






MOBILE-DERIVATIVE TECHNOLOGY

Steve Oberlin
CTO, Accelerated Computing





The Attack of the Killer Micros

E. D. Brooks III

Massively Parallel Computing Initiative
Lawrence Livermore National Laboratories

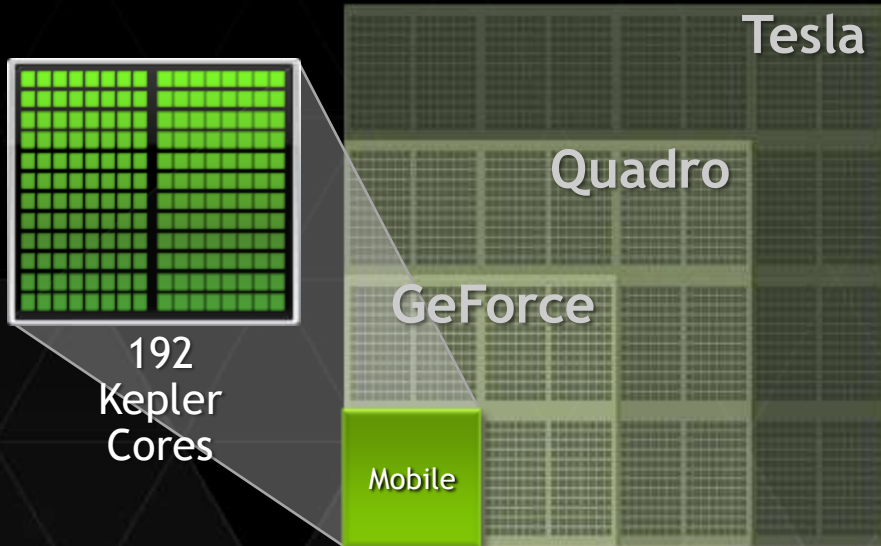

brooks@{maddog.llnl.gov, maddog.uucp}

Presented at: Supercomputing '89
Reno, Nov 13-17 1989

COMING SOON:
ATTACK OF THE KILLER SMARTPHONE

TEGRA K1


Unify GPU and Tegra Architecture



CUDA Enabled

JETSON TK1

Development Platform for Embedded Computer
Vision, Robotics, Medical



192 Kepler Cores · 326 GFLOPS

4 ARM A15 Cores

2 GB DDR3L

16-256 GB Flash

Gigabit Ethernet

CUDA Enabled

5-11 Watts

\$192

Available Now

“ATTACK OF THE KILLER SMARTPHONES?”

● K40 + CPU


- Peak SP: 4.2 TFLOPS
- SP SGEMM: ~3.8 TFLOPS
- Memory: 12 GB @ 288 GB/s
- Power:
 - GPU: 235 W
 - CPU + Mem: 150 W
 - Total: 385 W
- Perf/Watt: ~10 SP GFLOPS/W

● TK-1


- Peak SP: 326 GFLOPS
- SP SGEMM: ~290 GFLOPS
- Memory: 2 GB @ 14.9 GB/s
- Power:
 - GPU + CPU: <11 W (working hard)
 - **1/35 of K40 + CPU**
- **Perf/Watt: ~26 SP GFLOPS/W**

For the same power as K40 + CPU, you could have
10+ TFLOPS SP, 70 GB DRAM @ 500+ GB/s

TEGRA K1: NOW IN TWO FLAVORS



←→
Pin
Compatible



Quad A15 CPUs

32-bit

3-way Superscalar

Up to 2.3GHz

32K+32K L1\$

Dual Denver CPUs

64-bit

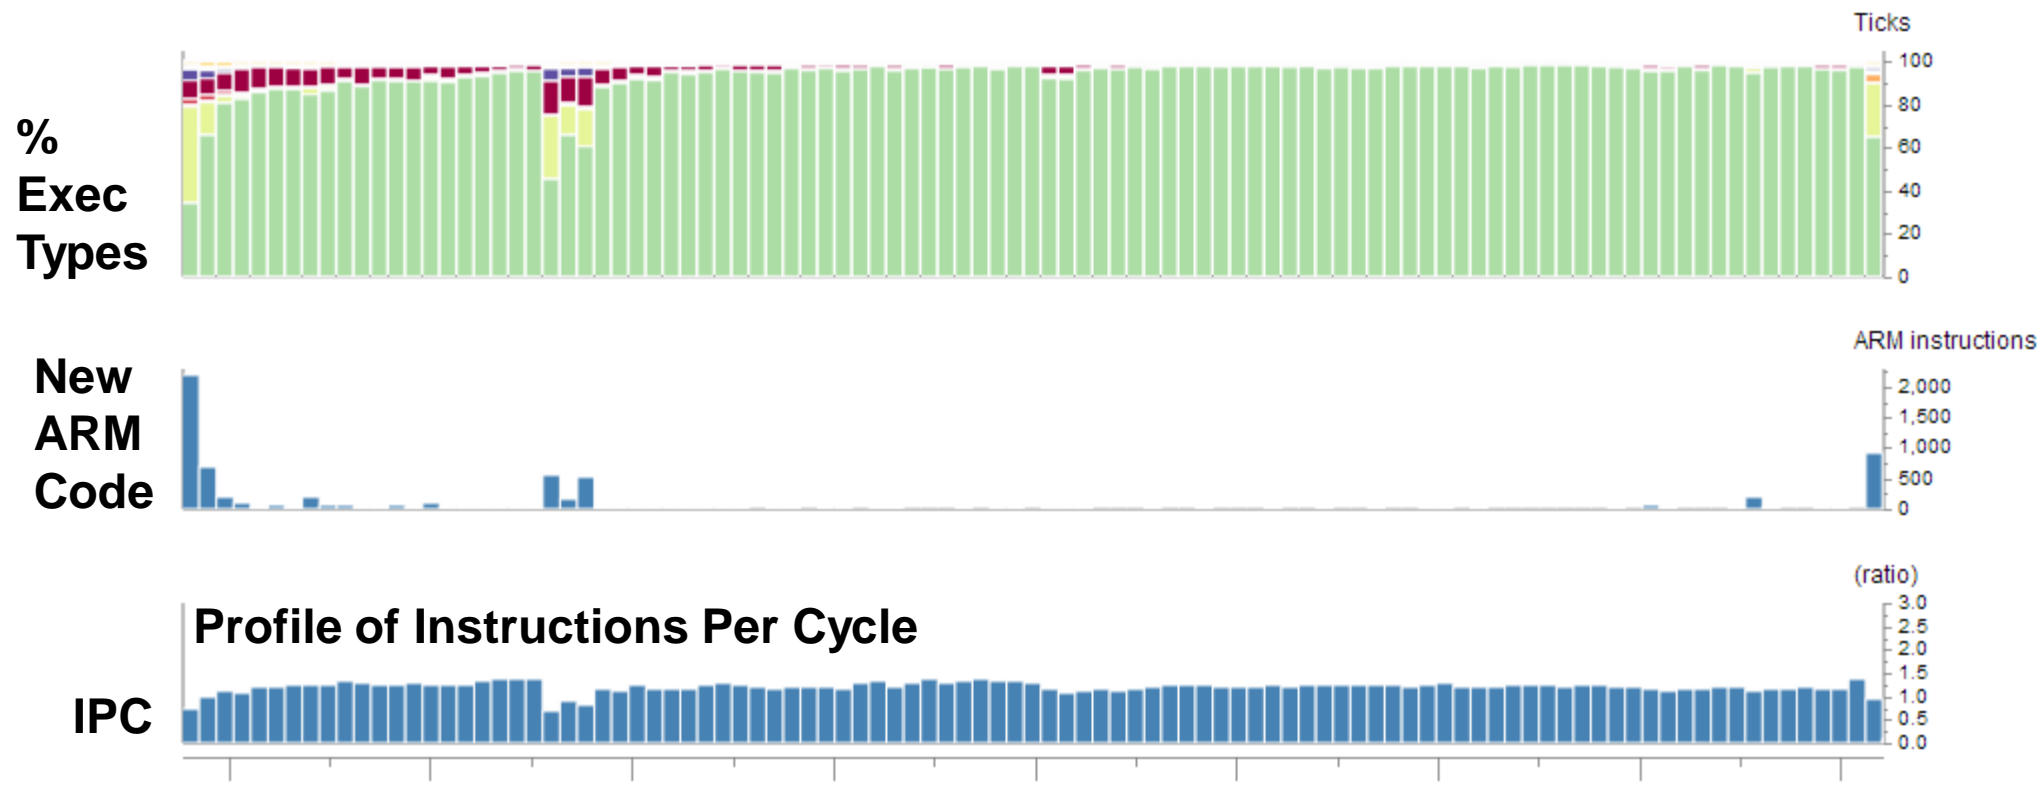
7-way Superscalar

Up to 2.5GHz

128K+64K L1\$

DYNAMIC CODE OPTIMIZATION: SPECINT EXAMPLE

← Full benchmark run →



Optimized uCode Execution

ARM Decoded Execution

Optimizer Execution

DENVER PERFORMANCE

