ESE532:
System-on-a-Chip Architecture

Day 27: April 26, 2017
Wrapup

Today

- What was course about
- Final
- Review w/ Simple Models
- Questions/Discussion
- Other Courses

2

Goal

- How to design/select/map to SoC to reduce Energy/Area/Delay.

3

Day 1

Outcomes

- Design, optimize, and program a modern System-on-a-Chip.
- Analyze, identify bottlenecks, design-space
- Decompose into parallel components
- Characterize and develop real-time solutions
- Implement both hardware and software solutions
- Formulate hardware/software tradeoffs, and perform hardware/software codesign

4

Day 1

Outcomes

- Understand the system on a chip from gates to application software, including:
  - on-chip memories and communication networks, I/O interfacing, RTL design of accelerators, processors, firmware and OS/ infrastructure software.
- Understand and estimate key design metrics and requirements including:
  - area, latency, throughput, energy, power, predictability, and reliability.

5

Day 2

Message for Day

- Identify the Bottleneck
  - May be in compute, I/O, memory, data movement
- Focus and reduce/remove bottleneck
  - More efficient use of resources
  - More resources
- Repeat

6

1

Abstract Approach

- Identify requirements, bottlenecks
- Decompose Parallel Opportunities
  - At extreme, how parallel could make it?
  - What forms of parallelism exist?
    - Thread-level, data parallel, instruction-level
- Design space of mapping
  - Choices of where to map, area-time tradeoffs
- Map, analyze, refine

7

SoC Designer
Hardware Building Blocks

- Computational blocks
  - Adders, multipliers, dividers, ALUs
- Registers
- Memory blocks
- Busses
- Multiplexers, Crossbars

- DMA Engines
- Processors
- I/O blocks
  - Ethernet, USB, PCI, DDR, Gigabit serial links

8

Final

9

Final

- Analysis
  - Bottleneck
  - Amdhal's Law Speedup
  - Computational requirements
  - Resource Bounds
  - Critical Path
  - Latency/throughput
- Model/estimate speedup, area, energy, yield

- From Code
- Forms of Parallelism
- Dataflow, SIMD, VLIW, hardware pipeline, threads
- Map/schedule task graph to (multiple) target substrates
- Memory assignment and movement
- Area-time points

10

Final

- Like midterm
  - Content
  - Style
  - Closed book, notes
  - Calculators allowed (encouraged)
- Monday, May 1 9am—11am
  - Towne 303 (here)
  - 2 hours (standard final exam period)

11

Review

12

Sequential Computation

- Computation requires a collection of operations
  - Arithmetic
  - Logical
 - Data storage/retrieval

$$T = \sum T_{op_i}$$

Computations in C

- Express computations in a programming language, such as C
- Can execute computation in many different ways
  - Sequential, parallel, spatial hardware

```
while(true) {
    for (i=0;i<16;i++)
        y[i]=a[i]*b[i];
    z[0]=y[0];
    for (i=1;i<16;i++)
        z[i]=z[i-1]+y[i];
}
```

Memory Characteristics

- Small memories
  - Fast
  - Low energy
  - Not dense
    - (high area per bit)
- Large memories
  - Slow
  - High energy
  - Dense (less area per bit)

Memory and Compute

- Computation involves both arith/ logical ops and memory ops

$$T = \sum T_{op_i}$$

$$T = \sum (op_i == mem) \times T_{mem} + \sum (op_i == alu) \times T_{alu}$$

$$T = N_{memop} \times T_{mem} + N_{alu} \times T_{alu}$$

Memory and Compute

- Computation involves both arith/ logical ops and memory ops
- Either can dominate
  - Be bottleneck

$$T = \sum T_{op_i}$$

$$T = N_{memop} \times T_{mem} + N_{alu} \times T_{alu}$$

Memory and Compute

- Timing
  - T_{mem}=10
  - T_{alu}=1
- N_{mpy}=
- N_{add}=
- N_{mem}=
- Total Time, T?

```
while(true) {
    for (i=0;i<16;i++)
        y[i]=a[i]*b[i];
    z[0]=y[0];
    for (i=1;i<16;i++)
        z[i]=z[i-1]+y[i];
}
```

$$T = N_{memop} \times T_{mem} + N_{alu} \times T_{alu}$$

(not count loop, indexing costs in this sequence)

3

Memory and Compute

- Where bottleneck?
  - $T_{mem}=10$
  - $T_{alu}=1$
- $N_{mpy}=16$
- $N_{add}=15$
- $N_{mem}=95$

```
while(true) {
    for (i=0;i<16;i++)
        y[i]=a[i]*b[i];
    z[0]=y[0];
    for (i=1;i<16;i++)
        z[i]=z[i-1]+y[i];
}
```

$$T = N_{memop} \times T_{mem} + N_{alu} \times T_{alu}$$

19

Data Reuse

- Reduce memory operations to large memory by storing in
  - Registers
  - Small memories

$$T = N_{memop} \times T_{mem} + N_{alu} \times T_{alu}$$

$$T = N_{smem} \times T_{smem} + N_{lmem} \times T_{lmem} + N_{alu} \times T_{alu}$$

$$N_{lmem} < N_{smem}$$

20

Memory and Compute

- a,b in large mem
- y,z in small mem
  - $T_{lem}=10$
  - $T_{sem}=1$
  - $T_{alu}=1$
- $N_{mpy}=16$
- $N_{add}=15$
- $N_{lmem}=?$
- $N_{smem}=?$

```
while(true) {
    for (i=0;i<16;i++)
        y[i]=a[i]*b[i];
    z[0]=y[0];
    for (i=1;i<16;i++)
        z[i]=z[i-1]+y[i];
}
```

$$T = N_{smem} \times T_{smem} + N_{lmem} \times T_{lmem} + N_{alu} \times T_{alu}$$

21

Memory and Compute

- a,b in large mem
- y,z in small mem
  - $T_{lem}=10$
  - $T_{sem}=1$
  - $T_{alu}=1$
- $N_{mpy}=16$
- $N_{add}=15$
- $N_{lmem}=32$
- $N_{smem}=63$

```
while(true) {
    for (i=0;i<16;i++)
        y[i]=a[i]*b[i];
    z[0]=y[0];
    for (i=1;i<16;i++)
        z[i]=z[i-1]+y[i];
}
```

- Performance impact?

$$T = N_{smem} \times T_{smem} + N_{lmem} \times T_{lmem} + N_{alu} \times T_{alu}$$

22

Task Parallel

- Can run unrelated tasks concurrently

$$T = \sum T_{op_i}$$

$$T = \min\left(\sum T_{op_i}(op_i \in Task1), \sum T_{op_i}(op_i \in Task2)\right)$$

Ideal: T(p) = T(1)/p

23

Data Parallel

- Can run same task on independent data in parallel

$$T = \sum T_{op_i}$$

Ideal: T(p) = T(1)/p

24

4

Data Parallel

- What's data parallel?

```
while(true) {
    for (i=0;i<16;i++)
        y[i]=a[i]*b[i];
    z[0]=y[0];
    for (i=1;i<16;i++)
        z[i]=z[i-1]+y[i];
}
```

Vector/SIMD

- Can perform same operation on a set of data items

$$T = \sum T_{op_i}$$

Ideal: T(VL) = T(1)/VL
(Vector Length)

Vector/SIMD

- Can perform same operation on a set of data items
 - Not everything vectorizable

$$T = \sum T_{op_i}$$

$$T = \left(\frac{1}{VL}\right)\sum T_{op_i}(vectorize(op_i)) + \sum T_{op_i}(\overline{vectorize(op_i)})$$

Vector/SIMD

- What's vectorizable?
- Speedup vector piece VL=4
  - (assume both memory and compute speedup)
- Overall speedup
- Amdahl's Law speedup for VL→infintiy?

```
while(true) {
    for (i=0;i<16;i++)
        y[i]=a[i]*b[i];
    z[0]=y[0];
    for (i=1;i<16;i++)
        z[i]=z[i-1]+y[i];
}
```

Dataflow Pipeline

- With no cycles in flowgraph
  - (no feedback)
  - (no loop carried dependencies)
- Producer and consumer can operate concurrently

$$T = \sum T_{op_i}$$

$$T = \min\left(\sum T_{op_i}(op_i \in \mathrm{Pr}od), \sum T_{op_i}(op_i \in Consume)\right)$$

Data Flow Pipeline

- Producer/consumer here?
- Time each
  - Assuming $T_{mem}=1$

```
while(true) {
    for (i=0;i<16;i++)
        y[i]=a[i]*b[i];
    z[0]=y[0];
    for (i=1;i<16;i++)
        z[i]=z[i-1]+y[i];
}
```

Spatial Pipeline

- Can build spatial pipeline of hardware operators to compute a dataflow graph

$$T = \sum T_{op_i}$$

- With no feedback: T=1
- With feedback loop length L: T=L

31

## Spatial Pipeline

- Pipeline taking one a[i], b[i] per cycle?

```
while(true)  {
    for (i=0;i<16;i++)
        y[i]=a[i]*b[i];
    z[0]=y[0];
    for (i=1;i<16;i++)
        z[i]=z[i-1]+y[i];
}
```

32

---

## Spatial Pipeline

- Full pipeline taking all 16 a's, b's in cycle?

- Latency?

```
while(true)  {
    for (i=0;i<16;i++)
        y[i]=a[i]*b[i];
    z[0]=y[0];
    for (i=1;i<16;i++)
        z[i]=z[i-1]+y[i];
}
```

33

Spatial Pipeline

34

VLIW

- Can control datapath to perform multiple, heterogeneous operations per cycle
  - Tune parallelism
- Op types: A, B, C $T = \sum T_{op_i}$
  - Ops N_A, N_B, N_C
  - Number of Hardware Units H_A, H_B, H_C
- $T_{RB} = max(N_A/H_A, N_B/H_B, N_C/H_C, \ldots) <= T_{VLIW}$
- $T_{CP} <= T_{VLIW}$

35

VLIW

- VLIW
  - 1 load/store (a,b)
    - Assume single cycle
  - 1 mpy
  - 1 add
- RB
- CP
- Schedule

```
while(true)  {
    for (i=0;i<16;i++)
        y[i]=a[i]*b[i];
    z[0]=y[0];
    for (i=1;i<16;i++)
        z[i]=z[i-1]+y[i];
}
```

36

6

VLIW

- VLIW
  - 4 load/store (a,b)
    - Assume single cycle
  - 2 mpy
  - 2 add
- RB
- CP
- Schedule

```
while(true) {
  for (i=0;i<16;i++)
    y[i]=a[i]*b[i];
  z[0]=y[0];
  for (i=1;i<16;i++)
    z[i]=z[i-1]+y[i];
}
```

37

Memory Bottleneck

- Memory can end up being bottleneck

$$T = \sum T_{op_i} = T_{comp} + T_{mem}$$

Ideal: $T(p) = T_{comp}(1)/p + T_{mem}$

38

Memory Bottleneck

- Pipeline to perform one multiply per cycle
- Memory supplies one data item from large memory (a,b) every 10 cycles
- Performance?
  - (number of cycles)

```
while(true) {
  for (i=0;i<16;i++)
    y[i]=a[i]*b[i];
  z[0]=y[0];
  for (i=1;i<16;i++)
    z[i]=z[i-1]+y[i];
}
```

39

Memory Bottleneck

- Memory can end up being bottleneck

$$T = \sum T_{op_i} = T_{comp} + T_{mem}$$

Ideal: $T(p) = T_{comp}(1)/p + T_{mem}/Membw$

40

Memory Bottleneck

- Invest in higher bandwidth
  - Wider memory
  - More memory banks
- Exploit data reuse
  - Smaller memories may have more bandwidth
  - Smaller memories are additional banks
    - More bandwidth

41

Communication

- Once parallel, communication may be bottleneck

$$T = \sum T_{op_i} = T_{comp} + T_{mem} + T_{comm}$$

- Ideal (like VLIW):
  $T <= max(T_{comp}/p, T_{mem}/membw, T_{comm}/netbw)$

42

7

Area

- Can model area of various solutions
  - Sum of component areas
  - (watch interconnect)

$$A = \sum A_i$$

- Cost proportional A/P_{chip}
- Without sparing
  - $P_{chip} = (P_{mm})^A$

$$P_{chip} = \prod P_i$$

- Sparing
  - Pmofn

Multi-Objective Optimization

- Many forms of parallelism
- Given fixed area (resources, energy)
  - Maximize performance
  - Select best architecture
- Given fixed performance goal
  - Find architecture that achieves
  - With minimum area (energy, cost)

Final

- Analysis
  - Bottleneck
  - Amdhal's Law Speedup
  - Computational requirements
  - Resource Bounds
  - Critical Path
  - Latency/throughput
- Model/estimate speedup, area, energy, yield

- From Code
- Forms of Parallelism
- Dataflow, SIMD, VLIW, hardware pipeline, threads
- Map/schedule task graph to (multiple) target substrates
- Memory assignment and movement
- Area-time points

Question/Discussion

Other Courses

Distinction

CIS240, 371, 501

- Best Effort Computing
  - Run as fast as you can
- Binary compatible
- ISA separation
- Shared memory parallelism
- Caching – automatic memory management
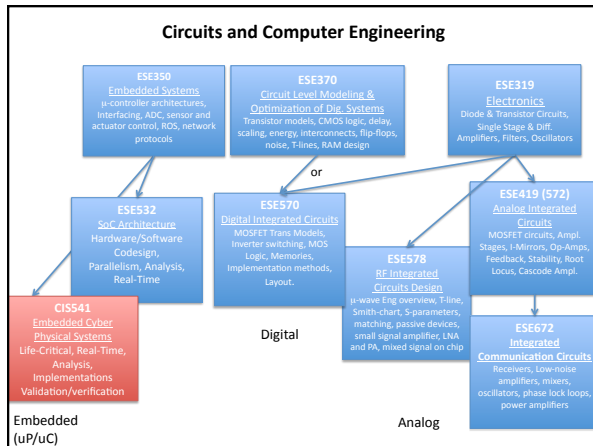- Superscalar
- Pipelined processor

ESE532

- Hardware-Software codesign
  - Willing to recompile, maybe rewrite code
  - Define/refine hardware
- Real-Time
  - Guarantee meet deadline
- Non shared-memory models

Circuits and Computer Engineering

Embedded (uP/uC) · Digital · Analog

Other Courses

- Security: CIS331, CIS551
- Networking: ESE407, CIS553
- GPGPU (graphics focus): CIS565

Message

- Any interesting, challenging computation will require both hardware and software
- SoC powerful implementation platform
  - Target pre-existing
  - Design customized for problem
  - Exploit heterogeneous parallelism
- Understand and systematically remove bottlenecks
  - Compute, memory, communicate, I/O

51

Admin

- Return Zed Boards → Detkin
- Final: Monday, May 1 9am—11am
  - Towne 303 (here)

52