The Hardware Software Interface

CSE 351 Summer 2024

Instructor: Ellis Haker

Teaching Assistants:

Naama Amiel Micah Chang Shananda Dokka Nikolas McNamee Jiawei Huang



Lecture Outline

- Introduction
- <u>Syllabus</u>/Logistics
- Binary and Numerical Representation

Introductions: Course Staff

Instructor: Ellis

 \odot BS/MS student

○ Research focus in Operating Systems

 \circ Long-time TA, first-time instructor!

•TAs:

Naama, Micah, Shananda, Nikolas, and Jiawei
Available in section, office hours, and on Ed

• More than anything, we want you to feel

 \odot Welcome in this space

 \odot Able to learn and succeed in this course

○ Comfortable reaching out if you need help or want change

Introductions: You!

- •~70 students
- •CSE majors, ECE majors, and more
 - \odot Majority of the class are non-majors
 - \odot Most of you will find almost everything in the course new
- Get to know each other! Help each other out!
 - $\ensuremath{\circ}$ Research suggests that group work supports learning
 - \odot Working well with others is a valuable life skill
 - \odot Diversity of perspectives expands your horizons
 - Take advantage of group work, where permissible, to learn, not just get a grade

Introductions: CSE 351



- You'll learn the key abstractions "under the hood"
 - How does your source code become something the computer understands?
 - What happens as your computer is executing one or more programs?

Some History

- Hardware started out quite primitive
 - Programmed with very basic instructions
 - Very tedious!
- Software was also very basic
 - Programs reflected the actual hardware they ran on
 - Programmer had to specify each step manually



Programmers working on the ENIAC, circa 1946.

Some History pt 2

- As time went on, programming became more abstract
 - Assembly language: basic set of instructions
 - Early high-level languages: C, FORTRAN, etc.
 - Data types, arrays, loops, etc.
 - Still closer to the hardware than modern languages
 - Now: Java, Python, etc.
 - Lots of convenient features!
 - Don't have to know much about the

hardware to program



Brian Kernighan and Dennis Ritchie, creators of the original C standard.

Layers of Computing

Software Applications (written in Java, Python, C, etc.)

Programming Languages & Libraries (e.g. Java Runtime Env, C Standard Lib)

OS/App Interface

Operating System (e.g. MacOS, Windows, Linux)

HW/SW Interface

Hardware (e.g. CPU, memory, disk, network, peripherals)

Course Perspective

- CSE 351 will make you a better programmer
 - Learn how software really works
 - Understand some of the abstractions that exist between hardware and software, why they exist, and how they build upon each other
- Why is this important?
 - Better debugging
 - Better basis for evaluation performance
- You might miss Java, but keep an open mind!
 - Take note of anything that piques your interest. There's lots of classes you can take to explore this space further

Syllabus/Logistics

Important Tools

• Course Website:

https://courses.cs.Washington.edu/courses/cse351/24su

 \odot Syllabus, schedule, lab specs, etc.

• Ed course

 \odot Discussion: announcements, ask and answer questions

 \odot Lessons: readings, lecture polls, homework

- Gradescope: lab submissions, take-home quizzes
- Canvas: surveys, grade book
- Panopto: lecture recordings



Find Ed Lessons in the top-right corner of Ed.

Quizzes

- Three throughout the quarter
- Take-home on Gradescope
- Will have ~1 week to complete

• More information to come when we get to the first quiz

Assignments

• Pre-lecture readings: Ed

○ Graded on completion, due 1pm before lecture

• Lecture polls: Ed

 \odot Conducted during class, graded on completion, due 1pm before the next lecture

• Homework: Ed

 \odot Unlimited attempts, due 1pm two lectures after it's assigned

• Labs: Course website/Gradescope (except for Lab0)

 \odot ~1-2 weeks per lab, due at 11:59pm

 Assignments need to be turned in <u>before</u> the time listed. If the clock reads 11:59pm when you turn in your lab, it's late!

Late Policies

• Labs can be turned in up to 2 days after the deadline

You get 5 free late days to use throughout the quarter
After your free late days are used up, we will deduct 10% per day

We cannot accept Ed assignments after the due date!

You only need to complete 80% of the readings and lecture polls for full credit
 Your lowest homework score will be dropped

If you have DRS accommodations that contradict these policies, I will contact you

 \odot Please email me if I don't reach out by the end of today

Extensions, Accommodations, Help

- •We're here to help! We want you to have the best experience possible
- Extenuating circumstances
 - \circ Students (and staff) face an extremely varied set of circumstances
 - We will try to be accommodating, but the earlier you reach out, the better
 - For formal accommodations go through Disability Resources for Students (DRS)
 - They can help you even for temporary disabilities (injury, etc.)
- Don't suffer in silence talk to a staff member!
 - Go to office hours
 - $\circ\,$ Ask on the Ed board
 - Request a <u>1-on-1 meeting</u>

Grade Breakdown

- Pre-lecture readings: 7%
- Homework: 20%
- Labs: 40%
 - $\circ\,$ Can be done individually or in pairs
- Quizzes: 30% (10% each)

• Participation: 3%

 $\circ\,$ Includes answering lecure polls, participation in section, office hours, and Ed $\,$

Administrivia/TODOs

- Explore the course website, and read the syllabus
- Make sure you can connect to the CSE Linux environment (attu/calgary)
- (Optionally) sign up for CSE 391: System and Software Tools
- Due **Friday** (6/21)
 - \circ Pre-course Survey (on Canvas), due 1pm
 - \circ HW0 (Ed) due 1pm
 - \circ RD1 & RD2 (Ed) due 1pm

• Due Monday (6/24)

- \circ RD3 (Ed) due 1pm
- \circ HW1 (Ed) due 1pm
- \circ Lab0 (Ed) due 11:59pm

• Reminder: no class on Wednesday due to holiday!

Binary

Base Definitions (Review) Base Definitions (Review)

• If we're in base *b*, that means we have *b* possible symbols to use

 $\circ\,$ In decimal (base 10), we have the digits 0-9 $\,$

- Each digit represents a power of b
 - \circ The rightmost digit always represents 1 (b^0), and it increases as we read left

Ex: compare the number written "351" in base 10 vs base 6

10 ² = 100	10 ¹ = 10	10 ⁰ = 1						
3	5	1						

In base 10, this numeral means we have 3 100s, 5 10s, and a 1.

$6^2 = 36$	6 ¹ = 6	6 ⁰ = 1						
3	5	1						

In base 6, this numeral means we have 3 36s, 5 6s, and a 1.

Base Conversions pt 1 (Review)

• To convert from base *b* to decimal, multiply each digit by its corresponding place value, then add them

Ex: find the decimal value of 246_8 (note, the subscript "8" denotes the base):

$8^2 = 64$	8 ¹ = 8	8 ⁰ = 1						
2	4	6						

So the total value is 2*64 + 4*8 +6*1 = <u>166</u>

Base Conversions pt 2 (Review)

• To convert a number *N* from decimal to base *b*:

- 1. Find the highest power *i* such that $b^i < N$
- 2. Divide *N* by *bⁱ* and round down to the nearest integer. The result is the *i*th digit of our final number
- 3. Take the remainder, and repeat step 2 with *i*-1 until you reach *i*=0

Ex: convert the number 123 to base 4:

$$4^{3} = 64, 123/64 = \underline{1}, \text{ remainder} = 59$$

 $4^{2} = 16, 59/16 = \underline{3}, \text{ remainder} = 11$
 $4^{1} = 4, 11/4 = \underline{2}, \text{ remainder} = 3$
 $4^{0} = 1, 3/1 = \underline{3}$
So the final answer is $\underline{1323_{4}}$
 $4^{0} = 1, 3/1 = \underline{3}$
 59
 $- \frac{1}{8}$
 $- \frac{1}{8}$
 $- \frac{1}{8}$
 $- \frac{1}{8}$

Review Question: Base Conversion



Binary

- Humans think about numbers in base 10, but computers "think" about numbers in base 2 (binary)
- A binary digit is called a bit
- A group of 4 bits is called a **nibble**
- A group of 8 bits is called a **byte**

Why Binary?

• Easy to store and read

 \odot For larger bases, noise would lead to inaccuracies

• Other bases are possible, but not yet viable

 \odot The Setun was a ternary (base 3) computer built in the USSR in 1958

○ UW research into DNA storage (base 4: A, C, G, T)

○ Quantum

Binary and Hex (Review)

- Binary is inconvenient for humans, so computer scientists often write numbers in **base 16** (hexadecimal) EX: 0,..., 9, A, B, ..., F, 10 in base 16 = 0,..., 9, 10, 11, 15, 16 in base 10
 - 16 digits: 0-9, A-F
 - Why hex?
 - Easy conversion, each hex digit = 4 bits!
 - 2 hex digits = 1 byte
- We use prefixes to denote common bases
 - \circ 0b = binary
 - \circ 0x = hex
 - No prefix = decimal Ο

Common Base Conversion (Review)

 hex -> binary: translate each digit according to chart, then drop leading 0s

• Ex: 0x2D = 0b0010 1101 = 0b101101

- binary -> hex: break into groups of 4 bits from right to left, add leading 0s if necessary, then translate
 Ex: 0b101101 = 0b0010 1101 = 0x2D
- Note: does not work for decimal conversions!
- You will probably have this table memorized by the end of the quarter :)

Base 10	Base 2	Base 16						
0	0000	0						
1	0001	1						
2	0010	2						
3	0011	3						
4	0100	4						
5	0101	5						
6	0110	6						
7	0111	7						
8	1000	8						
9	1001	9						
10	1010	А						
11	1011	В						
12	1100	С						
13	1101	D						
14	1110	E						
15	1111	F						

Review Questions: Binary, Hex, and Decimal

16·6

What is the *hex value* of 108?



Convert 0b100110110101101 to hex. = $0 \times 4 DAD$

Convert 0x3C9 to binary.

=01 0011 1100 1001

Numerical Encoding

• You can represent anything countable using numbers!

• But you need to agree on an encoding

• Computers store all data as a binary number

• Examples:

- \circ Decimal Integers: 0→0b0, 1→0b1, 2→0b10, etc.
- \circ English Letters: CSE→0x435345, yay→0x796179

○ Emojis: ♥ 0x0, ♀ 0x1, ♥ 0x2, ♥ 0x3, ♥ 0x4, ♥ 0x5

So What's it Mean?

A sequence of bits can have many meanings!

- Consider the hex sequence 0x4E6F21
 - Possible interpretations include:
 - The decimal number 5120257
 - The real number 7.203034*10⁻³⁹
 - The characters "No!"
 - The background color of this slide (sort of an olive green?)
- It's up to the program/programmer to decide how to interpret the sequence of bits

Binary Encoding Example - Colors

- RGB Red, Green Blue
 - Additive model, represent amount of each color of light
 - 1 byte (8 bits) for each color

Ex: **Blue = 0x0000FF**, **White = 0xFFFFF**

Dark Purple = 0x7030a0



Binary Encoding Example - Text

American Standard Information

Exchange (ASCII)

• Unicode (international)

<u>Dec</u>	H	(Oct	Cha	r	Dec	Hx	Oct	Html	Chr	Dec	Hx	Oct	Html	Chr	Dec	Hx	Oct	Html C	<u>hr</u>
0	0	000	NUL	(null)	32	20	040	&# 32;	Space	64	40	100	«#64;	0	96	60	140	`	1
1	1	001	SOH	(start of heading)	33	21	041	∉#33;	1.00	65	41	101	A	A	97	61	141	⊛#97;	a
2	2	002	STX	(start of text)	34	22	042	 <i>‱#</i> 34;	"	66	42	102	B	В	98	62	142	 ‰#98;	ь
3	3	003	ETX	(end of text)	35	23	043	∝# 35;	#	67	43	103	C	C	99	63	143	 ∉#99;	С
4	4	004	EOT	(end of transmission)	36	24	044	∉#36;	ę.	68	44	104	∝#68;	D	100	64	144	€#100;	: d
5	5	005	ENQ	(enquiry)	37	25	045	∝# 37;	*	69	45	105	 ∉#69;	E	101	65	145	e	e
6	6	006	ACK	(acknowledge)	38	26	046	 ∉38;	6	70	46	106	∝#70;	F	102	66	146	f	f
- 7	7	007	BEL	(bell)	39	27	047	 ∉39;	1	71	47	107	G	G	103	67	147	≪#103;	g
8	8	010	BS	(backspace)	40	28	050	∝#40;	(72	48	110	H	н	104	68	150	 %#104;	: h
9	9	011	TAB	(horizontal tab)	41	29	051	‰#4l;)	73	49	111	¢#73;	I	105	69	151	∝#105;	; i
10	A	012	LF	(NL line feed, new line)	42	2A	052	€#42;	*	-74	4A	112	¢#74;	J	106	6A	152	∝#106;	; j
11	В	013	VT	(vertical tab)	43	2B	053	+	+	75	4B	113	 ∉#75;	K	107	6B	153	 ∉#107;	; k
12	С	014	FF	(NP form feed, new page)	44	2C	054	∝#44;	1.	76	4C	114	L	L	108	6C	154	∝#108;	: 1
13	D	015	CR	(carriage return)	45	2D	055	∝#45;	- 1	77	4D	115	M	М	109	6D	155	∝#109;	: m
14	Ε	016	S0 -	(shift out)	46	2E	056	.	A U N	78	4E	116	∉ #78;	N	110	6E	156	€#110;	n
15	F	017	SI	(shift in)	47	2F	057	¢#47;		79	4F	117	 ∉#79;	0	111	6F	157	o	: O
16	10	020	DLE	(data link escape)	48	30	060	∝#48;	0	80	50	120	 ∉#80;	P	112	70	160	⊛#112;	p
17	11	021	DC1	(device control 1)	49	31	061	∝#49;	1	81	51	121	 ∉#81;	Q	113	71	161	¢#113;	a d
18	12	022	DC2	(device control 2)	50	32	062	 ∉\$0;	2	82	52	122	 ∉#82;	R	114	72	162	r	r
19	13	023	DC3	(device control 3)	51	33	063	3	3	83	53	123	 ∉#83;	S	115	73	163	s	: 3
20	14	024	DC4	(device control 4)	52	34	064	& # 52;	4	84	54	124	 ∉#84;	Т	116	74	164	t	t
21	15	025	NAK	(negative acknowledge)	53	35	065	∝# 53;	5	85	55	125	 ∉#85;	U	117	75	165	u	. u
22	16	026	SYN	(synchronous idle)	54	36	066	∝#54;	6	86	56	126	V	V	118	76	166	∝#118;	v
23	17	027	ETB	(end of trans. block)	55	37	067	 ∉#55;	7	87	57	127	 ∉#87;	W	119	77	167	w	. W
24	18	030	CAN	(cancel)	56	38	070	∝#56;	8	88	58	130	X	Х	120	78	170	∝#120;	: x
25	19	031	EM	(end of medium)	57	39	071	∝#57;	9	89	59	131	 ∉#89;	Y	121	79	171	y	Y Y
26	1A	032	SUB	(substitute)	58	ЗA	072	 ∉58;	÷	90	5A	132	 ∉#90;	Z	122	7A	172	∝#122;	z
27	1B	033	ESC	(escape)	59	ЗB	073	 ∉\$59;	2	91	5B	133	[[123	7B	173	∝#123;	: {
28	1C	034	FS	(file separator)	60	ЗC	074	∝#60;	<	92	5C	134	 ∉#92;	- N -	124	7C	174	∝#124;	: 1
29	1D	035	GS	(group separator)	61	ЗD	075	l;	=	93	5D	135	 ∉#93;	1	125	7D	175	∝#125;	; }
30	lE	036	RS	(record separator)	62	ЗE	076	 ∉62;	>	94	5E	136	 ∉#94;	^	126	7E	176	∝#126;	~
31	lF	037	US	(unit separator)	63	ЗF	077	 ∉#63;	2	95	5F	137	_	_	127	7F	177		DEL

Source: www.LookupTables.com

Binary Encoding Example - Files and Programs

- At the lowest level, all digital data is stored as bits!
- Layers of abstraction keep everything comprehensible to humans
 - Data/files are groups of bits interpreted by a program
 - Program is *also* a sequence of bits interpreted by the CPU

- You can see the hex representation of your files in a text editor try it!
 - Vim: %! xxd
 - Emacs: M-x hexl-mode

Summary

- All computer data is stored in **binary**
 - Humans think in **decimal**, have to convert between bases
 - Hex as a more human-readable base that's easy to convert
- Binary can represent anything!
 - Program needs to know how to interpret the bits

Some fun topics we will touch on (vote on Ed!)

- A. What is a GFLOP and why is it used in computer benchmarks?
- B. How and why does running many programs for a long time eat into your memory (RAM)?
- C. What is stack overflow and how does it happen?
- D. Why does your computer slow down when you run out of *disk* space?
- E. What was the flaw behind the original Internet worm, the Heartbleed bug, and the Cloudbleed bug?
- F. What is the meaning behind the different CPU specifications? (*e.g.*, # of cores, # and size of cache)