ESE532: System-on-a-Chip Architecture

Day 14: March 13, 2017
Design Space Exploration

Today

- Midterm
- Design-Space Exploration

Message

- The universe of possible implementations (design space) is large
  - Many dimensions to explore
- Formulate carefully
- Approach systematically

Midterm

Mea Culpa

- Too long
  - Too time constrained

Mismatch?

- What I think I'm teaching
- What you think you are learning

- Midterm focused on performance modeling

Agree?

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

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

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.

Agree?

- How to design/select/map to SoC to reduce Energy/Area/Delay?
  - Haven't touched Area/Energy much, yet, so initial focus on Delay, Performance

Reduce Delay

- What can we do to reduce delay?
- How know when we succeed?
  - Compile and run
  - How know what's wrong and how to fix if we don't get a speedup?
- Need to be able to model, estimate, and reason about performance.

Model to Select Design

- Homework shows you that implementing a design is expensive
- **Goal:** have a model to drive you to select good implementations to try

Model to Understand

- If not get the performance you expect, what went wrong?
  - Need an expectation
  - Need to drill down into components
    - Model expectations wrong?
      - Which part?
      - Refine?
    - Something went wrong in implementations?
      - What?
      - How fix to match model?

Model to Understand

- If not get the performance you expect, what went wrong?
- Very Powerful
- Not match → something I need to understand
  - Bugs in implementation, understanding
- Does match → Tells me
  - where to look
  - what to try to fix
  - when to give up (on a particular approach)

Expect

- Kinds of analysis on the exam underlies all design exploration and experiments

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

Second Half

- How to design/select/map to SoC to reduce Energy/Area/Delay
- Spend some more time understanding
  - Area
    - Started as we look at accelerators
    - …worry about fit in fixed array
    - Cost for custom design
  - Energy

Design-Space Exploration

Generic

Design Space

- Have many choices for implementation
  - Alternatives to try
  - Parameters to tune
  - Mapping options
- Our freedom to impact implementation costs
  - Area, delay, energy

Design Space

- Ideally
  - Each choice orthogonal axis in high-dimensional space
  - Want to understand points in space
  - Find one that bests meets constraints and goals
- Practice
  - Seldom completely orthogonal
  - Requires cleverness to identify dimensions
  - Messy, cannot fully explore
  - But…can understand, priorities, guide

Preclass 1

- What choices (design-space axes) can we explore in mapping a task to an SoC?
- What showed up in homework so far?

From Homework?

- Types of parallelism
- Mapping to different fabrics / hardware
- How manage memory, move data
- Levels of parallelism
- Pipelining, unrolling, II

Design-Space Choices

- Type of parallelism
- How decompose / organize parallelism
- Area-time points (level exploited)
- What resources we provision for what parts of computation
- Where to map tasks
- How schedule/order computations
- How synchronize tasks
- How represent data
- Where place data; how manage and move
- What precision use in computations

Generalize Continuum

- Encourage to think about parameters (axes) that capture continuum to explore
- Start from an idea
  - Maybe can compute with 8b values
  - Maybe can put dist computation on FPGA fabric
  - Move data in 1KB chunks
- Identify general knob
  - Tune intermediate bits for computation
  - How much of computation go on FPGA fabric
  - What is optimal data transfer size?

Finding Optima

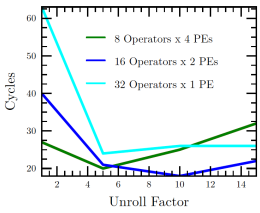

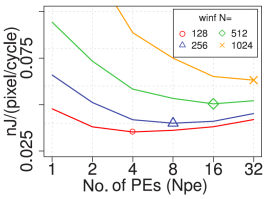

- Kapre, FPL 2009
- Kadric, TRETS 2016

Design Space Explore

- Think systematically about how might map the application
- Avoid overlooking options
- Understand tradeoffs

- Large design space
  →more opportunities to find good solutions
  Reduce bottlenecks

Elaborate Design Space

- Refine design space as you go
- Ideally identify up front
- Practice bottlenecks and challenges
  - will suggest new options / dimensions
    - If not initially expect memory bandwidth to be a bottleneck…
- Some options only make sense in particular sub-spaces
  - Bitwidth optimization not a big issue on the 64b processor
    - More interesting on vector, FPGA

Tools

- Sometimes tools will directly help you explore design space
  - What SDSoC/Vivado HLS support?
- Often they will not
  - What might you want that does not support?

Tools

- Sometimes tools will directly help you explore design space
  - Unrolling, pipelining, II
  - Some choices for data movement
  - Some loop transforms
  - Granularity to place on FPGA
- Often they will not
  - Need to reshape functions and loops
  - Data representations and sizes

Design-Space Exploration

Example FFT

Fourier Transform

- Identify spectral components
- Convert from Time-domain to Frequency-domain
  - E.g. tones from data samples
  - Central to audio coding – e.g. MP3 audio

$$Y[k] = \sum_{j=0}^{n-1} \left(X[j] e^{-2i\pi \frac{k}{n}} \right)$$

Fast-Fourier Transform (FFT)

- Efficient way to compute FT
- O(N*log(N)) computation

FFT

- Large space of FFTs
- Radix-2 FFT Butterfly

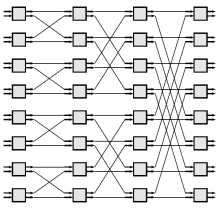

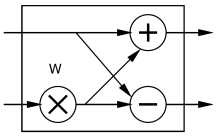

Basic FFT Butterfly

- Y0=X0+W(stage,butterfly)*X1
- Y1=X0-W(stage,butterfly)*X1
- Common sub expression, compute once: W(stage,butterfly)*X1

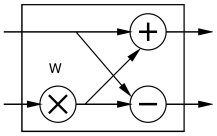

Preclass 2

- What parallelism options exist?
  - Single FFT
  - Sequence of FFTs

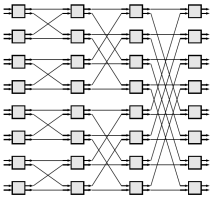

FFT Parallelism

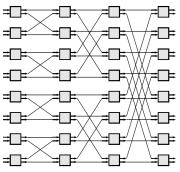

- Spatial
- Pipeline
- Streaming
- By column
  - Choose how many Butterflies to serialize on a PE
- By subgraph
- Pipeline subgraphs

Streaming FFT

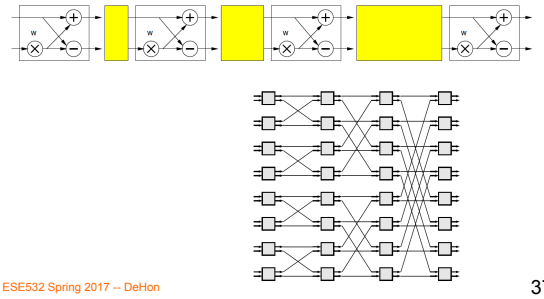

Common Subgraphs

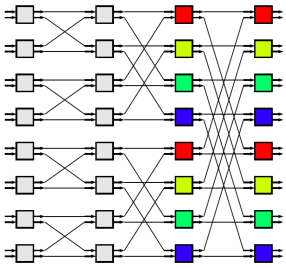

Preclass 3

- How large of a spatial FFT can implement with 220 multipliers?

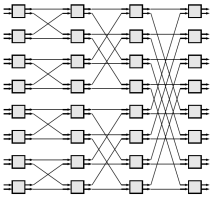

Bit Serial

- Could compute the add/multiply bit serially
  - One full adder per adder
  - W full adders per multiply
  - 50,000 LUTs
    - ~= 2500 bit-serial butterflies for W=16?
- Another dimension:
  - How much serialize word-wide operators

Accelerator Building Blocks

- What might we use as primitive, FFT-specific building blocks?

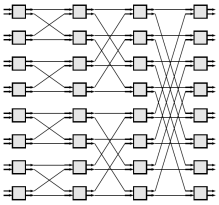

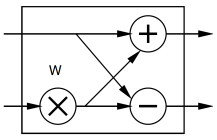

Processor Mapping

- How map butterfly operations to processors?
  - Implications for communications?

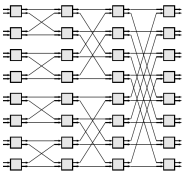

Preclass 4a

- How large local memory to communicate from stage to stage?

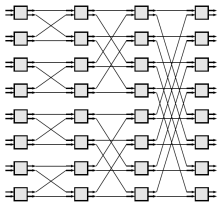

Preclass 4b

- How change evaluation order to reduce local storage memory?

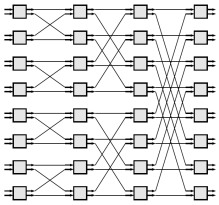

Communication

- How implement the data shuffle between processors or accelerators?
  - Memories / interconnect ?
  - How serial / parallel ?
  - Network?

Data Precision

- Input data from A/D likely 12b
- Output data, may only want 16b
- What should internal precision and representation be?

Number Representation

- Floating-Point
  - IEEE standard single (32b), double (64b)
    - With mantissa and exponent
    - …half, quad ….
- Fixed-Point
  - Select total bits and fraction
    - E.g. 16.8 (16 total bits, 8 of which are fraction)
      - Represent 1/256 to 256-1/256

Heterogeneous Precision

- May not be same in every stage
  - W factors less than 1
  - Non-fraction grows at most 1b per stage

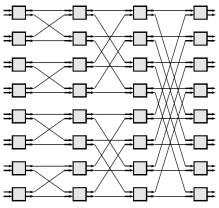

W/Twiddle factors

- Precompute and store in arrays
- Compute as needed
  - How? sin/cos hardware? CORDIC? Polynominal approximation?
- Specialize into computation
  - Many evaluate to 0, ±1, ±½, ….

FFT (partial) Design Space

- Parallelism
- Decompose
- Size/granularity of accelerator
  - Area-time
- Sequence/share
- Communicate
- Representation/precisions
- Twiddle

Big Ideas:

- Large design space for implementations
- Worth elaborating and formulating systematically
  - Make sure don't miss opportunities
- Think about continuum for design axes

Admin

- HW7 out → due Friday
  - Individual
- Working on getting Project ready