GPU Computing with CUDA
Lab 8 - N-body problem

Christopher Cooper
Boston University

August, 2011
UTFSM, Valparaíso, Chile

Objectives

- Implement an efficient N-body solver
- Compare timings of different implementations
- Optimizations:
  - Tiling
  - Loop unrolling

Efficient N-body solver

- N-body problem is basically a mat-vec multiplication
  - Matrix for kernel, vector for strength of each body
- We want to do this “matrix-free”
  - We will be using the 1/r kernel (look at the CPU kernel)
- Each thread will compute one row
  - If using shared memory, each thread will also load to it

Efficient N-body solver

- Steps
  - Implement in global memory
  - Use tiling with shared memory tiles with size of blocks
  - Unroll loops
  - What if number of elements is not a multiple of block size?
  - What is the optimal block size?

Efficient N-body solver

- My results (32768 bodies)
  - Tiled: 0.03137s ~370 FLOPS
  - Unrolled (32): 0.026s ~ 455 FLOPS
  - Bigger tiles: 0.025s ~ 470 FLOPS