

Lecture 8: The Cache Oblivious Approach

www.cs.illinois.edu/~wgropp



Designing for Memory Hierarchy

- Lets start with a simple 3-level memory model:
 - ◆ Registers
 - ◆ Cache
 - ◆ Main Memory
- The cache is composed of lines of L words; we normally assume that there are at least L lines (often many more)
- If a load or store from the CPU cannot be satisfied from the cache, a *cache miss* occurs.
- When a miss occurs, one line is loaded from memory into the cache (this line contains (a copy of) the referenced memory location)
- If the cache was already full, one line is evicted (written back out to memory if necessary (e.g., a write occurred to some data in the line))



Some Further Assumptions

- The cache is perfect:
 - ◆ Each line of memory can be placed anywhere in the cache. This is called a *fully associative* cache for reasons we will discuss later.
 - ◆ When a line is evicted, the cache makes the perfect choice, choosing the line that will not be needed for the longest time *in the future*.
 - In practice, the policy used is called *Least Recently Used (LRU)*: the cache line that has not been used (accessed) for the longest period of time is selected for eviction.



A New Complexity Model

- In addition to counting computational operations, count cache misses.
 - ◆ This is one step better than counting loads and stores, as loads and stores assumes a 2-level memory hierarchy – registers and main memory.
 - ◆ There are many other parameters that can be important, but following our strategy, we can use this model as a kind of bound on performance.
 - ◆ Specifically, we do *not* weight a cache miss – we will simply count them. This means that this model cannot be used to compare algorithms that trade work for cache misses.
 - ◆ We also assume that cache misses are not overlapped with other operations (we'll see later that this can be a significant limitation)
- However, with all of these restrictions, this is still a better model than just counting operations



Cache Oblivious Algorithms

- Our new complexity model is $W(n) + Q(n,L,Z)$ where $W(n)$ is the work, $Q(n,L,Z)$ are the number of cache misses, n is the “problem size”, L is the line size, and Z is the cache size.
- The goal is to minimize (or at least reduce) this cost relative to the simple algorithms that only consider $W(n)$.
- Lets simplify further: Are there algorithms that are independent of L and Z that are good approximations to optimal solutions according to this complexity model?
 - ◆ Call these *cache parameter oblivious algorithms*, or *cache oblivious algorithms* for short
 - ◆ The obviously aren't cache oblivious, but they are portable to any system for which this three-level memory model is a good approximation.



Divide and Conquer

- Consider the matrix examples: matrix multiply and transpose. We know that one approach that improves cache utilization is to subdivide each loop into blocks. These blocks must be chosen to fit in the relevant level of cache (in a model that is optimizing for a particular machine architecture)
- What if instead we applied *divide and conquer* and stopped at some “good” size?





Divide and Conquer: Transpose




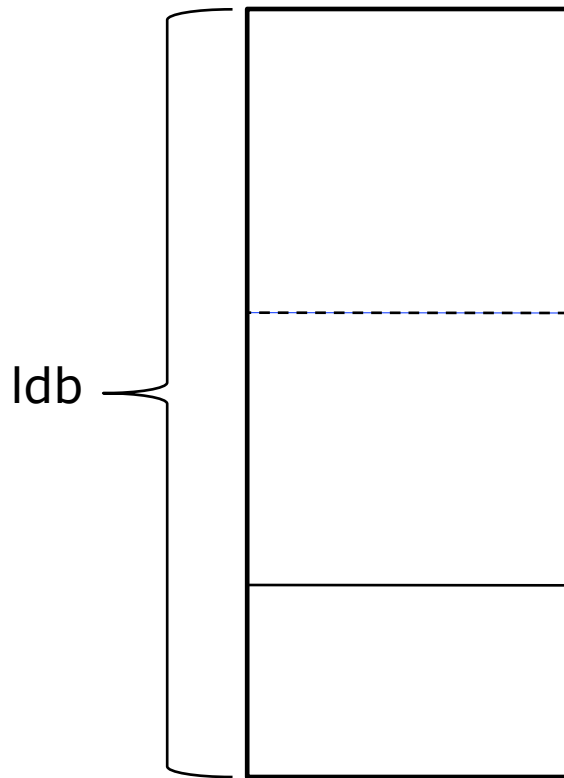
- By recursively subdividing (and keeping the submatrices reasonably square by always dividing the long dimension), we get to a block that will fit in cache, eliminating most (non-compulsory) cache misses.
- By picking a reasonably small size (say, 8k words) for the smallest block (where recursion is terminated), we should reduce the number of cache misses *without an explicit knowledge of the cache size* (other than “bigger than 16k”)



Details of CO Transpose



Details of CO Transpose



The CO Transpose Code

```
int transpose( double *a, int ndra, int nr, int nc, double *b, int ndr )
{
    if (nr < 32) {
        transposeBase( a, ndra, nr, nc, b, ndr );
    }
    else {
        /* subdivide the long side */
        if (nr > nc) {
            transpose( a, ndra, nr/2, nc, b, ndr );
            transpose( a + nr/2, ndra, nr-nr/2, nc, b+(nr/2)*ndr, ndr );
        }
        else {
            transpose( a, ndra, nr, nc/2, b, ndr );
            transpose( a+ndra*(nc/2), ndra, nr, nc-nc/2, b+nc/2, ndr );
        }
    }
}
```



Optimality

- The “cache oblivious” algorithm is asymptotically optimal (under the performance measurement of this model) within a factor of two
 - ◆ No data movement (involving the matrices) occurs until the minimum block size is reached. At this point, the “naïve” transpose algorithm may be used, because both input and output submatrices fit in cache. The number of misses is the number of compulsory misses, plus a small number if the blocks are not exact multiples of the line size
 - And they probably aren’t exact multiples – that’s part of the point of being cache oblivious

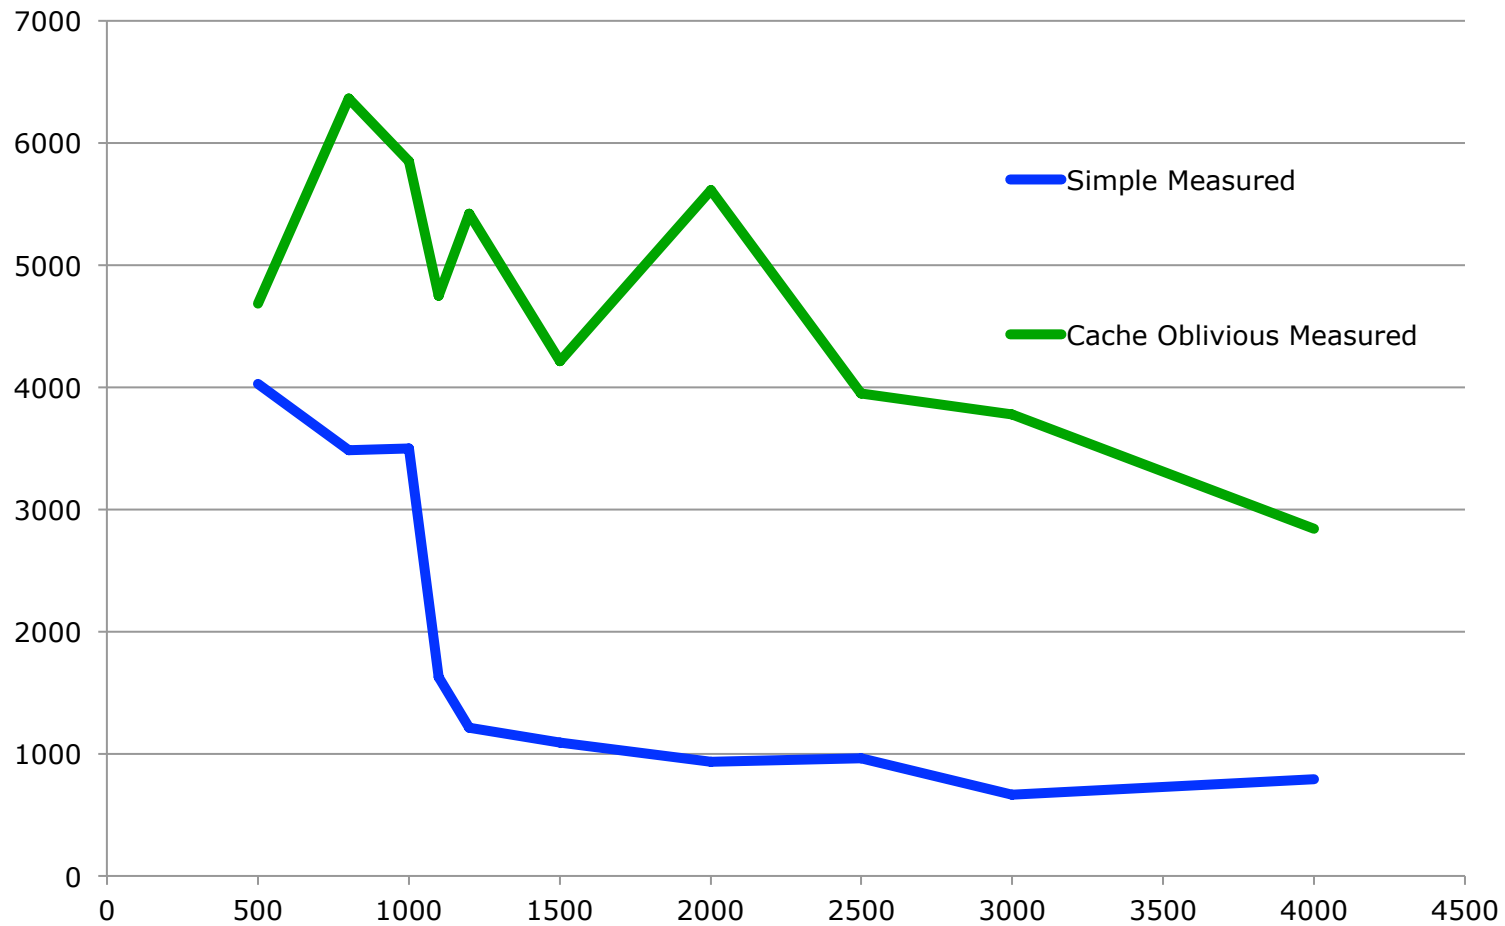


Some Results

Size	Naïve	Cache Oblivious
64	3000	2800
128	1900	3400
256	1900	4300
512	1200	1800
1024	320	1100
2048	180	1100
2049	780	1100
4096	150	980
4097	560	1100



More Recent Results



Observations

- Cache Oblivious approach provides more consistent performance with a relatively simple implementation
 - ◆ Less sensitive to layout than the naïve implementation
- However, may not preserve full L1-resident performance



Refinements and Issues

- Recursion, implemented as function calls, can be relatively expensive (a few 100's of instructions)
 - ◆ Manual implementation with a stack is relatively easy but tedious
- Cache utilization is approximate
- Performance model assumes that the cache is *fully associative* – any data item may be placed in any location. More on this later.



Some Comparisons

- From Cache-Oblivious Algorithms EXTENDED ABSTRACT Matteo Frigo Charles E. Leiserson Harald Prokop Sridhar Ramachandran (1999)
- Note cost: 30ns/2Flops, or 67 Mflops
 - ◆ Depending on platform, this is either good or bad, relative to the best available code
- This result shows that cache oblivious programs can be faster than the simple, naïve code, with only a small increase in complexity
- They do not compare with the best available implementations




Figure 5: Average time taken to multiply two $N \times N$ matrices, divided by N^3 .



Cache Oblivious Searching

- Consider a static set of keys that needs to be searched repeatedly, and assume that a binary tree is the preferred approach
- If the tree is built using memory allocated for each entry, linked with pointers, then each link will often cause a cache miss.
 - ◆ For suitably large set of keys, a cacheline is likely to be ejected before another key on that cacheline is accessed
 - ◆ Thus, $O(\log n)$ cache misses



A CO Data Structure

- A binary tree of height $\log_2 n$ fits in $n-1$ words.
- Organize the tree as a collection of subtrees, each of which is stored in consecutive locations
- $\log_2 L$ accesses are in a cache line of L keys, so the total number of cache misses is roughly $\log_2 n / \log_2 L$, where n is the number of keys.

