

§8.1–8.2: Low Level Storage

•*devo*

28 Oct 2005
CMPT14x
Dr. Sean Ho
Trinity Western University

Reminders:

- ***journals** in folder*
- ***Homework** due*

Welcome to CMPT 145!

- We've already covered one whole book:
 - You now know the basics of M2 programming!
- CMPT145 (book 2, ch8–12) dives deeper into advanced techniques:
 - Data storage and how I/O works
 - Cool datatypes: sets and records
 - Scope/visibility, exceptions
 - Software development techniques
 - Pointers and dynamic ADTs

What's on for today (8.1–8.2)

- Number bases:
 - Binary, hexadecimal, octal
- Units of measure of memory:
 - Bits, nibbles, bytes, words, pages
- Units of measure for hard disks:
 - Geometry, cylinders/heads/sectors
- SI units vs binary units

Ch8: Data storage and I/O

- As programmers, you're already expert **users** of various datatypes and file I/O
- Now we peek **under the hood** to see what the M2 compiler and the OS are really doing to implement these
- Every **VAR**iable we declare takes up space in **memory** (RAM):
 - How much **space** does each variable need?
 - How is our data **stored**?

Binary numbers



- At the lowest level, all computer data are stored using logical **bits**: each bit can be either 0 or 1
 - **High voltage** (1) vs. **low** voltage (0)
 - Most memory chips use a big bank of tiny **capacitors**: has charge (1) vs. no charge (0)
- We use groups of bits to **represent** data (numbers, characters, strings, etc.):
 - e.g., this pattern of eight **bits**: 0 1 0 0 0 0 1 1
 - ◆ Could represent the decimal **number** 35
 - ◆ Or it might represent the **character** “#”
 - ◆ Or something else – depends on how we **interpret** it

Number bases

- God gave us 10 fingers; so we often count in **base 10**:

- “5927” interpreted as a **decimal** number:

- ◆ 5 units of ($10^3 = 1000$)
- ◆ 9 units of ($10^2 = 100$)
- ◆ 2 units of ($10^1 = 10$)
- ◆ 7 units of ($10^0 = 1$)

- Counting in **binary** is similar:

- “0110” interpreted as a binary number:

- ◆ 0 unit of ($2^3 = 8$)
- ◆ 1 unit of ($2^2 = 4$)
- ◆ 1 unit of ($2^1 = 2$)
- ◆ 0 unit of ($2^0 = 1$)



Hexadecimal, octal

- **Hexadecimal** is base 16: we use 'A'..'F' to represent the “digits” ten, eleven, twelve, etc.
 - “BEEF” as a hexadecimal number:
 - ◆ B (11) units of ($16^3 = 4096$) \Rightarrow 45056
 - ◆ E (14) units of ($16^2 = 256$) \Rightarrow 3584
 - ◆ E (14) units of ($16^1 = 16$) \Rightarrow 224
 - ◆ F (15) units of ($16^0 = 1$) \Rightarrow 15
 - ◆ Total: BEEF (hex) \Rightarrow 48879 (dec)
- There's also **octal**, base 8:
 - only the digits 0..7 are used

Using bases in Modula-2

- Modula-2 has special **notation** for expressing numbers in hexadecimal and octal:
 - **Hexadecimal**: suffix “H”
 - **Octal**: suffix “B”
 - Can't start with a letter
- Can also use octal notation to specify a **character** by suffixing “C”: similar to CHR() but in octal

CONST

hexNum = 0BEEFH ;	(* 48879 *)
octNum = 115B ;	(* $1(8^2) + 1(8^1) + 5(8^0) = 77$ *)
charM = 115C ;	(* character #77 = 'M' *)

Bits, bytes, nibbles, words

- One hexadecimal digit can be represented by **four bits**: one **nibble**
- Two nibbles (**eight bits**) is called a **byte**
 - One byte can be used to store one **CHAR**
- A group of bytes can be used to represent one datum: this is called a **word**
 - Pentium CPUs generally use 4-byte words (**32 bits**)
 - Newer CPUs can use 8-byte words (**64 bits**)
 - Word is the smallest **unit of data** the machine can store or retrieve

Accessing memory



- A computer's **main memory** (generally, RAM) stores everything it needs to do its current tasks
- A location within memory is uniquely identified by its **address**
 - Most modern CPUs use 32-bit words to **store** memory addresses
 - This means there is a maximum of 2^{32} unique memory addresses (the **address space**)
 - If each location stores one byte of data, then there is 2^{32} bytes = 4GB of **addressable memory**

Units of measure

■ SI abbreviations:

- K = kilo = 1,000
- M = mega = 1,000,000
- G = giga = 1,000,000,000


■ When working with binary data:

- KB = kilobyte = 1,024 bytes = 2^{10} bytes
- MB = megabyte = 1,024,576 = 2^{20} bytes
- GB = gigabyte = 1,073,741,824 = 2^{30} bytes
- But hard drive manufacturers use SI abbrevs

Units of measure, cont.

- Kilobytes vs. kilobits:
 - KB = kilobyte = 1,024 bytes = $1024 * 8$ bits
 - Kb = kilobit = 1,024 bits
 - RAM chip manufacturers often use kilobits
- Also, in SI abbreviations,
 - M = mega = 10^6 : e.g., megawatt = 10^6 watt
 - m = milli = 10^{-3} : e.g., milliwatt = 10^{-3} watt
- But not everyone is consistent, so be careful

Storage



- A **page** of memory is generally 256 bytes
- A **sector** is a unit of disk storage, also commonly 256 bytes (but sometimes 512 bytes)
- A **block** of disk storage is usually 512 bytes
- Hard disks are made up of **platters**, accessed by magnetic **heads** on movable arms
- The platters have concentric tracks that (across all heads) make up **cylinders**
- Hard drive geometry is often expressed in **C/H/S**:
cylinders / # heads / # sectors per track

Summary of today (8.1–8.2)

■ Number bases:

- Binary
- Hexadecimal (0BEEFH)
- Octal (115B)

◆ Defining characters with octal: 115C

■ Units of measure of memory:

- Bits, nibbles, bytes, words, pages

■ Units of measure for hard disks:

- C/H/S geometry

■ SI units vs binary units, KB vs. Kb, etc.

TODO items

- No lab next week!
- Homework due next Wed: 8.13 #44
- Quiz ch8 next Fri
- Reading: through §8.5 for Mon