


Measuring and Modeling Data Movement Power on Real Hardware

Vignesh Adhinarayanan

Introduction

- Power and energy have become a first-class design constraint in all areas of computing

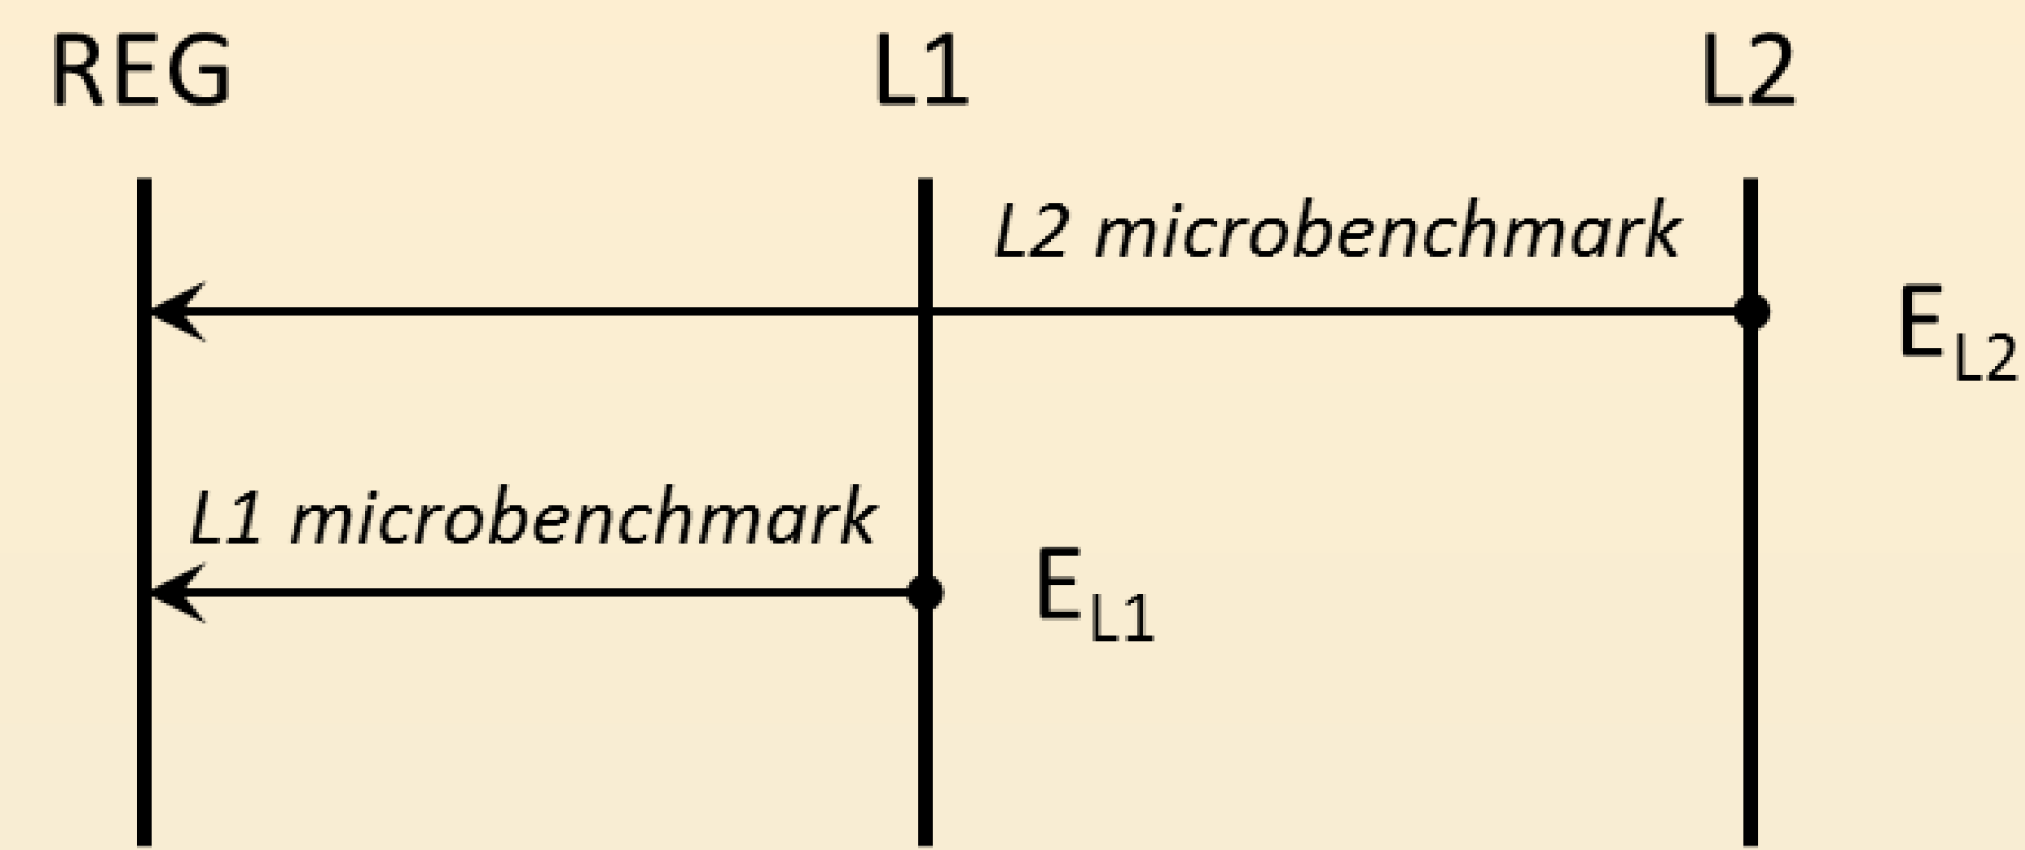


- Data movement thought to be a major source of power consumption

Objective

Measure and understand the nature of data-movement power on real hardware

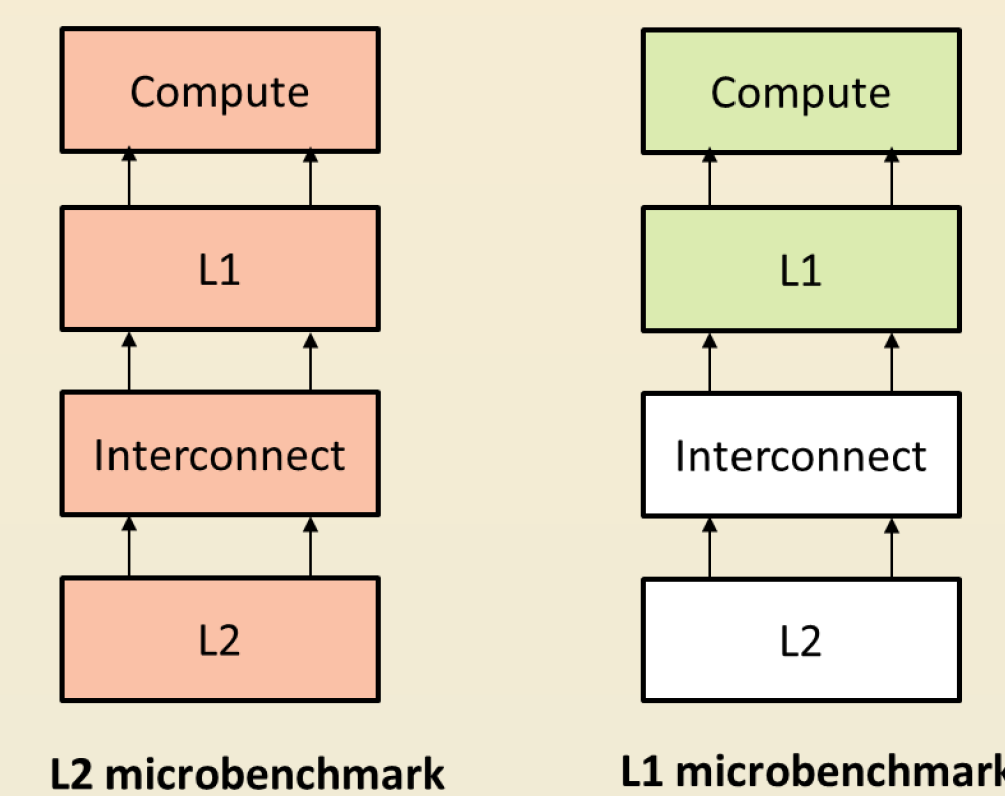
Related Work



G. Kestor et al., "Quantifying the energy cost of data movement in scientific applications," IISWC 2013

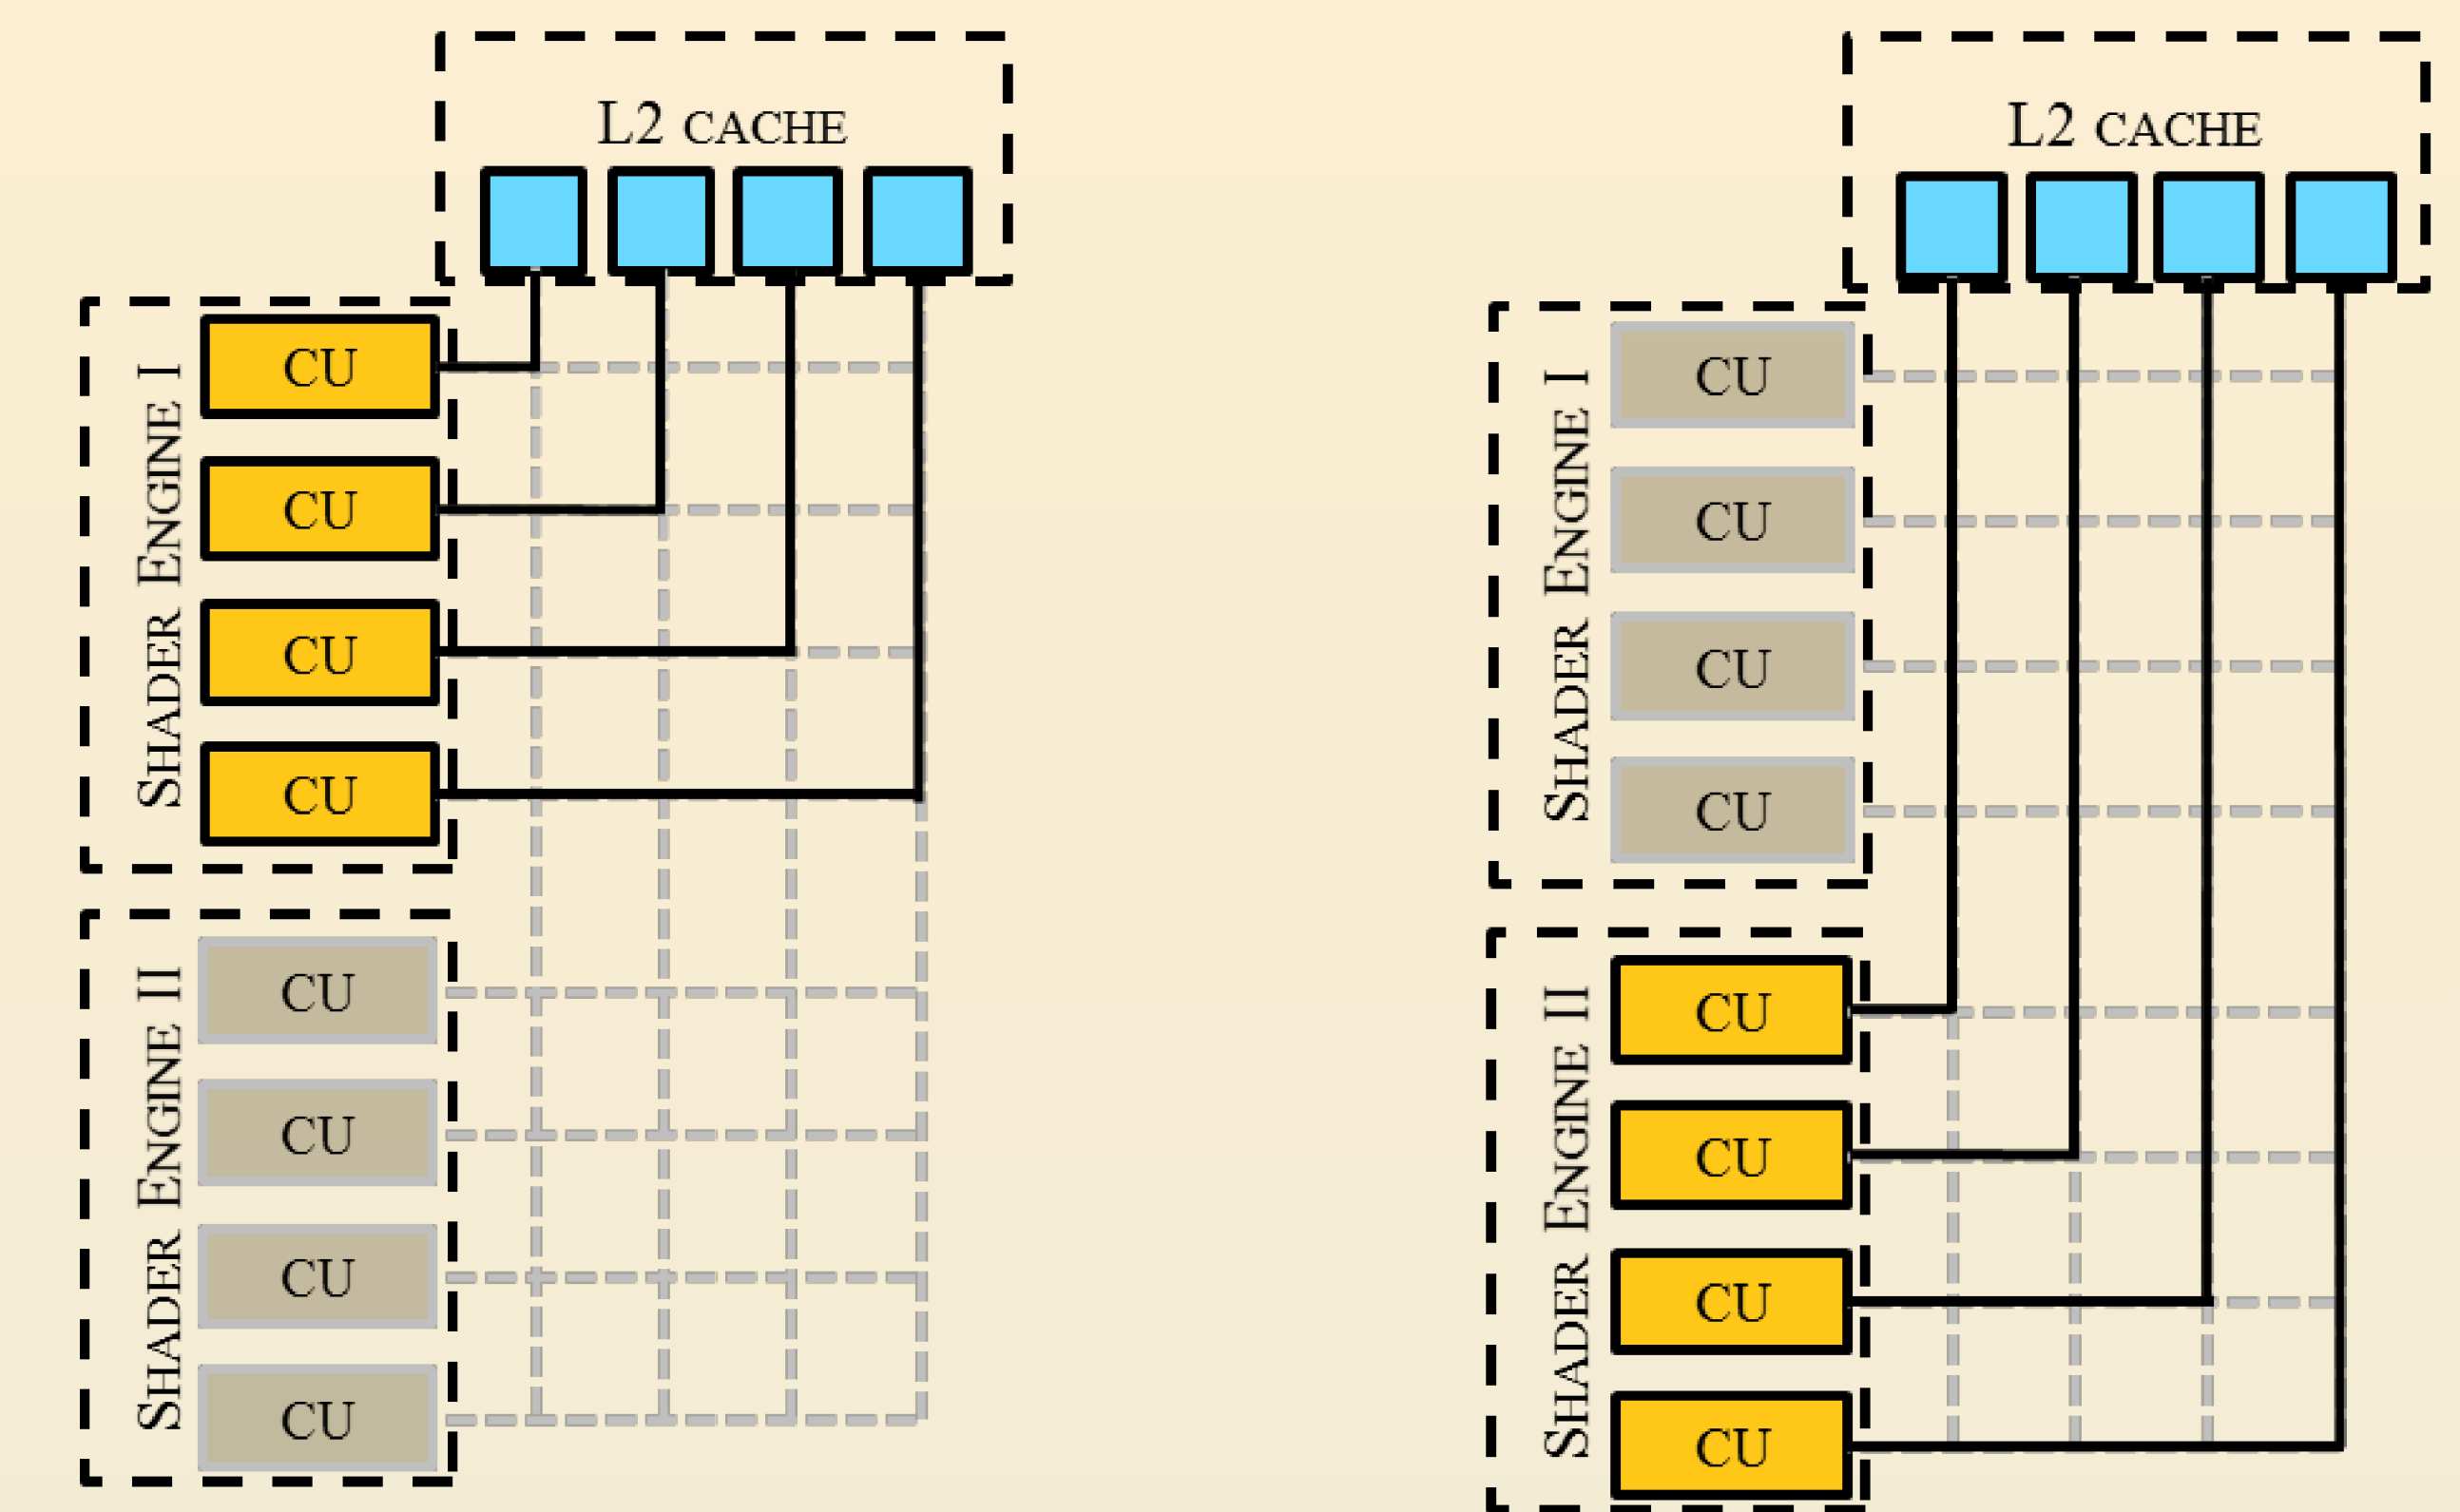
$$\text{Energy cost of moving data from L2 to L1} = E_{L2} - E_{L1}$$

Limitation



- L2 access power cannot be isolated

Our Approach



Short-path microbenchmark

Long-path microbenchmark

Design microbenchmarks based on data-movement distance as wire distance affects power

Challenges and Solutions

- OpenCL™ lacks native support to pin threads to programmer-specified compute units

```

(a) Initial OpenCL code snippet
__kernel void l2_read( __global float *data,
                    __global float *output) {
    int gid = get_global_id(0);
    if (gid >= 0 && gid <= 10) {
        // Read data from L2
    }
}

(b) Equivalent assembly code
s_min_w32    s0, s0, 0x0000ffff // 000000000014: 8300700 0000ffff
s_wd1_l32    s0, s16, s0        // 000000000015: 93000010
s_wd1_l32    s0, s0, s1        // 000000000020: 81000100
s_wd1_l32    w0, w0, s0, v0     // 000000000024: 68000000
s_wd1_l32    s0, s16, s2       // 000000000028: 81000210
s_cmp_gt_l32 s0, -1           // 000000000030: 8F000000
s_branch_scc label_0011      // 000000000030: 8F000004

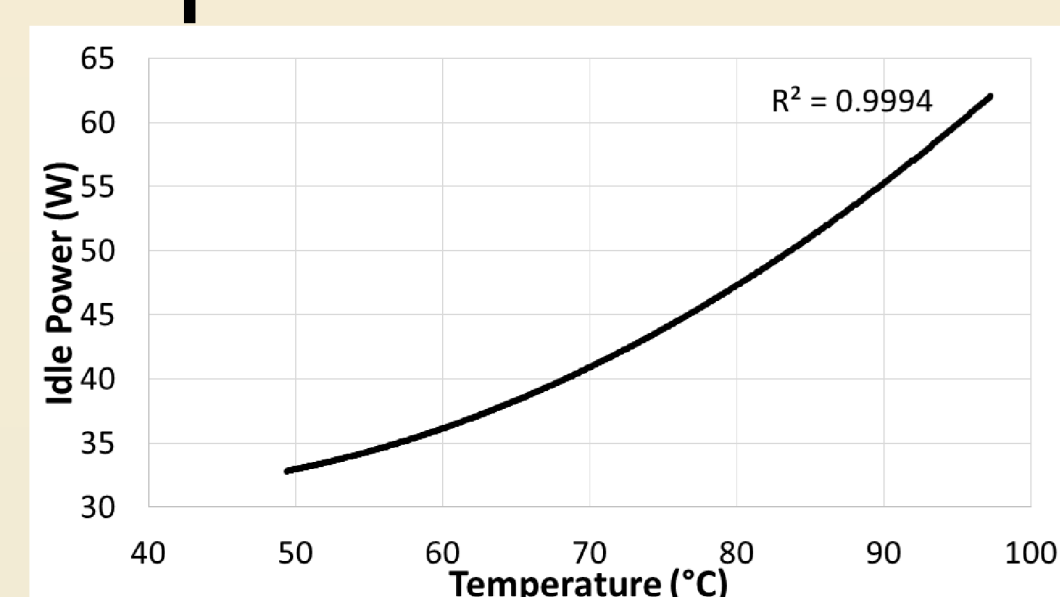
(c) Equivalent binary (in hex)
00 ff 80 83 ff ff 00 00
10 00 00 83 00 01 00 81
00 00 00 4a 10 02 00 81
00 c1 02 bf 04 00 84 bf
00 ff 04 bf 81 00 00 00
c1 80 01 85 01 00 82 bf

(d) Modified binary (in hex)
00 ff 80 83 ff ff 00 00
10 00 00 83 00 01 00 81
00 00 00 4a 04 04 32 00 89
00 c1 02 bf 04 00 84 bf
00 ff 04 bf 81 00 00 00
c1 80 01 85 01 00 82 bf

(e) Equivalent OpenCL code
__kernel void l2_read( __global float *data,
                    __global float *output) {
    int gid = get_global_id(0);
    if (gid >= 0 && gid <= 10) {
        // Read data from L2
    }
}
    
```


Solution: binary hacking

- Temperature of device during tests will affect the power consumption




Solution: Model idle power separately and subtract from measured power


Summary of Results



- Linear relationship between (i) distance and power
- (ii) toggle rate and power
- Asymmetrical power consumption for 0s and 1s



Optimization(s)



- On-chip data movement power reduces by 48% for an L1-L2 distance optimized layout
- More results in the paper