# 18-344: Computer Systems and the Hardware-Software Interface Fall 2023



#### **Course Description**

#### **Lecture 10: Design Space Exploration**

This course covers the design and implementation of computer systems from the perspective of the hardware software interface. The purpose of this course is for students to understand the relationship between the operating system, software, and computer architecture. Students that complete the course will have learned operating system fundamentals, computer architecture fundamentals, compilation to hardware abstractions, and how software actually executes from the perspective of the hardware software/boundary. The course will focus especially on understanding the relationships between software and hardware, and how those relationships influence the design of a computer system's software and hardware. The course will convey these topics through a series of practical, implementation-oriented lab assignments.

#### **Credit: Brandon Lucia**

#### Some details...

- HW3 was released last week, due Oct 3 (Thursday), 11:59 pm ET
- Lab 2 has been released today, due Oct 21 (Monday), 11:59 pm ET
- How is the class going in general? Anything you want to bring up?

## Today: Design Space Exploration

- Defining the design space of a hardware or software system
- Pareto Frontiers and optimizing within a design space
- Applied Performance Evaluation
  - Finding the best performing design under constraints

#### Defining a design space

- A design space is a set of possible incarnations of a system
- A design space is defined over a set of parameters
- A point in the design space is a concrete system with a concrete value for each of the design space's parameters
- Design spaces exist to allow systematic exploration of a collection of possible designs, like architectures.

#### Example: Branch Predictor Design Space



What are the parameters / dimensions of this branch predictor's design?

#### Example: Branch Predictor Design Space



**GHT** size BHT # entries GHT/PC hash func **BHT** entry size **BranchID** hash **BTB** # entries **BTB** assoc

These parameters are the dimensions of a design space vector

**GHT** size **BHT** # entries GHT/PC hash func BHT entry size **BranchID** hash **BTB** # entries BTB assoc

Sg Nb Hp Sb Hb Nt At

#### Example: Branch Predictor Design Space



#### Example: Branch Predictor Design Space



# Can find a good design by measuring points in the design space



# Can find a good design by measuring points in the design space



#### Is one of these better?









# Plotting the design space: Geometric view of figures of merit



FoM = "Feature of Merit", i.e. an attribute we care about.

Plotting many designs to study a tradeoff



#### Plotting many designs to study a tradeoff



#### Which points in this plot are optimal?



#### Pareto Optimality of Design Alternatives



**Pareto Optimality:** 

A design is optimal if no change leads to improvement in one dimension without a loss in at least one other dimension

Vilfredo Pareto



#### Pareto Optimality of Design Alternatives



# Design Consequence of Pareto Optimality

Never select designs other than at the frontier, at least without motivation outside of plot. Any design anywhere other than at the frontier can achieve the same or better performance at a lower cost w.r.t. the plotted dimensions.



#### Design Consequence of Pareto Optimality



## Worthwhile Options Are Along the Pareto Frontier



#### **Design Space Exploration**

Applied Performance Evaluation to find the best feasible system

- Define a system's important design parameters
- Define a system's figure(s) of merit
- Define a set of constraints on the feasibility of a binding of design parameters
- Choose a feasible parameter setting and measure its merit
- Iterate until satisfied:
  - If this system is better than the last one, keep it. If worse, discard it.
  - Choose a parameter and change it

#### **Design Space Exploration**

- Applied Performance Evaluation to find the best feasible system
  - Define a system's important design parameters
  - Define a system's figure(s) of merit
    - Define a set of constraints on the feasibility of a binding of design parameters
  - Choose a feasible parameter setting and measure its merit
  - Iterate until satisfied:
    - If this system is better than the last one, keep it. If worse, discard it.
    - Choose a parameter and change it

## Constraining your design space



## **Physical design constraints** Max BP power = 4mW Max BTB associativity = 2 Max memory (BTB+BHT) = 20kB

Designs candidates are often described as needing to "Make PPA":

- Power
- Performance
- Area

#### **Design Space Exploration**

- Applied Performance Evaluation to find the best feasible system
  - Define a system's important design parameters
  - Define a system's figure(s) of merit
    - Define a set of constraints on the feasibility of a binding of design parameters
    - Choose a feasible parameter setting and measure its merit
  - Iterate until satisfied:
    - If this system is better than the last one, keep it. If worse, discard it.
    - Choose a parameter and change it
    - Random restarts to search different sub-spaces



Change dimensions when way off frontier

Systematically fill out your design space



# Example of Design Space Optimization The Q100 Database Acceleration Architecture

#### Q100: The Architecture and Design of a Database Processing Unit

Lisa Wu Andrea Lottarini Timothy K. Paine Martha A. Kim Kenneth A. Ross Columbia University, New York, NY



#### **Cutting edge database query hardware accelerator**

- "GPU for SQL & Database operations"
- Architecture built up of a collection of special computing tiles in hardware
- Each tile runs a particular kind of database operation
- Tiles connected by configurable wires that can be set up to make circuits to do a database query
- (Includes one of the best design space explorations I've encountered in a research paper)

		Area		Power		<b>Critical Path</b>	D	esign Widtl	h (bits)	
	Tile	$mm^2$	% Xeon <sup>a</sup>	mW	% Xeon	ns	Record	Column	Comparator	<b>Other Constraint</b>
	Aggregator	0.029	0.07%	7.1	0.14%	1.95		256	256	
	ALU	0.091	0.21%	12.0	0.24%	0.29		64	64	
	BoolGen	0.003	0.01%	0.2	< 0.01%	0.41		256	256	
Functional	ColFilter	0.001	< 0.01%	0.1	< 0.01%	0.23		256		
	Joiner	0.016	0.04%	2.6	0.05%	0.51	1024	256	64	
	Partitioner	0.942	2.20%	28.8	0.58%	***3.17	1024	256	64	
	Sorter	0.188	0.44%	39.4	0.79%	2.48	1024	256	64	1024 entries at a time
	Append	0.011	0.03%	5.4	0.11%	0.37	1024	256		
Auviliany	ColSelect	0.049	0.11%	8.0	0.16%	0.35	1024	256		
Auxiliary	Concat	0.003	0.01%	1.2	0.02%	0.28		256		
	Stitch	0.011	0.03%	5.4	0.11%	0.37		256		

		A	rea	Po	wer	<b>Critical Path</b>	]	Design Widtl	h (bits)	
	Tile	$mm^2$	% Xeon <sup>a</sup>	milli ab 1	tile	ns	Record	Column	Comparator	<b>Other Constraint</b>
	Aggregator	0.000	number	ofeach	o a high-	1.95		256	256	
	ALU	Choose of	i -inal WO	rk they u		0.29		64	64	
	BoolGen	In the O	riginal to	n decide	now	0.41		256	256	
Functional	ColFilter	level sir	nulation	o more		0.23		256		
	Joiner	many ti	iles yield r	fit and	use that	0.51	1024	256	64	
	Partitioner	many	mance bei	nefil and	any of	***3.17	1024	256	64	
	Sorter	perform	er to bour	nd how "		2.48	1024	256	64	1024 entries at a time
	Append	d numb	tile they c	onsider	0.11%	0.37	1024	256		
Anviliany	ColSelect	0 each		8.0	0.16%	0.35	1024	256		
Auxillary	Concat	0.003	0.01%	1.2	0.02%	0.28		256		
	Stitch	0.011	0.03%	5.4	0.11%	0.37		256		

	Tile	<b>Area</b> mm <sup>2</sup> % X	eon <sup>a</sup>	Po	wer	Critic	Tile	Maximum Useful Count	"Tiny" Tile	Tile Counts Explored	her Constraint
Functional	Aggregator ALU BoolGen ColFilter Joiner Partitioner Sorter	<b>Choose a nu</b> In the origin level simula many tiles performan	mber o al work ation to yield no ice beno bound	of each they d decide o more efit and how m	o a high- how use that hany of		Aggregator ALU BoolGen ColFilter Joiner Partitioner Sorter	4 5 6 4 5 6	X X X X	4 1 5 6 4 1 5 1 6	entries at a time
Auxiliary	Append ColSelect Concat Stitch	0 each tile t 0.003 0.0 0.011 0.0	they con 1 % 01% 03%	8.0 1.2 5.4	0.11% 0.16% 0.02% 0.11%		Append ColSelect Concat Stitch	8 7 2 3	X X X X	8 7 2 3	

			Area		ower	<b>Critical Path</b>	D	esign Widtl		
(	Tile	$mm^2$	% Xeon <sup>a</sup>	mW	% Xeon	ns	Record	Column	Comparator	<b>Other Constraint</b>
	Aggregator	0.029	0.07%	7.1	0.14%	1.95		256	256	
Functional	ALU	0.091	0.21%	12.0	0.24%	0.20	lino	64	64	
	BoolGen	0.003	0.01%	0.2	< 0.01 %	asonable	baseline	256	256	
	ColFilter	0.001	< 0.01%	0	a area W.r	.t. a reasona	red to a	256		
	Joiner	0.016	0.04%	Estimat	e area i mal WO	rk, they compare	the Intel	256	64	
	Partitioner	0.942	2.20%	In the C	original ve	rver processor,	auded all	of 256	64	
	Sorter	0.188	0.44%	conter	nporary se	comparison in	+0	56	64	1024 entries at a time
	Append	0.011	0.03%	E5620	Xeon. The	wires, buffers, e	1024	256		
Annilian	ColSelect	0.049	0.11%	the co	onnecting	0.35	1024	256		
Auxinary	Concat	0.003	0.01%		0.02%	0.28		256		
	Stitch	0.011	0.03%	5.4	0.11%	0.37		256		

		Area		Power		<b>Critical Path</b>	Design Width (bits)			
2	Tile	$mm^2$	% Xeon <sup>a</sup>	mW	% Xeon	ns	Record	Column	Co pr	<b>Other Constraint</b>
Functional	Aggregator ALU BoolGen ColFilter Joiner Partitioner	0.029 0.091 0.003 0.001 0.016 0.942	0.07% 0.21% 0.01% <0.01% 0.04% 2.20%	7.1 12.0 0.2 0.1 2.6 28.8	0.14% 0.24% <0.01% 0.05% 0.58%	1.95 Want to mini In the origina of high-power 2 and all	al work, li er (10s of owed ark	ver mited the nu mW) units t pitrary count have <10mV	umber 0 0, 1, of "tiny" N.	
	Sorter	0.188	0.44%	39.4	0.79%	or 2, and	units that		64	1024 entries at a time
Auxiliary	Append ColSelect Concat Stitch	0.011 0.049 0.003 0.011	0.03% 0.11% 0.01% 0.03%	5.4 8.0 1.2 5.4	$0.11\% \\ 0.16\% \\ 0.02\% \\ 0.11\%$	0.37 0.35 0.28 0.37	1024 1024	256 256 256 256		

		Area		Power		<b>Critical Path</b>	Critical Path Design Width (bits)			n (bits)	
	Tile	$mm^2$	% Xeon <sup>a</sup>	mW	% Xeon	ns	R	ecord	Column	Comparator	<b>Other Constraint</b>
	Aggregator	0.029	0.07%	7.1	0.14%	1.95			256	256	
	ALU	0.091	0.21%	10		0.29			64	64	
	BoolGen	0.003	the tile lat	ency	ubat	0.41			256	256	
Functional	Colf	v limite	d by the	n mear	is that	0.23			256		
	J Frequein	volv nipe	elined desig	ac the r	naximum	0.51		1024	256	64	
	Pa Aggressi	very par	delay define	estine	as the	***3.17		1024	256	64	
	So the criti	cal part	(which is th	le same		2.48		1024	256	64	1024 entries at a time
	Ap switchir	ng delay	e design).	Grad	for Q100)	0.37		1024	256		
Auxiliary	Col. frequer	ncy of the	ways define	s freq.	0.10%	0.35		1024	256		
	Con (partit	ioner all	Waye	1.2	0.02%	0.28			256		
	Stite (Faire	0.011	0.03%	5.4	0.11%	0.37			256		

		Area		Power		<b>Critical Pat</b>	itical Path D		esign Widtl	n (bits)	
2	Tile	$mm^2$	% Xeon <sup>a</sup>	mW	% Xeon	n	s ]	Record	Column	Comparator	<b>Other Constraint</b>
	Aggregator	0.029	0.07%	7.1	0.14%	1.9	5		256	256	
	ALU	0.091	0.21%	12.0	0.240	0.2	9		64	64	
Functional	BoolGen	0.003	0.0107	DB he	nchmark	0.4	1		256	256	
	ColFilter	0.0	on standard			0.2	3		256		
	J simulate	e design	masurem	ents for	(TPC-H)	0.5	1	1024	256	64	
	Pa	run time	measure	ichmark		***3.1	7	1024	256	64	
	So. Collect	tion-Pro	cessing being	system	without	2.4	8	1024	256	64	1024 entries at a time
	Ap Transac	tresses	a database	-hing fro	om meriio	0.3	7	1024	256		
Auxiliary	Col. which a	bottlene	cked by let	8.0	0.16%	0.3	5	1024	256		
	Con being	0000	0.01%	1.2	0.02%	0.2	8		256		
	Stitch	0.011	0.03%	5.4	0.11%	0.3	7		256		

#### Q100 Pareto Frontier



Pareto plot from a research paper on the Q100 Database accelerator by Wu et al, ASPLOS 2014

- How did they select magenta points?
- What other points might they have selected?
- What is the value in seeing all these points?

- "Pareto Design" as used in the paper means the design that maximizes (runtime) performance per watt.
- Although there were designs with nominally better runtime, the goal of the paper was to select three options for further study. The two options with a nominally better runtime were only negligibly better but at a much higher cost in terms of energy, rendering them less interesting to the authors.

#### Results of Design Space Exploration

			Area		Power						
	Tiles	NoC	SBs	Total	Total	Tiles	NoC	SBs	Total	Total	
	$mm^2$	$mm^2$	$mm^2$	$mm^2$	% Xeon	W	W	W	W	% Xeon	
LowPower	1.890	0.567	0.520	2.978	7.0%	0.238	0.071	0.400	0.710	14.2%	
Pareto	3.107	0.932	0.780	4.819	11.3%	0.303	0.091	0.600	0.994	19.9%	
HighPerf	5.080	1.524	0.780	7.384	17.3%	0.541	0.162	0.600	1.303	26.1%	



#### Heat Plots Can Be Used to Explore 2D Space



**Figure 10.** Even with a LowPower design, the communication bandwidth for heat map, Pareto design maximum intramost connections exceed the provisioned connection bandwidth exhibit almost 6.3 GB/s NoC bandwidth, marked as X's identical behavior as HighPerf design.

Here heat plots are used to show the communication bandwidth needed between tiles and which design elements exceed a reference threshold.

# Q100 Takeaways / What did we just learn

- Practical application of design space exploration
- Defined design space based on tiles and connections between tiles
- Defined constraints and optimization goals based on power, area, frequency
- Runs experiments to produce Pareto Frontier with performance and power as main design dimension
- Final designs come from Pareto Frontier fast, balanced, low-power
- Compare design to characteristics of known baseline (Xeon)

## What to think about next?

- Miscellaneous (micro)architectural tricks & optimizations (future)
  - Super-scalar Out-of-Order
  - VLIW
  - Vector processors / SIMD
  - SIMT/GPU