### **Bottom-Up Approach Inspiration**





# **Bottom-Up Approach Inspiration**

- Robert Lee Moore
  - Mathematician
    - Penn (1911-1920)
    - U. of Texas (1920-1969)
  - Developed "Moore Method" or "Texas Method"
    - A unique style of teaching mathematics
    - Start with axioms, have class build up the framework during the semester
    - "Looking ahead" forbidden





# **Bottom-Up Approach Inspiration**

- Robert Lee Moore
  - Mathematician
    - Penn (1911-1920)
    - U. of Texas (1920-1969)
  - Developed "Moore Method" or "Texas Method"
    - A unique style of teaching mathematics
    - Start with axioms, have class build up the framework during the semester
    - "Looking ahead" forbidden

#### • Yale Patt

- Computer Architect
- Renowned teacher and researcher
- "Bits and Bytes to C and Beyond" textbook





- We're going to follow a "bottom-up" approach
  - We're going to start with the lowest-level primitives
  - Build up high abstractions using these primitives

- We're going to follow a "bottom-up" approach
  - We're going to start with the lowest-level primitives
  - Build up high abstractions using these primitives
- Why?
  - No-magic understanding
  - Gain performance intuition

- We're going to follow a "bottom-up" approach
  - We're going to start with the lowest-level primitives
  - Build up high abstractions using these primitives
- Why?
  - No-magic understanding
  - Gain performance intuition
- Discarded alternative #1: "top-down"
  - Start with high-level parallelism libraries and languages
    - As "black boxes"
  - Then try to explain how they work
    - And their odd performance characteristics

- We're going to follow a "bottom-up" approach
  - We're going to start with the lowest-level primitives
  - Build up high abstractions using these primitives
- Why?
  - No-magic understanding
  - Gain performance intuition
- Discarded alternative #1: "top-down"
  - Start with high-level parallelism libraries and languages
    - As "black boxes"
  - Then try to explain how they work
    - And their odd performance characteristics
- Discard alternative #2: "top-up"

- Need to "forget" (for now) some things
  - Put aside any thoughts of higher-level constructs
  - Forget anything you might know of
    - **TBB**, **Cilk**, **OpenMP**, Blocks, Grand Central Dispatch, Ct, Java JSR-166, ZPL, NESL, Fortress, X10, Chapel, StreamIT, Brook, CUDA, **OpenCL**, **memory consistency**, etc.

- Need to "forget" (for now) some things
  - Put aside any thoughts of higher-level constructs
  - Forget anything you might know of
    - **TBB**, **Cilk**, **OpenMP**, Blocks, Grand Central Dispatch, Ct, Java JSR-166, ZPL, NESL, Fortress, X10, Chapel, StreamIT, Brook, CUDA, **OpenCL**, **memory consistency**, etc.
- Need to actively engage with material
  - Engage in class discussion, in-class exercises
  - Spend time outside of class on material
  - Embrace the homework assignments
    - More gritty, more effort at the start

- Need to "forget" (for now) some things
  - Put aside any thoughts of higher-level constructs
  - Forget anything you might know of
    - **TBB**, **Cilk**, **OpenMP**, Blocks, Grand Central Dispatch, Ct, Java JSR-166, ZPL, NESL, Fortress, X10, Chapel, StreamIT, Brook, CUDA, **OpenCL**, **memory consistency**, etc.
- Need to actively engage with material
  - Engage in class discussion, in-class exercises
  - Spend time outside of class on material
  - Embrace the homework assignments
    - More gritty, more effort at the start
- Less spoon-feeding, more self-learning

- Need to "forget" (for now) some things
  - Put aside any thoughts of higher-level constructs
  - Forget anything you might know of
    - **TBB**, **Cilk**, **OpenMP**, Blocks, Grand Central Dispatch, Ct, Java JSR-166, ZPL, NESL, Fortress, X10, Chapel, StreamIT, Brook, CUDA, **OpenCL**, **memory consistency**, etc.
- Need to actively engage with material
  - Engage in class discussion, in-class exercises
  - Spend time outside of class on material
  - Embrace the homework assignments
    - More gritty, more effort at the start
- Less spoon-feeding, more self-learning
- Best way I know how to teach this course

#### Simple Parallel Work Decomposition

- Sequential code
  - for (int i=**0**; i<**SIZE**; i++):
    - calculate(i, ..., ..., ...)
- Parallel code, for each thread:
  - void each\_thread(int thread\_id):
    - segment\_size = SIZE / number\_of\_threads
    - assert(SIZE % number\_of\_threads == 0)
    - my\_start = thread\_id \* segment\_size
    - my\_end = my\_start + segment\_size
    - for (int i=my\_start; i<my\_end; i++)</pre>
      - calculate(i, ..., ...)
- Hey, its a parallel program!

- Sequential code
  - for (int i=**0**; i<**SIZE**; i++):
    - calculate(i, ..., ..., ...)
- Parallel code, for each thread:
  - void each\_thread(int thread\_id):
    - segment\_size = SIZE / number\_of\_threads
    - assert(SIZE % number\_of\_threads == 0)
    - my\_start = thread\_id \* segment\_size
    - my\_end = my\_start + segment\_size
    - for (int i=my\_start; i<my\_end; i++)</pre>
      - calculate(i, ..., ..., ...)
- Hey, its a parallel program!

- Sequential code
  - for (int i=**0**; i<**SIZE**; i++):
    - calculate(i, ..., ...)
- Parallel code, for each thread:
  - void each\_thread(int thread\_id):
    - segment\_size = SIZE / number\_of\_threads
    - assert(SIZE % number\_of\_threads == 0)
    - my\_start = thread\_id \* segment\_size
    - my\_end = my\_start + segment\_size
    - for (int i=my\_start; i<my\_end; i++)</pre>
      - calculate(i, ..., ...)

- Sequential code
  - for (int i=**0**; i<**SIZE**; i++):
    - calculate(i, ..., ..., ...)
- Parallel code, for each thread:
  - void each\_thread(int thread\_id):
    - segment\_size = SIZE / number\_of\_threads
    - assert(SIZE % number\_of\_threads == 0)
    - my\_start = thread\_id \* segment\_size
    - my\_end = my\_start + segment\_size
    - for (int i=my\_start; i<my\_end; i++)</pre>
      - calculate(i, ..., ..., ...)
- Hey, its a parallel program!

- Sequential code
  - for (int i=**0**; i<**SIZE**; i++):
    - calculate(i, ..., ...)
- Parallel code, for each thread:
  - int counter = 0 // global variable
  - void each\_thread(int thread\_id):
    - while (true)
      - int i = **atomic\_add**(&counter, 1)
      - if (i >= SIZE)
        - return
      - calculate(i, ..., ..., ...)
- Dynamic load balancing, but high overhead

- Sequential code
  - for (int i=**0**; i<**SIZE**; i++):
    - calculate(i, ..., ...)
- Parallel code, for each thread:
  - int counter = 0 // global variable
  - void each\_thread(int thread\_id):
    - while (true)
      - int i = **atomic\_add**(&counter, 1)
      - if (i >= SIZE)
        - return
      - calculate(i, ..., ..., ...)
- Dynamic load balancing, but high overhead



- Sequential code
  - for (int i=**0**; i<**SIZE**; i++):
    - calculate(i, ..., ..., ...)
- Parallel code, for each thread:
  - int counter = 0 // global variable
  - void each\_thread(int thread\_id):
    - while (true)
      - int i = **atomic\_add**(&counter, 1)
      - if (i >= SIZE)
        - return
      - calculate(i, ..., ..., ...)



- Sequential code
  - for (int i=**0**; i<**SIZE**; i++):
    - calculate(i, ..., ...)
- Parallel code, for each thread:
  - int counter = 0 // global variable
  - void each\_thread(int thread\_id):
    - while (true)
      - int i = **atomic\_add**(&counter, 1)
      - if (i >= SIZE)
        - return
      - calculate(i, ..., ..., ...)
- Dynamic load balancing, but high overhead





- Parallel code, for each thread:
  - int num\_segments = SIZE / GRAIN\_SIZE
  - int counter = 0 // global variable
  - void each\_thread(int thread\_id):
    - while (true)
      - int i = **atomic\_add**(&counter, 1)
      - if (i >= num\_segments)
        - return
      - my\_start = i \* GRAIN\_SIZE
      - my\_end = my\_start + GRAIN\_SIZE
      - for (int j=my\_start; j<my\_end; j++)</pre>
        - calculate(j, ..., ..., ...)



- Parallel code, for each thread:
  - int num\_segments = SIZE / GRAIN\_SIZE
  - int counter = 0 // global variable
  - void each\_thread(int thread\_id):
    - while (true)
      - int i = **atomic\_add**(&counter, 1)
      - if (i >= num\_segments)
        - return
      - my\_start = i \* GRAIN\_SIZE
      - my\_end = my\_start + GRAIN\_SIZE
      - for (int j=my\_start; j<my\_end; j++)</pre>
        - calculate(j, ..., ..., ...)
- Dynamic load balancing with lower (adjustable) overhead