ECE 2574: Data Structures and Algorithms - Link-Based Performance

C. L. Wyatt

Today we will look more citically at the relative advantages of array-based versus link-based implementations.

- Finish link-based bag implementation
- A benchmark task
- How to do timing
- Array-based performance
- Link-based performance

A singly-linked version of the ADT Bag

see in-class code

A benchmark task

- A benchmark is a relatively simple problem that can be used to experimentally test some aspect of an implementation.
- Lets define a benchmark to examine the difference in performance between array and link-based implementations of a list.
- Generate N uniform random positive integers sequentially, inserting them into a list so as to maintain a sorted ordering.

A container generic version of the benchmark

See in-calss code

Adding measurements to the benchmark

- time used
- space used

The array-based variant of the benchmark

See example runs

The single link-based variant of the benchmark

See example runs

What conclusions can we draw?

Memory Architectures

How much longer does it take (roughly) to fetch a value from RAM versus a CPU cache?

```
Latency Comparison Numbers
--------------------------
L1 cache reference                           0.5 ns
L2 cache reference                           7   ns
Main memory reference                      100   ns
```

from https://gist.github.com/jboner/2841832 Latency Numbers Every Programmer Should Know

What might we do to mitigate this behavior?

You can write your own memory management layer.

- This is required when there is no OS (bare metal).
- It can also be effective when there is an OS because your application might know more about memory access patterns and can optimize for those.

This is beyond the scope of this course but is typically covered in 3574.

Next Actions and Reminders

- Read CH pp. 159-171
- Take warmup before Fri 9/22 at noon.