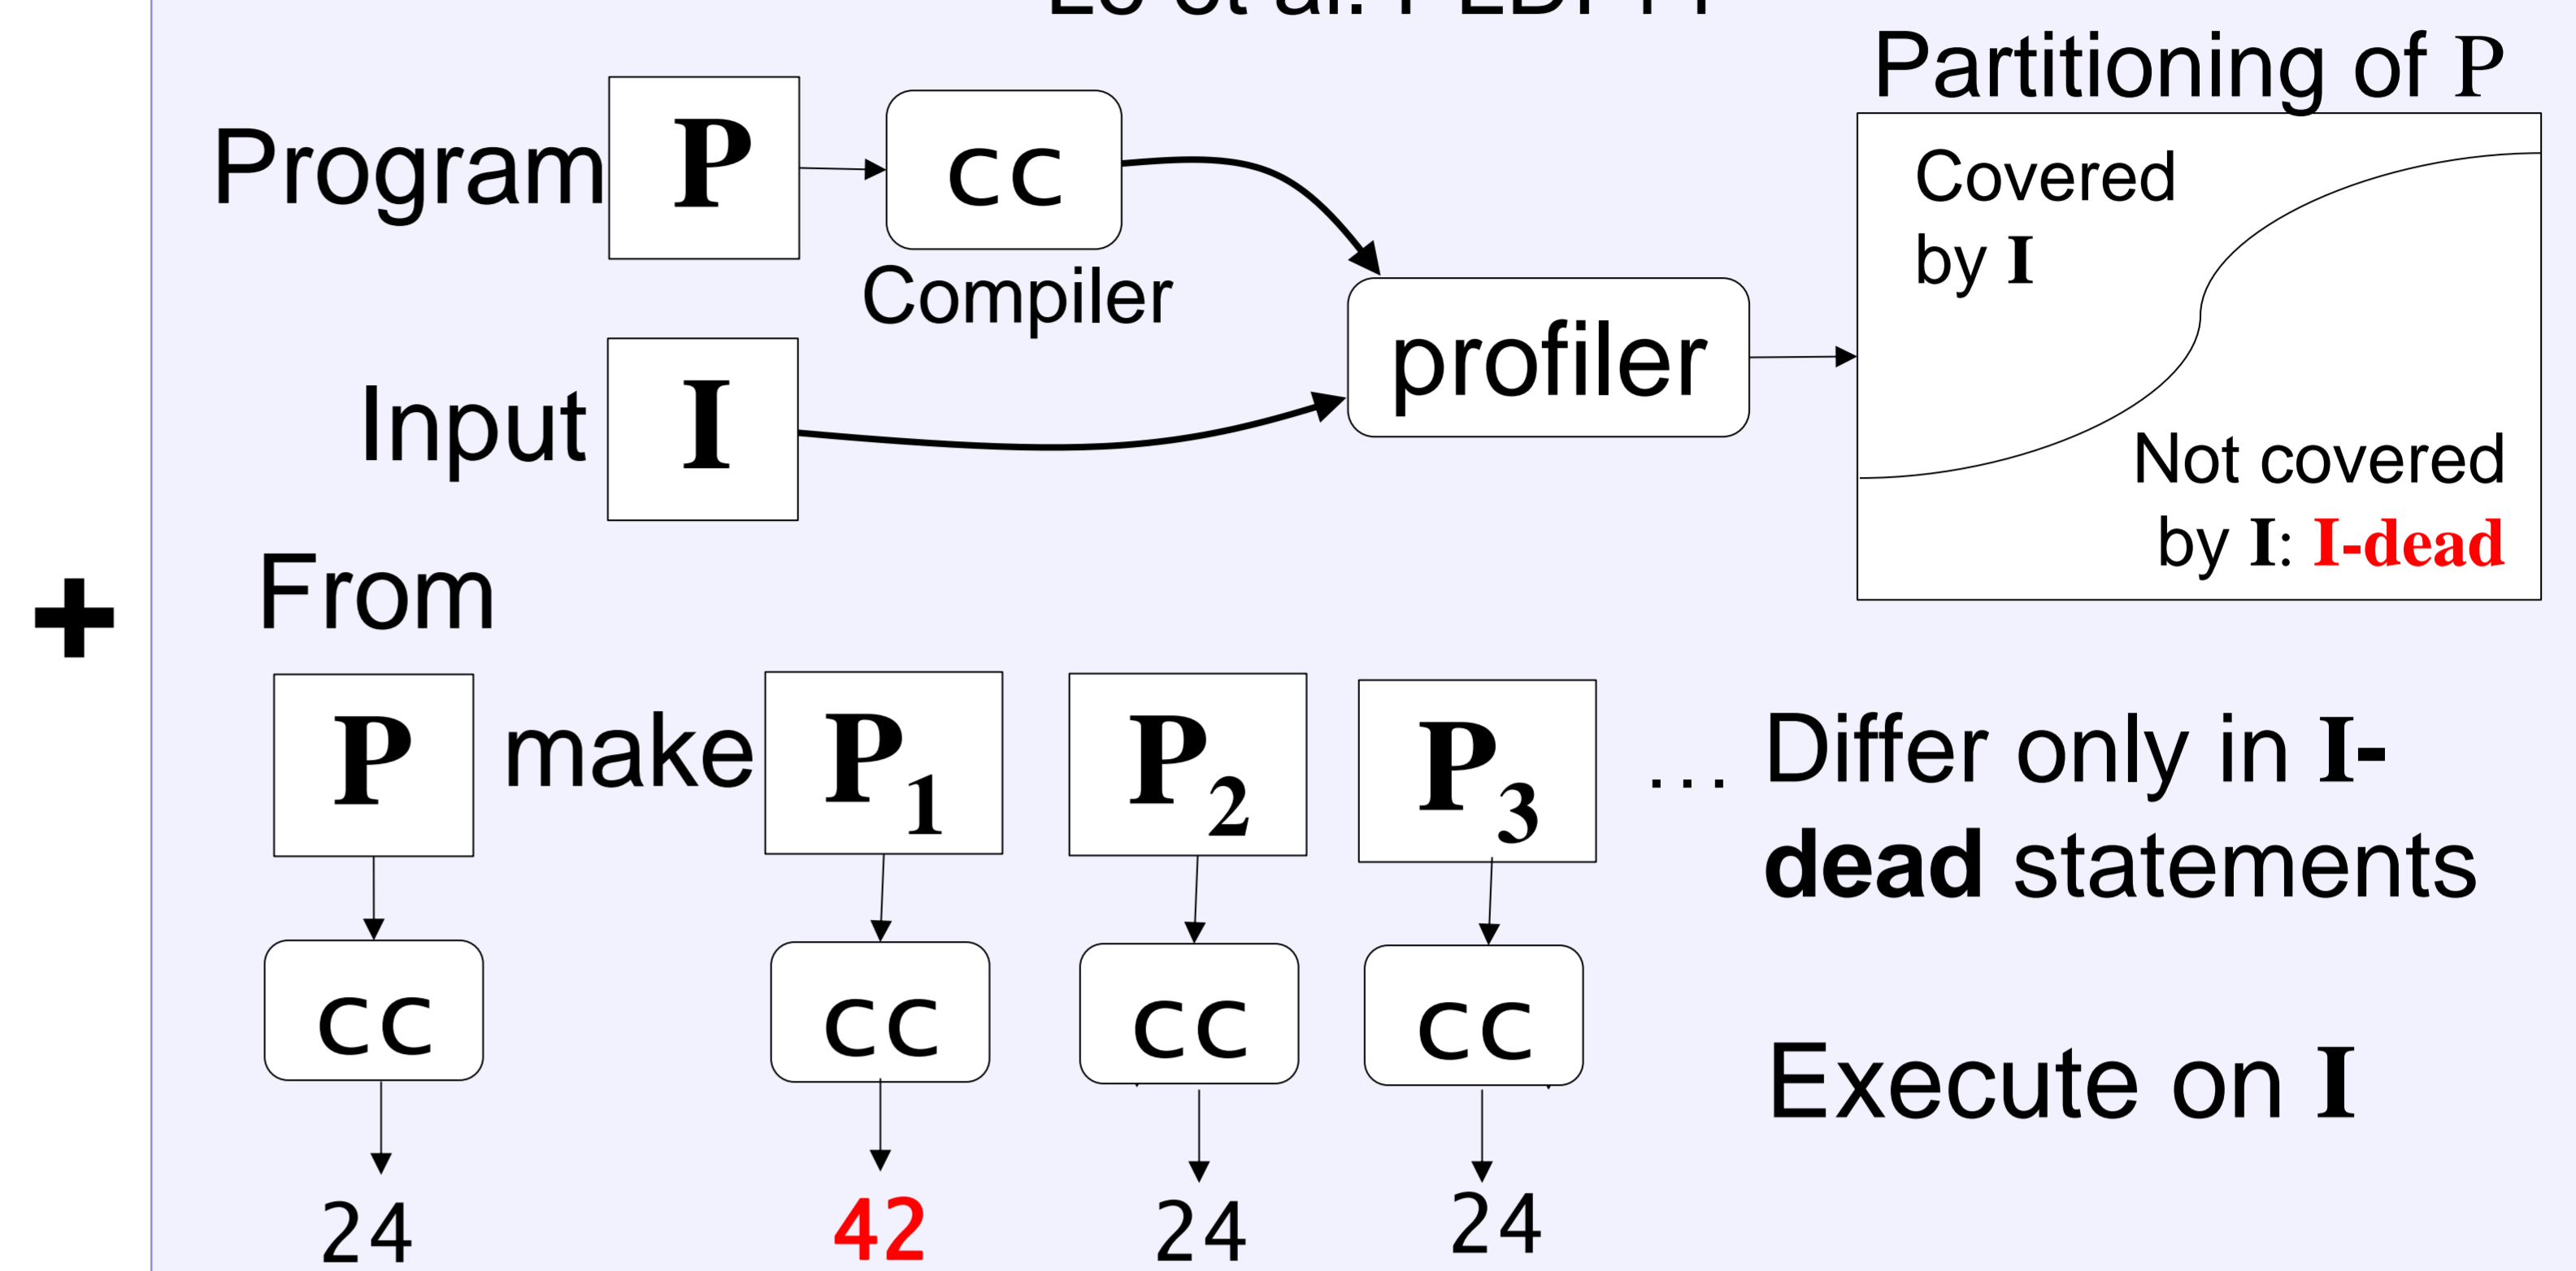
Yang et al. PLDI'11



Cross-check **multiple compilers** on random programs. **Mismatches** indicate **bugs**.

## Equivalence Modulo Inputs Testing

Le et al. PLDI'14



Cross-check **single compiler** using EMI variants. **Mismatches** indicate **bugs**.

...in a many-core domain: **OpenCL**. Our innovations include:

## Generating Random Kernels

Our **CLsmith** tool generates **concurrent-yet-deterministic** OpenCL kernels

**Vector mode:** exercises **vector operations**

```
uint4 v = clamp(a, b.xxyy, clz(c));
```

**Barrier mode:** orchestrates **race-free** shared memory communication via **barriers**

**Atomic mode:** exercises **read-modify-write** operations in a deterministic way

```
if(atomic_inc(c) == 42) {
    // Thread-insensitive code
}
```

Which thread enters: **non-deterministic**  
 Effect of section: **deterministic**

## Dead-By-Construction Injection

Equip kernel with extra Boolean parameter, **DEAD**  
 Inject code conditional on **DEAD**

```
kernel void BFS(boo1 DEAD, global int *q1, ...) {
    if(get_local_id(0) == 0)
        *local_q_tail = 0;
    barrier();
    if( tid < no_of_nodes)
    {
        int pid = q1[tid];
        g_color[pid] = BLACK;
        int cur_cost = g_cost[pid];
        struct Node cur_node = g_graph_nodes[pid];
        for(int i=cur_node.x; i < cur_node.y + cur_node.x; i++)
            if(DEAD) {
                // ARBITRARY code
            }
    }
}
```

At runtime, set **DEAD** to **false**

Injected code is **dead by construction**

Enables **EMI testing** *without* the need to detect I-dead code

Tested 21 OpenCL (device, compiler) combinations; major vendors + open source



Results show that RDT and EMI are **effective** in this domain

More than 50 distinct bugs found, many now **fixed** by vendors

**Test programs and full results:** [multicore.doc.ic.ac.uk/tools/CLsmith](http://multicore.doc.ic.ac.uk/tools/CLsmith)

Thanks to support from:



Undergraduate Research Opportunities Programme